QMATH summer school on the mathematics of entanglement via nonlocal games

Thomas Vidick

Lecture 1

Introduction to nonlocal games

1.1 Nonlocal games

A *nonlocal game* describes the actions of a *referee* interacting with two, or sometimes more, *cooperating players*.

Definition 1.1 (Nonlocal game)**.** A *nonlocal game* $\mathcal{G}$ is specified by two finite sets of *questions* $\mathcal{X}$ and $\mathcal{Y}$, two finite sets of *answers* $\mathcal{A}$ and $\mathcal{B}$, a distribution π on $\mathcal{X} \times \mathcal{Y}$ and a *game predicate* $V : \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B} \to \{0,1\}$. We write succinctly $\mathcal{G} = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, \pi, V)$ for this data.

The game is "played" as follows: The referee selects a pair of questions $(x,y) \in \mathcal{X} \times \mathcal{Y}$ according to π. They send x to the first player, usually referred to as "Alice," and y to the second player, "Bob." Alice responds with some $a \in \mathcal{A}$ and Bob with some $b \in \mathcal{B}$. The main restriction on how a, b are determined is that Alice and Bob may cooperate but should not communicate—we make this precise below. The players win the game if and only if $V(x,y,a,b) = 1$.

We generally consider that the game is repeated an arbitrarily large number of times and are interested in the largest fraction of games that Alice and Bob can win in the limit of infinitely many repetitions. Since the referee's actions are memoryless (each game instance is played independently), without loss of generality the optimal strategy is also memoryless and we only consider such strategies.

Let us formalize how a strategy for the players is represented. While in general a "strategy" may involve much thinking and "strategizing," for purposes of determining the player's success probability in the game all the relevant information is captured by a single list of numbers, namely for every pair of questions (x,y) and every pair of answers (a,b), what is the probability that the players answer (a,b) to (x,y). This is so important that we make it into a definition.

Definition 1.2 (Strategy)**.** Given a nonlocal game $\mathcal{G} = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, \pi, V)$ a *strategy* for the players in $\mathcal{G}$ is an element $p = (p(a,b|x,y))_{abxy} \in [0,1]^{\mathcal{A}\times\mathcal{B}\times\mathcal{X}\times\mathcal{Y}}$ such that for all $(x,y) \in \mathcal{X} \times \mathcal{Y}$, $\sum_{(a,b)\in\mathcal{A}\times\mathcal{B}} p(a,b|x,y) = 1$.

If *any* strategy of the form specified in the definition was allowed for the players, the study of nonlocal games would be rather boring.

Exercise 1.1. Given a game $\mathcal{G}$, say that a pair of questions $(x,y) \in \mathcal{X} \times \mathcal{Y}$ for $\mathcal{G}$ is *nontrivial* if $\pi(x,y) > 0$ and there is some $(a,b) \in \mathcal{A} \times \mathcal{B}$ such that $V(x,y,a,b) = 1$. For any game $\mathcal{G}$, give a simple expression for the maximum winning probability of players allowed to use any strategy satisfying the requirements of Definition 1.2.

Given a nonlocal game $\mathcal{G}$ and a strategy S for $\mathcal{G}$, the *success probability* of S in $\mathcal{G}$ is defined as

$$\omega(G;S) \;=\; \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b)\, p(a,b|x,y) \;.$$

Informally this quantity is the average, over the referee's choice of questions and the player's probabilistic strategy, that the players provide valid answers to the referee. (It coincides with the maximum fraction of games that can be won in the limit of infinite repetitions.) If one fixes a collection of possible strategies S then one can define an associated *value* $\omega(G;S)$ for the game, which is the supremum success probability achievable using strategies $S \in \mathcal{S}$:

$$\omega(G;\mathcal{S}) \;=\; \sup_{S\in\mathcal{S}} \omega(G;S) \;.$$

What makes the study of nonlocal games interesting is that different types of restrictions can be placed on the player's strategies. So far we only said that "the players should not communicate." So how do we model this intuitive requirement? We will see that there are multiple natural ways of doing this, leading to a rich theory of values (i.e. maximum success probabilities) for nonlocal games. We start by giving arguably the three most natural possibilities.

1.1.1 Classical and non-signaling strategies

The most natural class of strategies are classical strategies. Informally, a strategy is "classical" if it can be written as a convex combination of product strategies.

Definition 1.3 (Classical strategy)**.** A strategy p for a nonlocal game $\mathcal{G}$ is called *classical* if there exists a probability space (Ω,μ) and for all a,x and b,y measurable functions $p_A(a|x,\cdot), p_B(b|y,\cdot) : \Omega \to [0,1]$ such that for all x,y,ω,

$$\sum_a p_A(a|x,\omega) \;=\; \sum_b p_B(b|y,\omega) = 1 \;,$$

and

$$p(a,b|x,y) \;=\; \int_\omega p_A(a|x,\omega) p_B(b|y,\omega) d\mu(\omega) \;.$$

The classical value of a game is denoted by $\omega(G) \;=\; \omega(\mathcal{G};\mathcal{S}_{class})$, where $\mathcal{S}_{class}$ denotes the set of all classical strategies.

Classical strategies are the most restrictive class of strategies we will consider. At the opposite end of the spectrum we have the most permissive class of strategies, non-signaling strategies, which informally correspond to those strategies which satisfy no other restriction but that of respecting causality between the players' answers and their questions (i.e. one player's answer cannot depend on the other player's question).

Definition 1.4 (Non-signaling strategy)**.** A strategy p for a nonlocal game $\mathcal{G}$ is called *non-signaling* if for every a,b,x,x',y,y' it holds that

$$\sum_a p(a,b|x,y) = \sum_a p(a,b|x',y) \quad \text{and} \quad \sum_b p(a,b|x,y) = \sum_b p(a,b|x,y') \;.$$

Exercise 1.2.

1. Show that any non-signaling strategy that is also deterministic is classical.
2. For $\mathcal{X} = \mathcal{Y} = \mathcal{A} = \mathcal{B} = \{0,1\}$, give an example of a strategy that is non-signaling but is not classical.

If you gave some thought to the previous exercise, you might have discovered that it is not so easy! Coming up with a non-signaling strategy is not hard, but how do we argue that it is *not* classical? An observation that may help you is that the sets of classical strategies, and the set of non-signaling strategies, are both closed convex sets. To show that a point p in the non-signaling set is not classical it suffices to find a linear function $\lambda : [0,1]^{\mathcal{X}\times\mathcal{Y}\times\mathcal{A}\times\mathcal{B}} \to \mathbb{R}$ such that $\lambda \cdot p > \sup_{q\in\mathcal{S}_{class}} \lambda \cdot q$. Any such linear function can be represented by its coefficients $\lambda = (\lambda_{xyab})$, which we may without loss of generality assume normalized such that $\|\lambda\|_1 = 1$. Now let $\pi(x,y) = \sum_{a,b} |\lambda_{xyab}|$ and $V(x,y,a,b) = \text{sign}(\lambda_{xyab})$. Then $\mathcal{G}_\lambda = (\mathcal{X},\mathcal{Y},\mathcal{A},\mathcal{B},\pi,V)$ is a nonlocal game and moreover for any strategy p, its success probability in $\mathcal{G}_\lambda$ is exactly

$$\sum_{abxy} \pi(x,y)V(x,y,a,b)p(a,b|x,y) \;=\; \lambda \cdot p \;.$$

This discussion leads us to an important insight: games are in a certain sense dual to strategies, and in particular games provide a means to separate different sets of strategies. To show that there exists a non-signaling but not classical strategy it suffices to design a game $\mathcal{G}$ such that there is a non-signaling strategy in $\mathcal{G}$ that has a higher success probability than any classical strategy.

Example 1.1 (CHSH game)**.** In the CHSH game $\mathcal{G}_{\text{CHSH}}$ questions and answers are a single bit each: $\mathcal{X} = \mathcal{Y} = \mathcal{A} = \mathcal{B} = \{0,1\}$. The distribution π is uniform on $\mathcal{X}\times\mathcal{Y}$, and the game predicate is $V(x,y,a,b) = 1$ if $a \oplus b = x \wedge y$ and $V(x,y,a,b) = 0$ otherwise.

Exercise 1.3.

1. Show that there is a classical strategy which succeeds in the game $\mathcal{G}_{\text{CHSH}}$ with probability $3/4$.
2. Show that there is a non-signaling strategy which succeeds in $\mathcal{G}_{\text{CHSH}}$ with probability 1.
3. Show that $3/4$ is best achievable for classical strategies. *[Hint: first consider classical deterministic strategies. Such a strategy is represented by* 4 *bits only.]*

So far we have introduced classical strategies, non-signaling strategies, and showed that (as soon as the number of questions and answers per player is at least 2) the latter can lead to strictly higher success probabilities. Next we turn to quantum strategies.

1.1.2 Quantum strategies

Informally, a quantum strategy in a nonlocal game is "any strategy that can be implemented locally using the laws of quantum mechanics." Unfortunately (or fortunately?) there is not a single possible mathematical interpretation of this sentence, even within quantum mechanics. We will return to this difficulty later on. For now, let us stick to the standard interpretation in quantum computing, which is to represent locality using a tensor product of Hilbert spaces. More precisely, in quantum mechanics we associate a Hilbert space with each player, $\mathcal{H}_A$ for Alice and $\mathcal{H}_B$ for Bob, such that the joint Hilbert space is $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$. The joint state of Alice and Bob is a unit vector (also called *state*) $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$. When Alice (resp. Bob) receives a question x (resp. y) she (he) may perform an arbitrary measurement, or POVM (for positive operator-valued measure), on her (his) system to determine her (his) answer.

Definition 1.5. Given a separable Hilbert space $\mathcal{H}$, a POVM on $\mathcal{H}$ is a collection $\{A_i\}_{i\in I}$ where I is an arbitrary index set such that for all $i \in I$, A_i is positive semidefinite on $\mathcal{H}$ and $\sum_i A_i = \text{Id}$. A POVM is called *projective* if all POVM elements A_i are projections.

For each $x \in \mathcal{X}$ let $\{A_{x,a}\}_{a\in\mathcal{A}}$ be a POVM on $\mathcal{H}_A$ associated with Alice's question x, and similarly for each $y \in \mathcal{Y}$ let $\{B_{y,b}\}_{b\in\mathcal{B}}$ be a POVM on $\mathcal{H}_B$. The Born measurement rule is that together, the state $|\psi\rangle$ and the POVM lead to the strategy

$$p(a,b|x,y) \;=\; \langle\psi|A_{x,a}\otimes B_{y,b}|\psi\rangle \;. \tag{1.1}$$

It is a good sanity-check to verify that this indeed defines a valid strategy according to Definition 1.2. For future reference we record it in a definition.

Definition 1.6 (Quantum (tensor) strategy)**.** A strategy p is a *quantum (tensor) strategy* if there exists (separable) Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, a unit vector $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$, and for every x, y POVM $\{A_{x,a}\}_a$ on $\mathcal{H}_A$ and $\{B_{y,b}\}_b$ on $\mathcal{H}_B$ such that (1.1) holds. We let $\mathcal{S}_{quant}$ denote the set of all quantum tensor strategies.

The quantum value of a game is denoted by $\omega^*(G) = \omega(\mathcal{G};\mathcal{S}_{quant})$.

While in general a strategy may use any POVM, it is sometimes convenient to restrict attention to strategies in which all measurements are projective. Naimark's dilation theorem makes this possible. We state here the theorem in a form adapted to our setting.

Theorem 1.7 (Naimark)**.** *Let $\{A_i\}$ be a POVM on a Hilbert space $\mathcal{H}$ and $|\psi\rangle \in \mathcal{H}$. Then there exists a space $\mathcal{H}'$, a state $|\psi'\rangle \in \mathcal{H}'$, and a projective measurement $\{A'_i\}$ on $\mathcal{H}\otimes\mathcal{H}'$ such that for any i,*

$$\sqrt{A_i}|\psi\rangle\langle\psi|\sqrt{A_i} = \mathrm{Tr}_{\mathcal{H}'}\left(A'_i|\psi\rangle\langle\psi|\otimes|\psi'\rangle\langle\psi'|A'_i\right) \;.$$

The theorem guarantees that post-measurement states, and hence outcome probabilities, are preserved.

Exercise 1.4. Prove Naimark's theorem. State and prove a version of the theorem that simultaneously "dilates" multiple POVM $\{A_{xa}\}_{a\in\mathcal{A}}$ acting on the same Hilbert space $\mathcal{H}$. Verify that the theorem can be applied in the bipartite setting to construct an equivalent projective strategy from any quantum tensor strategy.

At first it may not be so easy to see how quantum strategies compare to classical and non-signaling strategies. The following exercise suggests a few simple observations.

Exercise 1.5.

1. Verify that any quantum strategy is non-signaling.
2. Show that any classical strategy is a quantum strategy. *[Hint: first, show this for deterministic strategies. Don't forget to show it for randomized strategies!]*
3. A quantum state $|\psi\rangle \in \mathcal{H}_A\otimes\mathcal{H}_B$ is said *entangled* if it cannot be written as $|\psi\rangle = |\psi_A\rangle\otimes|\psi_B\rangle$ for some $|\psi_A\rangle \in \mathcal{H}_A$, $|\psi_B\rangle \in \mathcal{H}_B$. Show that any quantum strategy using a state $|\psi\rangle$ that is *not* entangled is classical. Give an example of a quantum strategy using an entangled state $|\psi\rangle$ that is nevertheless classical.[1]

The exercise shows that quantum strategies lie "between" classical and non-signaling strategies. Do they coincide with either? We will soon show that the answer is no, and in fact the CHSH game from Example 1.1 can be used to separate all three sets. Before we do this we investigate a class of games called XOR games which contains the CHSH game. These games have the advantage that quantum strategies for them can always be written in a particularly simple form. Studying this will allow us to develop insights in the strength and limitations of quantum strategies.

[1] It is worth noting fact that quantum algorithms may be more powerful than classical algorithms is irrelevant here, because each prover on their own is always allowed arbitrary large computational power.

1.2 XOR games

XOR games are the simplest class of games that has been extensively studied. Informally, a game $\mathcal{G}$ is an XOR game if $\mathcal{A} = \mathcal{B} = \{-1, 1\}$ and the game predicate $V(x, y, a, b)$ depends only on x, y and the product ab.[2] Setting some notation, in an XOR game

1. The referee selects a pair of questions $(i, j) \in \{1, \ldots, m\} \times \{1, \ldots, n\}$ according to a distribution π.
2. The referee sends i to Alice and j to Bob. Alice and Bob reply with signs $a_i, b_j \in \{\pm 1\}$ respectively.
3. The game predicate takes the form $V(i, j, a, b) = \frac{1}{2}(1 + a_i b_j c_{ij})$, for some $c_{ij} \in \{\pm 1\}$. Note that this always lies in $\{0, 1\}$ and only depends on the product $a_i b_j$, as desired.

Given an arbitrary XOR game $\mathcal{G}$, we can represent it succinctly using the matrix $G \in \mathbb{R}^{m \times n}$ with coefficients $G_{ij} = \pi(i, j)c_{ij}$. This matrix satisfies $\|G\|_1 = 1$. Conversely, from any $G \in \mathbb{R}^{m \times n}$ with $\|G\|_1 = 1$ we can construct an XOR game $\mathcal{G}$ by setting $\pi(i, j) = |G_{ij}|$ and $c_{ij} = \text{sign}(G_{ij})$.

Observe that the CHSH game from Example 1.1 is an XOR game with $m = n = 2$, π uniform, and $c_{00} = c_{01} = c_{10} = +1$ and $c_{11} = -1$. The associated matrix is $G_{\text{CHSH}} = \frac{1}{4}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$.

1.2.1 Bias of an XOR game

Suppose that a strategy p for XOR game $\mathcal{G}$ succeeds with probability $\omega = \frac{1}{2} + \frac{1}{2}\beta$, where $\beta \in [-1, 1]$. Because of the special form we assumed for $V(\cdot)$, by flipping all of Bob's answers in the strategy, we obtain a new strategy p' that succeeds with probability $\omega' = \frac{1}{2} - \beta$. This means that the maximum success probability in $\mathcal{G}$, where the maximum is taken over any set of strategies that is closed under flipping all of Bob's answers, is always of the form $\frac{1}{2} + \frac{1}{2}\beta$ for some $0 \leq \beta \leq 1$. We call β the *bias* of the game.

Definition 1.8. Given an XOR game $\mathcal{G}$, define its *classical bias* $\beta(\mathcal{G})$ as twice the maximum success probability of classical strategies in the game, minus 1. Similarly, define the *quantum bias* $\beta^*(G)$ and *non-signaling bias* $\beta^{ns}(G)$ from quantum (tensor) and non-signaling strategies respectively.

The classical, quantum and non-signaling biases admit nice expressions which we now give.

Lemma 1.9. *For any XOR game $\mathcal{G}$,*

$$\beta(\mathcal{G}) = \max_{a_i, b_j \in \{\pm 1\}} \sum_{i,j} G_{ij} a_i b_j \,. \tag{1.2}$$

Proof. By definition of the set of classical strategies, extreme points are deterministic strategies. A classical deterministic strategy is specified by a collection of signs $a_i, b_j \in \{\pm 1\}$ representing Alice and Bob's answers to questions i, j respectively. The probability that (i, j) is selected as questions and Alice and Bob answer correctly is exactly $\frac{1}{2}\pi(i, j)(1 + c_{ij} a_i b_j)$. Summing over all (i, j) and using the definition of G_{ij} gives the result. □

Example 1.2. The bias of the CHSH game can be evaluated exactly:

$$\beta(\mathcal{G}_{\text{CHSH}}) = \max_{a_i, b_j \in \{\pm 1\}} \sum_{i,j} \pi(i, j) c_{ij} a_i b_j = \max \frac{1}{4}(a_1 b_1 + a_1 b_2 + a_2 b_1 - a_2 b_2) = \frac{1}{2} \,.$$

[2] If we represent the outcomes as elements of $\{0, 1\}$, then the product becomes an XOR $a \oplus b$, justifying the name "XOR game." We find the $\{-1, 1\}$ notation more convenient.

So the maximum success probability of classical strategies in the game is

$$\omega(G) \;=\; \frac{1}{2} + \frac{1}{2}\,\beta(G) \;=\; \frac{3}{4}\,.$$

Now we turn to the non-signaling bias.

Exercise 1.6. Say that a game is *nontrivial* if all pairs of questions with $\pi(x,y) > 0$ are nontrivial (see Exercise 1.1). For a nontrivial XOR game $\mathcal{G}$, show that the non-signaling bias satisfies $\beta^{ns}(\mathcal{G}) = 1$.

1.2.2 The quantum bias

Finally, let's consider quantum tensor strategies. Using that answers in an XOR game are always binary we can represent each player's strategy as a family of *observables*.

Definition 1.10. A binary observable is a Hermitian matrix $X \in \mathbb{C}^{d\times d}$ such that $X = X^\dagger$ and $X^2 = \mathbb{I}$ (in other words, all eigenvalues of X are in $\{\pm 1\}$). A binary observable X can always be decomposed as $X = X^0 - X^1$ for two projections X^0, X^1 such that $X^0 + X^1 = \mathrm{Id}$, i.e. $\{X^0, X^1\}$ is a projective measurement.

Let $(|\psi\rangle, \{A_{i,a}\}, \{B_{j,b}\})$ be a quantum strategy for an XOR game G. Fixing a pair of questions (i,j), the product of the outcomes (a,b) returned by this strategy to (i,j) has expectation

$$\begin{aligned}
\mathrm{E}[a\cdot b] &= \Pr\big((a,b)=(0,0)\big) + \Pr\big((a,b)=(1,1)\big) - \Pr\big((a,b)=(0,1)\big) - \Pr\big((a,b)=(1,0)\big)\\
&= \langle\psi|A_{i,0}\otimes B_{j,0}|\psi\rangle + \langle\psi|A_{i,1}\otimes B_{j,1}|\psi\rangle - \langle\psi|A_{i,0}\otimes B_{j,1}|\psi\rangle - \langle\psi|A_{i,1}\otimes B_{j,0}|\psi\rangle\\
&= \langle\psi|A_i\otimes A_j|\psi\rangle \;\in\; [-1,1]\,,
\end{aligned}$$

where $A_i = A_{i,0} - A_{i,1}$ and $B_j = B_{j,0} - B_{j,1}$. Note that A_i, B_j are Hermitian and satisfy $\|A_i\|, |B_j\| \le 1$. Moreover any Hermitian X can be written as $X = X_0 - X_1$ where $X_0 + X_1 = \mathrm{Id}$ by letting $X_0 = \frac{1}{2}(\mathrm{Id} + X)$ and $X_1 = \frac{1}{2}(\mathrm{Id} - X)$. Therefore, the quantum bias $\beta^*(\mathcal{G})$ is

$$\beta^*(\mathcal{G}) \;=\; \sup_{|\psi\rangle,\{A_i\},\{b_j\}} \sum_{i,j} G_{ij}\,\langle\psi|A_i\otimes B_j|\psi\rangle\,, \tag{1.3}$$

where the supremum is taken over all states $|\psi\rangle \in \mathcal{H}_A\otimes\mathcal{H}_B$ and Hermitian A_i, B_j such that $\|A_i\|, \|B_j\| \le 1$. Observe that (1.3) is linear in any given eigenvalue of any given A_i or B_j, all other parameters (including the eigenbasis) being fixed. This means that the supremum is achieved at A_i, B_j that are observables. We just proved the following lemma.

Lemma 1.11. *For any XOR game $\mathcal{G}$,*

$$\beta^*(\mathcal{G}) \;=\; \sup_{|\psi\rangle,\{A_x\},\{B_y\}} \sum_{i,j} G_{ij}\,\langle\psi|A_x\otimes B_y|\psi\rangle\,, \tag{1.4}$$

where the supremum is taken over all spaces $\mathcal{H}_A, \mathcal{H}_B$, states $|\psi\rangle \in \mathcal{H}_A\otimes\mathcal{H}_B$ and families of observables $\{A_x\}$, $\{B_y\}$.

A priori the supremum in (1.4) is needed, because there is no limit to the dimension of the spaces $\mathcal{H}_A$, $\mathcal{H}_B$; in general these may even be infinite-dimensional. You can verify that if we restrict the dimension of both spaces to 1, then the only states are $|\psi\rangle = (e^{i\theta})$ for some real angle θ, the only observables are

$X, Y \in \{\pm 1\}$, and the quantum bias reduces to the classical bias. If we restrict the state to the *maximally entangled state* in dimension $d \geq 1$

$$|\phi_d\rangle \;=\; \frac{1}{\sqrt{d}} \sum_i \; |i\rangle|i\rangle \;,$$

then we obtain a particularly simple formula, since in this case

$$\langle \phi_d | X \otimes Y | \phi_d \rangle \;=\; \frac{1}{d} \mathrm{Tr}(XY^T) \;. \tag{1.5}$$

We will slightly abuse notation and write $\langle X, Y \rangle = \mathrm{Tr}(XY^\dagger)$ to denote the matrix trace inner product.

Example 1.3. Consider

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad B_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad A_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad B_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix} .$$

Then you can verify that A_0, A_1, B_0, B_1 are observables. Moreover,

$$\frac{1}{2}\langle A_0, B_0 \rangle = \frac{1}{2}\langle A_0, B_1 \rangle = \frac{1}{2}\langle A_1, B_0 \rangle = \frac{\sqrt{2}}{2} \;, \quad \text{and} \quad \frac{1}{2}\langle A_1, B_1 \rangle = -\frac{\sqrt{2}}{2} \;.$$

Plugging these calculations into the definition of the CHSH game (Example 1.1), we see that these observables, together with a maximally entangled state in dimension $d = 2$, achieve a bias of

$$\frac{1}{4} \cdot 4 \cdot \frac{\sqrt{2}}{2} = \frac{\sqrt{2}}{2} \approx 0.73 \;,$$

which is strictly larger than the bias $1/2$ that you proved optimal for classical players in Exercise 1.3. In particular, quantum strategies are strictly more powerful than classical strategies.

The example of the CHSH game shows that quantum players can sometimes strictly outperform their classical peers. This already has a pretty neat (and, arguably, deep) consequence: it is possible to use the CHSH game as a "statistical test for information-theoretic randomness"! Indeed, any strategy that succeeds in the test with probability larger than $\frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$ (a simple condition to verify) cannot be a classical strategy, and in particular it cannot be a deterministic strategy. Thus, any pair of isolated devices (representing the players) that generate an input-output behavior that leads to a sufficiently high success probability in the game is necessarily randomness-generating. Note that this kind of randomness is very different from "pseudo-randomness": there is no question of computational power here, and the guarantees provided by the test are information-theoretic!

1.3 An SDP formulation of the quantum bias and Tsirelson's bound

Is there any limit to how well quantum players can do in the CHSH game, or more generally in an XOR game? Recall that the optimal bias for quantum players is given by (1.4):

$$\beta^*(\mathcal{G}) = \sup_{\substack{A_i, B_j \in \mathbb{C}^{d \times d} \\ A_i^2 = B_j^2 = \mathbb{I} \\ A_i = A_i^\dagger \\ B_j = B_j^\dagger}} \sum_{i,j} G_{ij} \cdot \langle \psi | A_i \otimes B_j | \psi \rangle \leq \sup_{\substack{\vec{u}_i, \vec{v}_j \in \mathbb{C}^{d^2} \\ \|\vec{u}_i\| = \|\vec{v}_j\| = 1}} \sum_{i,j} G_{ij} \vec{u}_i \cdot \vec{v}_j = \mathrm{SDP}(\mathcal{G}) \;. \tag{1.6}$$

Here the inequality holds because we can set $\vec{u}_i = A_i \otimes \mathrm{Id}\,|\psi\rangle$ and $\vec{v}_j = \mathrm{Id} \otimes B_j|\psi\rangle$. Under such choice, one can verify that

$$\|\vec{u}_i\| = \|\vec{v}_i\| = 1, \qquad u_i \cdot v_j = \langle\psi|E_i \otimes F_j|\psi\rangle\ .$$

The expression on the right-hand side of (1.6) is nice because it is directly analogous to the expression (1.2) for the classical bias: the only difference is that we now optimize over inner products of unit vectors in any dimension, instead of just products of ± 1 values. Although we will not show explicitly why, those of you familiar with semidefinite programs will easily recognize that $\mathrm{SDP}(\mathcal{G})$ is, indeed, an SDP.[3] In particular, the quantity can be approximated to within $\pm\varepsilon$ in time polynomial in n, m, and $\log(1/\varepsilon)$.

Note that (1.6) is only an upper bound. How good is it? Let's first use it to prove *Tsirelson's theorem*, which states that the lower bound of $\frac{\sqrt{2}}{2}$ on the quantum bias of the CHSH game obtained in Exercise 1.3 is tight.

Theorem 1.12 (Tsirelson)**.** *For $\mathcal{G}_{CHSH}$ the CHSH game, it holds that $\beta^*(\mathcal{G}_{CHSH}) \leq \frac{\sqrt{2}}{2}$.*

Proof. For the CHSH game we can write

$$\begin{aligned}
\mathrm{SDP}(\mathcal{G}_{CHSH}) &= \sup_{\substack{\vec{u}_i,\vec{v}_j \in \mathbb{C}^{d^2} \\ \|\vec{u}_i\|=\|\vec{v}_j\|=1}} \frac{1}{4}\big(\vec{u}_0 \cdot \vec{v}_0 + \vec{u}_1 \cdot \vec{v}_0 + \vec{u}_0 \cdot \vec{v}_1 - \vec{u}_1 \cdot \vec{v}_1\big) \\
&= \sup_{\substack{\vec{u}_i,\vec{v}_j \in \mathbb{C}^{d^2} \\ \|\vec{u}_i\|=\|\vec{v}_j\|=1}} \frac{1}{4}\big(\vec{u}_0 \cdot (\vec{v}_0 + \vec{v}_1) + \vec{u}_1 \cdot (\vec{v}_0 - \vec{v}_1)\big) \\
&= \sup_{\substack{\vec{v}_j \in \mathbb{C}^{d^2} \\ \|\vec{v}_j\|=1}} \frac{1}{4}\big(\|\vec{v}_0 + \vec{v}_1\| + \|\vec{v}_0 - \vec{v}_1\|\big) \\
&\leq \sup_{\substack{\vec{v}_j \in \mathbb{C}^{d^2} \\ \|\vec{v}_j\|=1}} \frac{\sqrt{2}}{4}\big(\|\vec{v}_0 + \vec{v}_1\|^2 + \|\vec{v}_0 - \vec{v}_1\|^2\big)^{1/2} \\
&= 2\frac{\sqrt{2}}{4}\ .
\end{aligned}$$

Here for the third line we used that for any nonzero $\vec{v}$ the supremum over unit $\vec{u}$ of $\vec{u} \cdot \vec{v}$ is $\|\vec{v}\|$, achieved at $\vec{u} = \vec{v}/\|\vec{v}\|$; the fourth line is the Cauchy-Schwarz inequality; the last expands the squares and uses that $\vec{v}_0$ and $\vec{v}_1$ are unit vectors. □

Tsirelson's proof of his theorem was a bit different: he worked directly with the operators that define a quantum tensor strategy, and considered the square of the game value. We will revisit his proof a little later. Theorem 1.12 shows that the bound (1.6) is tight for the CHSH game. What is amazing is that it is *always* tight, for any XOR game! This is can be shown by using a beautiful trick of Tsirelson's.

Exercise 1.7. Show that given a vector solution to $\mathrm{SDP}(\mathcal{G})$ it is always possible to find a quantum strategy that achieves exactly the same value. *[Hint: Consider Hermitian matrices $C_1, \ldots, C_d \in \mathbb{C}^{D\times D}$ that square to identity and pairwise anti-commute. For any vector u, consider $u \mapsto C(u) = \sum_i u_i C_i$. What can you say about $C(u)$? And about $\langle\phi_D|C(u) \otimes C(v)|\phi_D\rangle$?]*

[3] Semidefinite programs have real, not complex, variables. It is a simple exercise to rewrite (1.6) so that the supremum on the right-hand side is only taken over real vectors. *[Hint: stack up the real and imaginary parts!]*

The fact that (1.6) is an equality has amazing consequences for the study of XOR games. In particular, it has the following implications.

- The right-hand size of (1.6) can be expressed as a semidefinite program, hence the maximum expected payoff of quantum players can be computed efficiently. As we will see later this fact markedly distinguishes XOR games from more general nonlocal games.
- The proof of Tsirelson's theorem in Exercise 1.7 is explicit, hence an optimal quantum strategy can always be found efficiently. Moreover, there is always an optimal strategy in dimension $2^{\lfloor \min(n,m)/2 \rfloor}$.
- The following exercise shows that, in XOR games, quantum players can only ever achieve a payoff that is a constant factor larger than the optimal classical payoff.

Exercise 1.8. Grothendieck's inequality states that there exists a universal constant $K_G^{\mathbb{R}} \in \mathbb{R}$ such that for any integer n and any $M = (M_{ij}) \in \mathbb{R}^{n \times n}$,

$$\sup_{\substack{d,\, \vec{u}_i, \vec{v}_j \in \mathbb{C}^d \\ \|\vec{u}_i\|, \|\vec{v}_j\| \leq 1}} \left| \sum_{i,j} M_{ij}\, \vec{u}_i \cdot \vec{v}_j \right| \leq K_G \max_{x_i, y_j \in [-1,1]} \left| \sum_{i,j} Mij\, x_i y_j \right|.$$

The constant $K_G^{\mathbb{R}}$ is known to satisfy $K_G^{\mathbb{R}} \leq 1.782\ldots$. Furthermore, if $M = (M_{ij}) \in \mathbb{C}^{n \times n}$ and supremum on the right-hand side is taken over all complex $x_i, y_j \in \mathbb{C}$ such that $|x_i|, |y_j| \leq 1$ then the inequality holds with an improved constant $K_G^{\mathbb{C}} < K_G^{\mathbb{R}}$ such that $K_G^{\mathbb{C}} \leq 1.405\ldots$.

- What is the best constant K such that $\beta^*(\mathcal{G}) \leq K\beta(\mathcal{G})$, for any XOR game $\mathcal{G}$?

Remark 1.13. There are many interesting games that are not XOR games. A good example which we study below is the *Magic Square game*. This game is a "pseudo-telepathy" game, which means that the quantum value is 1 (there is a perfect quantum strategy), while the classical value is strictly below 1.

Exercise 1.9. Suppose that $\mathcal{G}$ is an XOR game such that $\beta^*(\mathcal{G}) = 1$. Show that $\beta(\mathcal{G}) = 1$.

1.4 Binary Linear System Games

We introduce a second family of nonlocal games that has been extensively studied.

A Binary Linear System (BLS) is specified by a collection of *variables* $v_1, \ldots, v_n$ ranging in $\{\pm 1\}$ and a collection of m *equations* $(E_1, c_1), \ldots, (E_m, c_m)$ such that for each $j \in \{1, \ldots, m\}$, $E_j \subseteq \{1, \ldots, n\}$ and $c_j \in \{\pm 1\}$. Informally, this is interpreted as the constraint $\prod_{\ell \in E_j} v_\ell = c_j$. In particular, we say that the BLS is *classically satisfiable* if there exists a ± 1 assignment to the v_j such that all constraints are satisfied.

Given a BLS $\{(E_j, c_j)\}$, we associate to it a game which is played as follows. We give an "operational" description, from which the question distribution π and game predicate V can easily be recovered. Questions to Alice are indexed by equations, $j \in \{1, \ldots, m\}$, and questions to Bob are indexed by variables, $i \in \{1, \ldots, n\}$. The referee selects a j uniformly at random, and an $i \in E_j$ uniformly at random. She sends j to Alice and i to Bob. Upon receiving question j, Alice is expected to return an assignment to all variables appearing in E_j; so, $\mathcal{A} = \{-1, 1\}^k$ where k is the maximum number of variables appearing in any given constraint (which we assume to be a constant for simplicity). Upon receiving question i, Bob is expected to return an assignment to variable v_i, so $\mathcal{B} = \{-1, 1\}$. The game predicate $V(i, j, a, b) = 1$ if and only if the assignment a satisfies E_j, i.e. $\prod_\ell a_\ell = c_j$ ("equation check"), and moreover it is consistent with b, i.e. the value that a assigns to variable i matches b ("consistency check").

Let's see an example.

Example 1.4 (Magic Square Game)**.** The magic square game is based on a BLS with $n = 9$ and $m = 6$. It is most easily visualized by arranging the 9 variables in a square,

$$\begin{array}{ccc} v_1 & v_2 & v_3 \\ v_4 & v_5 & v_6 \\ v_7 & v_8 & v_9 \end{array}.$$

The 6 constraints are that the product of all variables in a row or column should equal $+1$, *except* for the last column where the product should equal -1.

Exercise 1.10. Determine the classical value of the Magic Square game. *[Hint: there are* 6×3 *questions in total. What is the best classical solution to the BLS that you can write down?]*

BLS games are interesting because there is a close connection between *perfect* quantum strategies for them and *operator solutions* to the underlying system of equations, which extends the obvious connection between classical strategies and classical solutions. This connection in some cases allows us to determine if there is a perfect quantum strategy in the game, and in other cases allows us to show that deciding whether this is the case can be very hard.

Definition 1.14. Given a BLS (E, c), an *operator solution* to it is a collection $A_1, \ldots, A_n$ of binary observables such that for every $j \in \{1, \ldots, m\}$, the collection $\{A_\ell\}_{\ell \in E_j}$ pairwise commute and $\prod_{\ell \in E_j} A_\ell = c_j \,\mathrm{Id}$.

Note that the definition only requires that any two A_ℓ, $A_{\ell'}$ such that both ℓ, ℓ' appear in the same constraint C_j must commute. This does not imply that all the A_ℓ mutually commute because commutation is not a transitive relation.

Exercise 1.11. Verify that the following is an operator solution to the Magic Square game. Here σ_X, σ_Z and $\sigma_Y = i\sigma_X\sigma_Z$ are the Pauli matrices.

$$\begin{array}{ccc} I \otimes \sigma_Z & \sigma_Z \otimes I & \sigma_Z \otimes \sigma_Z \\ \sigma_X \otimes I & I \otimes \sigma_X & \sigma_X \otimes \sigma_X \\ \sigma_X \otimes \sigma_Z & \sigma_Z \otimes \sigma_X & \sigma_Y \otimes \sigma_Y \end{array} \tag{1.7}$$

It is not too hard to show that for any BLS, an operator solution immediately translates into a perfect quantum strategy for it. Thus the preceding exercise shows that the Magic Square has a perfect quantum strategy; combined with Exercise 1.10 it follows that the Magic Square is a pseudo-telepathy game.

Lemma 1.15. *Suppose given an operator solution* $Y_1, \ldots, Y_n$ *to a BLS* (E, c) *such that each* Y_j *is a binary observable on a finite-dimensional Hilbert space* $\mathcal{H}$*. Then the following strategy succeeds with probability* 1 *in the BLS game:*

- *The players share the* maximally entangled *state*

$$|\phi_d\rangle = \frac{1}{\sqrt{d}} \sum_i |i\rangle \otimes |i\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B \,, \tag{1.8}$$

  where d *is the dimension of* $\mathcal{H}$*, each of* $\mathcal{H}_A$ *and* $\mathcal{H}_B$ *is a copy of* $\mathcal{H}$*, and* $\{|i\rangle\}$ *an orthonormal basis for it.*[4]

[4]The maximally entangled state is a natural generalization of the EPR pair which can be defined on any tensor product of (finite-dimensional) isomorphic Hilbert spaces.

- *On question j, Alice sequentially measures the observables $Y_{j_1}, Y_{j_2}, \ldots, Y_{j_k}$ on her share of $|\psi\rangle$, where $j_1, \ldots, j_k$ are the elements of E_j. She obtains outcomes $a_1, \ldots, a_k$ that she returns as her answer.*
- *On question $i \in \{1 \ldots, n\}$ Bob measures the observable Y_i^T on his share of $|\psi\rangle$. He obtains an outcome $b \in \{\pm 1\}$ that he returns as his answer.*

Proof. First we note that the strategy described in the lemma is valid: since by definition of an operator solution the observables $Y_{j_1}, Y_{j_2}, \ldots, Y_{j_\ell}$ always commute it is possible for Alice to measure them simultaneously.

The main ingredient of the proof is the relation (1.5). Using this relation it is a matter of direct calculation to verify that the players' answers always satisfy the verifier's checks in the game. In more detail,

- For the consistency check, we note that the probability that the two players return consistent answers on question (j,k) is
$$\frac{1}{2} + \frac{1}{2}\langle\psi| Y_{j_k} \otimes Y_{j_k}^T |\psi\rangle \;=\; \frac{1}{2} + \frac{1}{2}\frac{1}{d}\mathrm{Tr}\big(Y_{j_k}^2\big) \;=\; 1\,,$$
where the first equality follows from (1.5) and the second holds since Y_{j_k} is a binary observable so $Y_{j_k}^2 = \mathrm{Id}$.
- For the equation check, we note that the probability that Alice's answers satisfy the check for the j-th equation is
$$\frac{1}{2} + \frac{c_j}{2}\langle\psi| Y_{j_1} \cdots Y_{j_k} \otimes \mathrm{Id}\, |\psi\rangle \;=\; \frac{1}{2} + \frac{c_j}{2}\langle\psi| c_j\, \mathrm{Id} \otimes \mathrm{Id}\, |\psi\rangle \;=\; 1\,,$$
where the first equality holds since $Y_{j_1} \cdots Y_{j_k} = c_j\, \mathrm{Id}$ by definition of an operator solution.

□

Remark 1.16. The reader will have noticed that in Lemma 1.15 we carefully added the assumption that the operator solution is finite-dimensional, and indeed this seems necessary for the state $|\psi\rangle$ to be well-defined. It is possible to show that infinite-dimensional operator solutions to a BLS correspond to *commuting-operator* strategies for the associated game, and conversely; this correspondence is established in [CLS17]. Commuting-operator strategies are a strict superset of tensor-product strategies

Combining Lemma 1.15 with the operator solution to the Magic Square given by (1.7) we obtain a perfect strategy for the Magic Square game that uses two qubits per player, and two EPR pairs shared between them. Since we saw that the Magic Square does not have a perfect classical strategy this strategy gives us another example of a non-signaling strategy that is not classical.

The following converse to Lemma 1.15 is shown in [CM14].

Lemma 1.17. *Suppose given a BLS (E, c) and a strategy $(|\psi\rangle, A, B)$ for the associated game that succeeds with probability* 1*. Then the BLS has a finite-dimensional operator solution.*

Proof. We give the proof for the special case of the Magic Square game, as the general case is similar. We start with the modeling step: a strategy $(|\psi\rangle, A, B)$ for the magic square game is given by a bipartite state $|\psi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ for finite-dimensional $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ as well as the following measurements. For the first player (Alice), for each row or column j there is a 9-outcome projective measurement $\{A_{ja} : a \in \{\pm 1\}^3\}$ on $\mathcal{H}_\mathsf{A}$. For the second player (Bob), for each variable (square) i there is an observable B_i on $\mathcal{H}_\mathsf{B}$. Note that here we assumed that the measurements made by each player are projective, which is without loss of generality by applying Naimark's theorem and enlarging the spaces $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ if necessary.

To each of Alice's questions we can associate three observables that correspond to the three bits of her answer. For example, for question $j = 1$ (first row) we can define

$$A_1 = \sum_{a_1,a_2,a_3 \in \{\pm 1\}} a_1 \, A_{1 a_1 a_2 a_3} \ , \quad A_2 = \sum_{a_1,a_2,a_3 \in \{\pm 1\}} a_2 \, A_{2 a_1 a_2 a_3} \ , \quad A_3 = \sum_{a_1,a_2,a_3 \in \{\pm 1\}} a_3 \, A_{3 a_1 a_2 a_3} \ .$$

We can similarly proceed to define $A_4, \ldots, A_9$ from the rows and $A'_1, \ldots, A'_9$ from the columns. Next we show that success with probability 1 in the consistency checks implies that

$$\forall i \in \{1, \ldots, 9\} \ , \qquad A_i = A'_i = B_i^T \ . \tag{1.9}$$

Take for example the consistency check on question $(1,2)$ (first row to Alice, second entry to Bob). It is easy to show that success in that check implies that

$$\langle \psi | A_2 \otimes B_2^T | \psi \rangle = 1 \ . \tag{1.10}$$

We use the following claim.

Claim 1.18. *Suppose that $|\psi\rangle$ is a bipartite state and A, B observables such that $\langle \psi | A \otimes B | \psi \rangle = 1$. Let $|\psi\rangle = \sum_t \lambda_t |u_t\rangle |v_t\rangle$ be the Schmidt decomposition of $|\psi\rangle$, with $\lambda_t > 0$ for all t and $\{|u_t\rangle\}$ and $\{|v_t\rangle\}$ orthonormal families. Let $S_A = Span\{|u_t\rangle\} \subseteq \mathcal{H}_A$ and $S_B = Span\{|v_t\rangle\} \subseteq \mathcal{H}_B$. Then S_A is stable by A and S_B is stable by B. Moreover, letting A_S denote the matrix of the restriction of A to S_A in the basis $\{|u_t\rangle\}$ and similarly for B, it holds that $A_S = B_S^T$.*

Proof sketch. Let $K = \sum_t \lambda_t |u_t\rangle\langle v_t|$. Then the equality $\langle \psi | A \otimes B | \psi \rangle = 1$ implies that $|\psi\rangle = A \otimes B |\psi\rangle$, because A, B have norm at most 1. Rearranging indices, this equality is equivalent to $K = AKB^T$. Identifying left and right eigenspaces we see that A and B must each preserve the eigenspaces of K associated with any given eigenvalue. Thus $AKB^T = K$ decomposes in block form $\oplus_\lambda A_\lambda B_\lambda^T = \mathrm{Id}_\lambda$, where for each block we indicated with a subscript λ the restriction of each operator to the eigenspace of K associated with eigenvalue λ. This shows the claim. □

Using Claim 1.18 and the implications of the form (1.10) for the consistency checks, (1.9) follows, where the operators and the transpose should be understood to be written with respect to the Schmidt bases of $|\psi\rangle$. To conclude we claim that $B_1^T, \ldots, B_9^T$ (precisely, their restriction to the support of $|\psi\rangle$ on $\mathcal{H}_B$) are an operator solution to the Magic Square. Commutation in each row or column follows from (1.9) and the definition of the A_y (which by definition commute by rows) and A'_y (by columns). The constraints follow from the fact that e.g. for the first row, $\langle \psi | A_1 A_2 A_3 \otimes \mathrm{Id} \, | \psi \rangle = +1$, which using Claim 1.18 implies that $A_1 A_2 A_3 = \mathrm{Id}$ and hence $B_1^T B_2^T B_3^T = \mathrm{Id}$. (Of course we could remove the transpose signs and still have a valid solution.) □

1.5 Motivation from complexity and cryptography

In theoretical computer science nonlocal games are known as *two-player*, or more generally *multiplayer*, games. Multiplayer games arose independently in complexity theory and cryptography around the 1980s. Their study can be motivated from the following vantage points.

- *Hardness of approximation for constraint satisfaction problems.* We can consider a natural generalization of BCS games to the case of an arbitrary constraint satisfaction problem. Consider the example of 3SAT, and the associated "clause-vs-variable" game. This game is parametrized by a 3SAT formula

which is known to all parties. Once the game starts, the referee selects a clause $C = x \wedge y \wedge z$ (some of the variables may be negated) uniformly at random, as well as a variable $w \in \{x, y, z\}$ appearing in C, again uniformly at random. She sends the triple $\{x, y, z\}$ to Alice, and the single variable w to Bob. Each player is expected to return an assignment to the variables he or she was asked about. The players win if and only if Alice's assignment satisfies the clause C, and Bob's assignment is consistent with Alice's on the variable they were asked in common.

Exercise 1.12. Relate the maximum success probability in the clause-vs-variable game to the largest number of clauses of φ that can be simultaneously satisfied by any assignment. Your relation need not be perfectly tight, but it should at least imply that the maximum success probability is 1 if and only if the formula is satisfiable.

From this construction we get an important consequence: Since 3SAT is NP-hard, it follows that in general the maximum classical success probability in a game specified explicitly (i.e. through a table specifying explicitly the distribution on questions and the truth table for valid answer tuples) is NP-hard to compute.[5] In fact, it follows from the PCP theorem that the maximum success probability is not only hard to compute exactly, but even to approximate within a sufficiently small constant factor. The language of games plays an important role in Dinur's proof of the PCP theorem [Din07], and it has been instrumental in many reductions deriving hardness of approximation for combinatorial problems. It can also be a useful perspective when studying rounding techniques for linear programming (LP) or semidefinite programming (SDP) relaxations of constraint satisfaction problems.

Recall that the SDP formulation given in Section 1.3 implies that $\beta^*(\mathcal{G})$ can be computed in polynomial time. So, for XOR games the maximum success probability of quantum tensor strategies can be computed exactly in polynomial time. What about more general types of clause-vs-variable games? In [KKM$^+$11, IKM09] it was shown, with quite some work, that, for some games, the quantum value remains NP-hard. Much more surprising, and even harder, is that there are families of nonlocal games, indeed BCS games, such that deciding if the quantum bias equals 1 or not is an *undecidable* problem [Slo19]! Unfortunately we will not be able to cover this result here. We return to complexity-theoretic questions around the quantum value in the fourth lecture.

- *Interactive protocols in cryptography.* In cryptography, games play a role as building blocks in *interactive protocols*, where the players are usually referred to as *provers*. A famous game in this context is the two-prover commitment protocol by Ben-Or et al. [BOGKW88]. This protocol was introduced to show that all languages in NP have two-prover interactive proofs with perfect zero-knowledge. Technically the protocol gives rise to a two-round game: the referee first interacts with the first prover (commit phase), and then with the second prover (reveal phase). Many kinds of games arise in cryptography, with the players sometimes exchanging messages between themselves, some players being trusted ("oracles") and others not, etc.

[5] The reason that we need to clarify how the game is specified is that complexity is always a function of the input size. The size of a 3SAT formula is the number of variables and constraints; for the reduction to imply hardness we need the size of the game to be roughly of the same order as this.

Lecture 2

Rigidity

Interestingly, many nonlocal games have the property that the optimal quantum strategy for the players is essentially unique. This phenomenon, called *rigidity*, can be leveraged to devise classical tests (the game) that verify that arbitrary quantum devices (the players) perform specific operations. This opens up a whole new world of possibilities, from the certification of information-theoretic randomness to "device-independent" security proofs in cryptography to protocols for delegated computation; we will touch on some of these topics in the fourth lecture.

Let's look back at the computation of SDP($\mathcal{G}_{CHSH}$) in the proof of Theorem 1.12. If we try to make all inequalities tight, then we don't have much choice. For simplicity, restrict the supremum to real vectors. First of all, for the third line we need to have $\vec{u}_0 = \frac{\vec{v}_0+\vec{v}_1}{\|\vec{v}_0+\vec{v}_1\|}$ and $\vec{u}_1 = \frac{\vec{v}_0-\vec{v}_1}{\|\vec{v}_0-v\vec{v}_1\|}$. So the only freedom is in choosing $\vec{v}_0$ and $\vec{v}_1$. But then to have equality in the application of the Cauchy-Schwarz inequality in the fourth line we need $\|\vec{v}_0 + \vec{v}_1\| = \|\vec{v}_0 - \vec{v}_1\|$, which requires $\vec{v}_0 \cdot \vec{v}_1 = 0$. Conversely, you can check that any two unit vectors $\vec{v}_0$ and $\vec{v}_1$ that are orthogonal will achieve the optimum (provided $\vec{u}_0, \vec{u}_1$ are defined from $\vec{v}_0, \vec{v}_1$ as indicated above). So the only freedom we have is which orthonormal pair to choose for $\vec{v}_0, \vec{v}_1$. However, note that any two orthonormal pairs are related by an orthogonal transformation, and the value of SDP($\mathcal{G}$) is invariant under any orthogonal rotation of the v_j, provided the $\vec{u}_i$ are rotated in the inverse direction. So this last degree of freedom is unavoidable, and we have completely characterized the set of optimal vector solutions.

Exercise 2.1. Suppose $\vec{u}_0, \vec{u}_1, \vec{v}_0, \vec{v}_1$ are real unit vectors that achieve a value of $\frac{\sqrt{2}}{2} - \varepsilon$, for some small $\varepsilon > 0$, in SDP($\mathcal{G}_{\text{CHSH}}$). Is the pair $(\vec{v}_0, \vec{v}_1)$ necessarily close to an orthonormal pair? If so, how would you measure distance to an orthonormal pair?

The preceding exercise asks you to suggest a formulation of "rigidity" of the CHSH game at the level of vectors. Our goal in this lecture is to develop the tools for making similar statements directly at the level of the quantum strategy, i.e. the players' observables and shared quantum state. A result of our investigations will be a theorem stating that the quantum strategy for the CHSH game introduced in Example 1.3 is unique up to local rotations. But we'll go much further than the CHSH game and develop techniques that can be used to show rigidity statements for large classes of games, including all BLS games.

2.1 Approximate group representations

We first make a little detour through the theory of group representations. For d-dimensional matrices A, B and σ such that σ is positive semidefinite, write

$$\langle A, B\rangle_\sigma = \mathrm{Tr}(A^* B\sigma)\ ,$$

where we use B^* to denote the conjugate-transpose. This is an extension of our earlier notation for the matrix inner product, which is recovered for $\sigma = \mathrm{Id}$. If σ is the totally mixed state, then we obtain a dimension-normalized variant of the trace inner product. We will also write $\|A\|_\sigma = \langle A, A\rangle_\sigma^{1/2}$. This is a semi-norm, and it is a norm if σ is invertible. Note that if $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ is any state whose reduced density on the first system equals σ, then $\|A\|_\sigma = \|(A \otimes I)|\psi\rangle\|$

Given an arbitrary finite group G (not necessarily abelian), a group representation of G is a map $f : G \to U_d(\mathbb{C})$, the group of $d \times d$ unitary matrices, such that f is a homomorphism: for any $x, y \in G$, $f(x^{-1}y) = f(x)^* f(y)$, where we used * to denote the conjugate transpose (which, for unitary matrices, corresponds to taking the inverse). The following definition introduces a notion of *approximate* group representation.

Definition 2.1. Given a finite group G, an integer $d \geq 1$, $\varepsilon \geq 0$, and a d-dimensional positive semidefinite matrix σ with trace 1, an (ε, σ)-representation of G is a function $f : G \to U_d(\mathbb{C})$, the unitary group of $d \times d$ matrices, such that

$$\mathrm{E}_{x,y\in G}\, \Re\big(\big\langle f(x)^* f(y), f(x^{-1}y)\big\rangle_\sigma\big) \ \geq\ 1 - \varepsilon\ , \tag{2.1}$$

where the expectation is taken under the uniform distribution over G.

Because $\mathrm{Tr}(\sigma) = 1$ and f is valued in the unitary matrices, (2.1) is equivalent to

$$\mathrm{E}_{x,y\in G}\, \big\|f(x)^* f(y) - f(x^{-1}y)\big\|_\sigma^2 \ \leq\ 2\,\varepsilon\ . \tag{2.2}$$

Remark 2.2. The condition (2.1) in the definition is closely related to the Gowers U^2 norm

$$\|f\|_{U^2}^4 \ =\ \mathrm{E}_{xy^{-1}=zw^{-1}}\, \big\langle f(x)f(y)^*, f(z)f(w)^*\big\rangle_\sigma.$$

While a large Gowers norm implies closeness to an affine function, we are interested in testing homomorphisms. The condition (2.1) will arise naturally from our calculations in the next section.

2.2 The Gowers-Hatami theorem

There are many possible notions for approximate group representation. The most often considered one replaces the norm in (2.2) by the operator norm. A famous theorem of Kazhdan [Kaz82] shows that all amenable groups are stable for the operator norm, i.e. any approximate representation for that norm is proportionately close to an exact representation. Here we are interested in a norm that generalizes the (dimension-normalized) Frobenius norm. In particular, this norm is not sub-multiplicative and hence Kazhdan's proof does not apply. Nevertheless, Gowers and Hatami [GH15] showed that also in the case of Definition 2.1 an approximate group representation can always be "rounded" to a nearby exact representations. We state and prove a slightly more general (but quantitatively weaker) variant of their result.

Theorem 2.3 (Gowers-Hatami)**.** *Let G be a finite group, $\varepsilon \geq 0$, and $f : G \to U_d(\mathbb{C})$ an (ε, σ)-representation of G. Then there exists a $d' \geq d$, an isometry $V : \mathbb{C}^d \to \mathbb{C}^{d'}$, and a representation $g : G \to U_{d'}(\mathbb{C})$ such that*

$$\mathrm{E}_{x \in G} \left\| f(x) - V^* g(x) V \right\|_\sigma^2 \leq 2\,\varepsilon.$$

Gowers and Hatami limit themselves to the case of $\sigma = d^{-1} I_d$, which corresponds to the dimension-normalized Frobenius norm. In this scenario they in addition obtain a tight control of the dimension d', and show that one can always take $d' = (1 + O(\varepsilon))d$ in the theorem. We will see a much shorter proof than theirs (our proof is inspired from the mroe general argument in [DCOT17]) that does not seem to allow to recover this estimate.

Note that Theorem 2.3 does not in general hold with $d' = d$. The reason is that it is possible for G to have an approximate representation in some dimension d, but no exact representation of the same dimension: to obtain an example of this, take any group G that has all non-trivial irreducible representations of large enough dimension, and create an approximate representation in e.g. dimension one less by "cutting off" one row and column from an exact representation. The averaging over all dimensions induced by the matrix σ used to weigh the norm $\|\cdot\|_\sigma$ will in general barely notice this, but it will be impossible to "round" the approximate representation obtained to an exact one without modifying the dimension.

Exercise 2.2. Prove Theorem 2.3 for the case where G is the single-qubit Weyl-Heisenberg group, which is the 8-element matrix group $\mathcal{P}$ generated by the Pauli σ_X and σ_Z matrices. *[Hint: Consider $V : \mathbb{C}^d \to \mathbb{C}^{d'} \otimes \mathbb{C}^2$, where $\mathbb{C}^{d'} \simeq \mathbb{C}^d \otimes \mathbb{C}^2$, defined by*

$$V|\varphi\rangle = \frac{1}{2}\big((\mathrm{Id}_d \otimes \mathrm{Id}_2 + A_0 \otimes \sigma_X + A_1 \otimes \sigma_Z + A_0 A_1 \otimes \sigma_X \sigma_Z) \otimes \mathrm{Id}_2\big)(|\varphi\rangle \otimes |\phi_2\rangle)\,,$$

where $|\varphi\rangle$ is an arbitrary state in $\mathbb{C}^d$, $|\phi_2\rangle$ is an EPR pair on the last two copies of $\mathbb{C}^2$, and $A_0 = f(\sigma_X)$, $A_1 = f(\sigma_Z)$ act on $\mathbb{C}^d$.]

The main ingredient for the proof is an appropriate notion of Fourier transform over non-abelian groups. Given an irreducible representation $\rho : G \to U_{d_\rho}(\mathbb{C})$, define

$$\hat{f}(\rho) \;=\; \mathrm{E}_{x \in G}\, f(x) \otimes \overline{\rho(x)}. \tag{2.3}$$

In case G is abelian, we always have $d_\rho = 1$, the tensor product is a product, and (2.3) reduces to the usual definition of Fourier coefficient. The only properties we will need of irreducible representations is that they satisfy the relation

$$\sum_\rho d_\rho \,\mathrm{Tr}(\rho(x)) \;=\; |G|\delta_{xe}\,, \tag{2.4}$$

for any $x \in G$, where here δ_{xe} is the Kronecker δ. Note that plugging in $x = e$ (the identity element in G) yields $\sum_\rho d_\rho^2 = |G|$.

Proof of Theorem 2.3. Our first step is to define an isometry $V : \mathbb{C}^d \to \mathbb{C}^d \otimes (\oplus_\rho \mathbb{C}^{d_\rho} \otimes \mathbb{C}^{d_\rho})$ by

$$V :\; u \in \mathbb{C}^d \;\mapsto\; \bigoplus_\rho\, d_\rho^{1/2} \sum_{i=1}^{d_\rho} \big(\hat{f}(\rho)(u \otimes e_i)\big) \otimes e_i,$$

where the direct sum ranges over all irreducible representations ρ of G and $\{e_i\}$ is the canonical basis. Note what V does: it "embeds" any vector $u \in \mathbb{C}^d$ into a direct sum, over irreducible representations ρ, of a

d-dimensional vector of $d_\rho \times d_\rho$ matrices. Each (matrix) entry of this vector can be thought of as the Fourier coefficient of the corresponding entry of the vector $f(x)u$ associated with ρ. The fact that V is an isometry follows from the appropriate extension of Parseval's formula:

$$\begin{aligned} V^*V &= \sum_\rho d_\rho \sum_i (I \otimes e_i^*)\hat{f}(\rho)^*\hat{f}(\rho)(I \otimes e_i) \\ &= \mathrm{E}_{x,y} f(x)^* f(y) \sum_\rho d_\rho \sum_i (e_i^* \rho(x)^T \overline{\rho(y)} e_i) \\ &= \sum_\rho \frac{d_\rho^2}{|G|} I = I, \end{aligned}$$

where for the second line we used the definition (2.3) of $\hat{f}(\rho)$ and for the third we used (2.4) and the fact that f takes values in the unitary group.

Next define

$$g(x) = \bigoplus_\rho \big(I_d \otimes I_{d_\rho} \otimes \rho(x)\big),$$

a direct sum over all irreducible representations of G (hence itself a representation). Lets' first compute the "pull-back" of g by V: following a similar calculation as above, for any $x \in G$,

$$\begin{aligned} V^*g(x)V &= \sum_\rho d_\rho \sum_{i,j} (I \otimes e_i^*)\hat{f}(\rho)^*\hat{f}(\rho)(I \otimes e_j) \otimes e_i^*\rho(x)e_j) \\ &= \mathrm{E}_{z,y} f(z)^* f(y) \sum_\rho d_\rho \sum_{i,j} (e_i^*\rho(z)^T\overline{\rho(y)}e_j)\big(e_i^*\rho(x)e_j\big) \\ &= \mathrm{E}_{z,y} f(z)^* f(y) \sum_\rho d_\rho \mathrm{Tr}\big(\rho(z)^T\overline{\rho(y)}\rho(x)^T\big) \\ &= \mathrm{E}_{z,y} f(z)^* f(y) \sum_\rho d_\rho \mathrm{Tr}\big(\rho(z^{-1}yx^{-1})\big) \\ &= \mathrm{E}_z f(z)^* f(zx), \end{aligned}$$

where the last equality uses (2.4). It then follows that

$$\mathrm{E}_x \big\langle f(x), V^*g(x)V \big\rangle_\sigma = \mathrm{E}_{x,z} \mathrm{Tr}\big(f(x)f(zx)^*f(z)\sigma\big).$$

This relates correlation of f with V^*gV to the quality of f as an approximate representation and proves the theorem. □

2.3 Rigidity for the CHSH game

Now let's see what this all has to do with the CHSH game. We will use our newfound theory of approximate group representations to prove the following theorem, originally due to [SW88] (with slightly weaker bounds; see also [MYS12] for the $O(\sqrt{\varepsilon})$ dependence).

Theorem 2.4. *Let $\varepsilon > 0$, and suppose that a strategy for the players in the CHSH game, using a bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ and observables A_0, A_1 for Alice and B_0, B_1 for Bob, achieves a bias at least $\sqrt{2}/2 - \varepsilon$ in the game. Then there are local isometries $V_A, V_B : \mathbb{C}^d \to \mathbb{C}^2 \otimes \mathbb{C}^{d'}$ such that*

$$\big\| V_A \otimes V_B |\psi\rangle - \frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big) \otimes |\psi'\rangle \big\|^2 = O(\sqrt{\varepsilon}), \tag{2.5}$$

and

$$\big\|(V_A \otimes V_B)(A_0 \otimes \mathrm{Id})|\psi\rangle - (\sigma_X \otimes \mathrm{Id})\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \otimes |\psi'\rangle\big\| \;=\; O\big(\sqrt{\varepsilon}\big)\;, \tag{2.6}$$

and a similar relation holds with A_0 replaced by A_1 and σ_X replaced by σ_Z. Moreover, analogous relations hold for Bob's observables.

It is important to note what the theorem says, and what it does not say. It does not say that the state $|\psi\rangle$ shared by the players must be close to an EPR pair — it says that, *up to local rotations*, the state must be close to an EPR pair *tensored with an ancilla state*. Since local unitaries have no effect on the Schmidt coefficients, it does imply that the original state shared by the players have Schmidt coefficients that can be split into two roughly even batches — and in particular, that there are at least two of them.

The theorem also does not say anything about the observables A_0, A_1 themselves. Eq. (2.6) only talks about the action of the observable *on the state*. This is inevitable, as the game only "observes" this action. In particular, it is perfectly possible for A_0 to look like something completely arbitrary in a portion of space in which the reduced density of $|\psi\rangle$ on Alice's space is zero, or very small. But the theorem does say that, in terms of "observable consequences" only, the action of A_0 on $|\psi\rangle$ is comparable to the action of σ_X on one half of an EPR pair. Although this may sound relatively weak, we will later see that this fact can be extremely useful in applications.

Proof. For the first step of the proof we follow Tsirelson's argument showing a bound of $\frac{\sqrt{2}}{2}$ on the quantum value of the CHSH game. Tsirelson's idea was to consider the square

$$\begin{aligned}(A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1)^2 &= ((A_0 + A_1) \otimes B_0 + (A_0 - A_1) \otimes B_1)^2 \\ &= 4\,\mathrm{Id} \otimes \mathrm{Id} + (A_1A_0 - A_0A_1) \otimes (B_0B_1 - B_1B_0)\;.\end{aligned} \tag{2.7}$$

This last term has operator norm at most 8, and as a result the CHSH operator $A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1$ has operator norm at most $\sqrt{8}$; Tsirelson's bound on the quantum value of CHSH follows by dividing by 4 and observing that the players' choice of entangled state cannot beat the largest eigenvector.

Furthermore, under the assumption that the strategy achieves a bias of at least $\sqrt{2}/2 - \varepsilon$ in the game, using $|\langle\psi|X|\psi\rangle|^2 \leq \langle\psi|XX^*|\psi\rangle$ for any Hermitian X, the left-hand side of (2.7), when evaluated on $|\psi\rangle$, must be at least $8(1-\varepsilon)^2$. Applying the Cauchy-Schwarz inequality it follows that both conditions

$$\mathrm{Tr}((A_1A_0 + A_0A_1)^2\rho_A) = O(\varepsilon) \qquad \text{and} \qquad \mathrm{Tr}((B_0B_1 + B_1B_1)^2\rho_B) = O(\varepsilon)$$

must hold, where ρ_A and ρ_B are the reduced density matrices of $|\psi\rangle$ on Alice and Bob respectively. Using the notation introduced in Section 2.1, this says that

$$\|A_0A_1 + A_1A_0\|_{\rho_A}^2 = O(\varepsilon)\;, \tag{2.8}$$

i.e. A_0 and A_1 approximately commute.

Now here comes the key observation. Consider the 8-element group $\mathcal{P}$ generated by the Pauli matrices σ_X and σ_Z, i.e.

$$\mathcal{P} \;=\; \pm\big\{\,\mathrm{Id},\,\sigma_X,\,\sigma_Z,\,\sigma_X\sigma_Z\big\}\;.$$

Then we claim that A_0 and A_1 induce an approximate representation of $\mathcal{P}$, by setting

$$f(\pm\,\mathrm{Id}) = \pm\,\mathrm{Id}, \quad f(\pm\sigma_X) = \pm A_0, \qquad f(\pm\sigma_Z) = \pm A_1, \quad f(\pm\sigma_X\sigma_Z) = \pm A_0A_1\;.$$

Note that this is a legal definition, since A_0, A_1, and A_0A_1 are all unitary. Moreover, using only (2.8) and the fact that A_0 and A_1 are observables, it is immediate to verify that the conditions of Theorem 2.3 are satisfied, i.e. f is an $(O(\varepsilon), \rho_A)$-representation of $\mathcal{P}$.

Applying the theorem, there must exist an exact representation of $\mathcal{P}$ to which f is close. However, the representation theory of $\mathcal{P}$ is not complicated. It has four 1-dimensional representations, but all of them map $-\mathrm{Id}$ to 1, so they cannot be close to f. Hence we are left with the unique irreducible 2-dimensional representation of $\mathcal{P}$, which is precisely given by the Pauli matrices! With a little more careful work, this shows (2.6). Moreover, we can apply the same considerations to Bob's operators B_0 and B_1, except that here we will choose to rotate the representation so that it sends σ_X to the Hadamard matrix H and σ_Z to the matrix G. This is another valid representation; it is the same as the standard one, but rotated.

Finally we need to verify the condition on $|\psi\rangle$. Note that by assumption

$$\frac{1}{4}\langle\psi|A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1|\psi\rangle \geq \frac{\sqrt{2}}{2} - \varepsilon \,,$$

which using (2.6) implies

$$\frac{1}{4}\langle\psi|(V_A \otimes V_B)^\dagger\big((\sigma_X \otimes H + \sigma_X \otimes G + \sigma_Z \otimes H - \sigma_Z \otimes G) \otimes \mathrm{Id}\,\big)(V_A \otimes V_B)|\psi\rangle \geq \frac{\sqrt{2}}{2} - O(\varepsilon) \,,$$

i.e.

$$\frac{1}{2}\langle\psi|(V_A \otimes V_B)^\dagger\big((\sigma_X \otimes \sigma_X + \sigma_Z \otimes \sigma_Z) \otimes \mathrm{Id}\,\big)(V_A \otimes V_B)|\psi\rangle \geq 1 - O(\varepsilon) \,. \tag{2.9}$$

Now observe that $\frac{1}{2}(\sigma_X \otimes \sigma_X + \sigma_Z \otimes \sigma_Z)$ is an observable with a single eigenvalue 1, with associated eigenvector $|\phi_2\rangle$, and all other eigenvalues equal to 0 or -1. Therefore, (2.9) implies

$$\big\|\big(\langle\phi_2| \otimes \mathrm{Id}\,\big)(V_A \otimes V_B)|\psi\rangle\big\|^2 \geq 1 - O(\varepsilon) \,,$$

which means that $(V_A \otimes V_B)|\psi\rangle = |\phi_2\rangle \otimes |\psi'\rangle + |\psi''\rangle$ for some sub-normalized states $|\psi'\rangle$ and $|\psi''\rangle$ such that $\||\psi''\rangle\|^2 = O(\varepsilon)$. This proves the theorem. □

2.4 Rigidity for the Magic Square game

We now investigate rigidity for the Magic Square game. As for the CHSH game, we will show that any near-optimal strategy has a specific structure by exhibiting (anti-)commutation relations that must be satisfied by any such strategy in the game.

2.4.1 The exact case

As a warm-up, let's investigate the structure of operator solutions to the BLS system that underlies the Magic Square game.

Lemma 2.5. *Suppose given an operator solution $Y_1, \ldots, Y_9$ to the magic square BLS. Then Y_2 and Y_4 anti-commute.*

Proof. We first rewrite the product Y_2Y_4 by rows to obtain

$$\begin{aligned} Y_2Y_4 &= Y_1Y_3 \cdot Y_6Y_5 \\ &= Y_1 \cdot (-Y_9) \cdot Y_5 \,, \end{aligned}$$

where the second line is by the last column constraint. Next we write the product Y_4Y_2 by columns:

$$\begin{aligned} Y_4Y_2 &= Y_1Y_7 \cdot Y_8Y_5 \\ &= Y_1 \cdot Y_9 \cdot Y_5 \,, \end{aligned}$$

where the second line is by the last row constraint. Combining both equations it follows that $Y_2Y_4 = -Y_4Y_2$, as claimed. □

Note that the same proof shows that any two observables that are not in the same row or column must anti-commute. Using furthermore the commutation conditions which are imposed for each row and column, it is possible to obtain the following full characterization of operator solutions, which we leave as an exercise.

Exercise 2.3. Let $Y_1, \ldots, Y_9$ be an operator solution to the Magic Square system. Show that there is an orthonormal basis with respect to which

$$\begin{array}{lll} Y_1 = (I_2 \otimes \sigma_Z) \otimes \mathrm{Id} & Y_2 = (\sigma_Z \otimes I_2) \otimes \mathrm{Id} & Y_3 = (\sigma_Z \otimes \sigma_Z) \otimes \mathrm{Id} \\ Y_3 = (\sigma_X \otimes I_2) \otimes \mathrm{Id} & Y_4 = (I_2 \otimes \sigma_X) \otimes \mathrm{Id} & Y_5 = (\sigma_X \otimes \sigma_X) \otimes \mathrm{Id} \,, \\ Y_7 = (\sigma_X \otimes \sigma_Z) \otimes \mathrm{Id} & Y_8 = (\sigma_Z \otimes \sigma_X) \otimes \mathrm{Id} & Y_9 = (\sigma_Y \otimes \sigma_Y) \otimes \mathrm{Id} \end{array}$$

where I_2 denotes the identity on $\mathbb{C}^2$ and Id is the identity on $\mathbb{C}^d$ for some d.

The following lemma is immediate from the proof of Lemma 1.17 and Lemma 2.5.

Lemma 2.6. *Suppose that a quantum tensor correlation*

$$p(a,b|x,y) \;=\; \langle \psi | A_{x,a} \otimes B_{y,b} | \psi \rangle$$

perfectly satisfies the referee's tests in the Magic Square game. Let S_B denote the support of the reduced density ρ_B of $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ on $\mathcal{H}_B$. Then the observables $B_1, \ldots, B_9$ stabilize S_B, and their restriction to that space form an operator solution to the Magic Square.

Proof. The first part of the lemma follows from the proof of Lemma 1.17. By Lemma 2.5 the observables associated to $y = 2$ and $y = 4$ anti-commute. □

As a last step we show that we can also characterize the entangled state used by any strategy. Interestingly, this characterization comes as a consequence of the characterization of the observables, which we obtained without talking much about the state. This is based on the following general lemma. A version of the lemma was already used implicitly at the end of the proof of Theorem 2.4.

Lemma 2.7. *Let $|\psi\rangle_{ABE} \in (\mathbb{C}^2)_A^{\otimes n} \otimes (\mathbb{C}^2)_B^{\otimes n} \otimes \mathcal{H}_E$ be such that for every $i \in \{1, \ldots, n\}$ it holds that*

$$\left(\sigma_{X,i}\right)_A \otimes \left(\sigma_{X,i}\right)_B |\psi\rangle_{ABE} \;=\; \left(\sigma_{Z,i}\right)_A \otimes \left(\sigma_{Z,i}\right)_B |\psi\rangle_{ABE} \;=\; |\psi\rangle_{ABE} \,,$$

where the Pauli operators act on the i-th copy of $\mathbb{C}^2$ in register A and B respectively. Then $|\psi\rangle_{ABE} = |\phi_2\rangle_{AB}^{\otimes n} \otimes |aux\rangle$, for some state $|aux\rangle$ on $\mathcal{H}$.

Proof. Note that $\sigma_X \otimes \sigma_X$ and $\sigma_Z \otimes \sigma_Z$ commute, hence are simultaneously diagonalizable. The proof immediately follows from the observation that the only simultaneous eigenvalue-1 eigenstate of $\sigma_X \otimes \sigma_X$ and $\sigma_Z \otimes \sigma_Z$ is the EPR pair $|\phi_2\rangle$. □

Exercise 2.4. Show that the conclusion of Lemma 2.7 holds under the following weaker assumption: $|\psi\rangle_{\mathsf{ABE}} \in (\mathbb{C}^2)_{\mathsf{A}}^{\otimes n} \otimes \mathcal{H}_{\mathsf{B}}^{\otimes n} \otimes \mathcal{H}_{\mathsf{E}}$ with $\mathcal{H}_{\mathsf{B}}$ arbitrary, and for every $i \in \{1, \ldots, n\}$,

$$\left(\sigma_{X,i}\right)_{\mathsf{A}} \otimes \left(X_i\right)_{\mathsf{B}} |\psi\rangle_{\mathsf{ABE}} \;=\; \left(\sigma_{Z,i}\right)_{\mathsf{A}} \otimes \left(Z_i\right)_{\mathsf{B}} |\psi\rangle_{\mathsf{ABE}} \;=\; |\psi\rangle_{\mathsf{ABE}} \,,$$

with X_i and Z_i arbitrary binary observables on $\mathcal{H}_{\mathsf{B}}$. *[Hint: Make a careful use of Claim 1.18]*

2.4.2 Consequences

The characterization of operator solutions from Exercise 2.3 together with Lemma 1.17 and Lemma 2.7 have some nice consequences. First of all, they imply that the Magic Square game tests not one, but two qubits: any perfect strategy must have a 4-qubit entangled state, two qubits per player, and Bob's observables specify two qubits, e.g. B_2 and B_4 for the first and B_1 and B_5 for the second. We even have access to more: for example, we know that when Bob is asked question 9 the observable he applies is $\sigma_Y \otimes \sigma_Y$.

Another consequence of the characterization has to do with the problem of randomness certification. At this point we know that, in any perfect strategy, whenever Bob is asked question 2 he measures the first qubit of an EPR pair in the standard basis. This has the following implications:

1. The answer reported by Bob on question 2 (and, in fact, on *any* question) is a uniformly random bit. In particular, no deterministic strategy can succeed in the game! We knew this already, because deterministic strategies are classical. As such, any game for which quantum strategies can succeed with strictly higher probability than classical strategies can serve as a "test for randomness".

2. More importantly, the randomness that is generated by Bob at each execution of the game is "fresh" and "private". What we mean by this is that Bob's random bit is (1) independent of any information at the verifier's side, including Bob's question, and (2) uncorrelated to the environment. Indeed, since Bob's bit is the result of a measurement of half an EPR pair, the only party that can obtain correlated information is Alice, who holds the other half of the EPR pair. By the rigidity theorem this EPR pair *must* be in control of Alice: she needs it for them to succeed in the game. Therefore the verifier has the guarantee that the bit she obtains (1) cannot have been "planted" *a priori* in the devices, and (2) cannot be learned, even partially, by any third party distinct from A and B, even if the party could a priori have kept entanglement with the devices—this is because, using the notation of Lemma 2.7, the third party would only at best have access to the entirety of system E, which is uncorrelated with AB.

These observations are important for cryptography, where the use of high-quality randomness that is uncorrelated from any possible eavesdropper or adversary is an essential resource. Indeed, the observations we just made form the basis for the so-called "device-independent" analysis of quantum cryptography protocols.

An important drawback of our analysis so far is that it is limited to the case of perfect strategies, i.e. strategies that succeed with probability 1 in the game. In practice one may only reasonably assume, after multiple executions of the game, that a given strategy succeeds with some probability that is close to one, $1-\varepsilon$ for some $\varepsilon \geq 0$ that can be made small but not 0. In the next section we discuss how the results can be extended to that case.

2.4.3 The approximate case

The following theorem gives the flavor of an approximate version of the lemmas from the previous section. It is taken from [CS17], where more general statements are shown for any BLS that satisfies appropriate conditions. (A similar result specialized to the case of the Magic Square game is shown in [WBMS16].)

Theorem 2.8. *Suppose that a strategy $(|\psi\rangle, A, B)$ succeeds with probability $1-\varepsilon$ in the Magic Square game, for some $\varepsilon \geq 0$. Then there are isometries $V_D : \mathcal{H}_D \to \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathcal{H}_{D'}$ for $D \in \{A, B\}$ such that*

$$\left\| V_A \otimes V_B |\psi\rangle_{AB} - |\phi_2\rangle \otimes |\phi_2\rangle \otimes |aux\rangle \right\|^2 = O\left(\sqrt{\varepsilon}\right) ,$$

for some state $|aux\rangle$ *on* $\mathcal{H}_{A'} \otimes \mathcal{H}_{B'}$*, and*

$$\big\| \mathrm{Id}_A \otimes \big(V_B B_2 - (\sigma_Z \otimes \mathrm{Id} \otimes \mathrm{Id}) V_B\big) |\psi\rangle_{AB} \big\|^2 = O(\varepsilon) \ , \tag{2.10}$$

$$\big\| \mathrm{Id}_A \otimes \big(V_B B_4 - (\sigma_X \otimes \mathrm{Id} \otimes \mathrm{Id}) V_B\big) |\psi\rangle_{AB} \big\|^2 = O(\varepsilon) \ , \tag{2.11}$$

and similar relations hold for the remaining seven observables on Bob's side.

The proof of Theorem 2.8 follows from Theorem 2.3 in a similar way as the proof of Theorem 2.4.

Proof. For the first step of the proof we follow the proof of Lemma 2.5, except that all equalities must be made approximate equalities. Assume without loss of generality that Alice's strategy is specified by six 4-outcomes projective measurements $\{A_{i,a_1,a_2,a_3}\}$, where $a_1, a_2, a_3 \in \{\pm 1\}$ range over the four possible assignments that satisfy the constraint associated with the i-th row ($i \in \{1,2,3\}$) or $(i-3)$-th column ($i \in \{4,5,6\}$). (We can assume that Alice always returns a valid assignment because she knows that she will lose if not, so we do not even need to include a symbol for such evidently wrong answers in the game.)

For any $i \in \{1,\ldots,9\}$, in addition to Bob's observable B_i associated with question i we define two observables from Alice's strategy, obtained by recording her answer associated to entry i when she is asked the unique row containing i — call this observable C_i — or the unique column containing i — call this observable C'_i. Formally, $C_1 = \sum_{a_1,a_2,a_3 \in \{\pm 1\}} a_1 A_{1,a_1,a_2,a_3}$, and similar relations can be used to define each C_i and C'_i. Due to the parity constraints enforced on Alice's answers it is always the case that $C_1 C_2 C_3 = +\mathrm{Id}, \ldots, C_3 C_6 C_9 = -\mathrm{Id}$.

By definition, success of the strategy in the game implies 18 equations of the form

$$\langle \psi | C_i \otimes B_i | \psi \rangle \ \geq \ 1 - 18\varepsilon \qquad \text{and} \qquad \langle \psi | C'_i \otimes B_i | \psi \rangle \ \geq \ 1 - 18\varepsilon \ , \tag{2.12}$$

for all $i \in \{1,\ldots,9\}$. Indeed, each such equation represents the probability that the players return valid answers, conditioned on each of the 18 possible pairs of questions in the game. These relations allow us to mimic the proof of Lemma 2.5, as follows:

$$\begin{aligned}
B_2 B_4 |\psi\rangle &\approx C_4 \otimes B_2 |\psi\rangle \\
&= C_5 C_6 \otimes B_2 |\psi\rangle \\
&\approx C_5 C_6 C_2 \otimes \mathrm{Id} \, |\psi\rangle \\
&= C_5 C_6 C_3 C_1 \otimes \mathrm{Id} \, |\psi\rangle \\
&\approx C_5 C_6 C_3 \otimes B_1 |\psi\rangle \\
&\approx C_5 C_6 \otimes B_1 B_3 |\psi\rangle \\
&\approx C_5 C'_6 \otimes B_1 B_3 |\psi\rangle \\
&\approx C_5 C'_6 C'_3 \otimes B_1 |\psi\rangle \\
&= C_5 C'_9 \otimes B_1 |\psi\rangle \ ,
\end{aligned}$$

where here we use the notation $|u\rangle \approx |v\rangle$ to mean $\||u\rangle - |v\rangle\|^2 = O(\varepsilon)$. Here, each of the approximations is obtained by bounding the squared norm of the difference using the Cauchy-Scharz inequality and the required relation; for example, for the first approximation we write

$$\begin{aligned}
\big\| (\mathrm{Id} \otimes B_2 B_4 - C_4 \otimes B_2) |\psi\rangle \big\|^2 &= \big\| (\mathrm{Id} \otimes B_2)(\mathrm{Id} \otimes B_4 - C_4 \otimes \mathrm{Id}) |\psi\rangle \big\|^2 \\
&= \big\| (\mathrm{Id} \otimes B_4 - C_4 \otimes \mathrm{Id}) |\psi\rangle \big\|^2 \\
&= 2 - 2\langle \psi | B_4 \otimes C_4 | \psi \rangle \\
&\leq 36\varepsilon \ ,
\end{aligned}$$

where the derivation uses that B_2, B_4 and C_4 are Hermitian and square to identity. Using a similar chain of approximations starting from $B_4 B_2 |\psi\rangle$, we conclude that $\| \mathrm{Id} \otimes \{B_2, B_4\} |\psi\rangle \| = O(\sqrt{\varepsilon})$. It follows that B_2 and B_4 induce an approximate representation of the group $\mathcal{P}$ introduced in Exercise 2.2 by setting

$$f(\pm \mathrm{Id}) = \pm \mathrm{Id}, \quad f(\pm\sigma_Z) = \pm B_2, \qquad f(\pm\sigma_X) = \pm B_4, \quad f(\pm\sigma_X\sigma_Z) = \pm B_4 B_2 \,.$$

Note that this is a legal definition, since B_2, B_4, and $B_2 B_4$ are all unitary. Moreover, using only the approximate anti-commutation and the fact that B_2 and B_4 are observables it is immediate to verify that the conditions of Theorem 2.3 are satisfied, i.e. f is an $(O(\varepsilon), \rho_{\mathsf{B}})$-representation of $\mathcal{P}$, where ρ_{B} denotes the reduced density of $|\psi\rangle$ on $\mathcal{H}_{\mathsf{B}}$.

Applying the theorem, there must exist an exact representation g of $\mathcal{P}$ to which f is close. The proof then concludes exactly as the proof of Theorem 2.4. □

Lecture 3

The Pauli braiding test

In this lecture we build up our work on nonlocal games and rigidity completed thus far and extend it to design a family of games $(G_n)_{n\geq 1}$ such that near-optimal strategies in G_n require n, as opposed to 1 or 2, EPR pairs. There are multiple ways in which this can be done. A natural approach would be to consider the parallel repetition of, for example, the Magic Square game. That is, we simply play n copies of the game in parallel and accept if all answers are valid. It can be shown that perfect or even near-optimal success probabilities in this game require strategies that make use of n (or $2n$ in the case of the Magic Square) EPR pairs, see e.g. [CN16]. However, this approach has multiple drawbacks. Firstly, the number of questions and answers in the parallel-repeated game scales exponentially with n. Secondly, for all results known the "robustness" degrades with n, i.e. if the success probability is $1-\varepsilon$ then closeness to n EPR pairs is only up to $\mathrm{poly}(n)\cdot\varepsilon^c$ for some small constant $c>0$. Both of these aspects are problematic in applications to complexity.

In this lecture we give a different method, which is based on leveraging the connection with approximate representation theory we discovered in the previous lecture, and gets around the second issue–there is no dependence of the robustness on n. Furthermore, the total number of answers in the game remains a constant. However, the number of questions still scales exponentially. At the end of the lecture we will describe a further improvement that uses only a polynomial number of questions and is based on an "efficient" variant of stability.

Based on our analysis of the CHSH and Magic Square games, to obtain a game that tests for n EPR pairs we should go through the following steps: (i) design a game such that success in the game requires the players to share observables that satisfy all relations that we expect from elements of the group $\mathcal{P}_n$ generated by n mutually commuting pairs of anti-commuting observables $(X_1,Z_1),\ldots,(X_n,Z_n)$, (ii) apply Theorem 2.3 to obtain closeness of the strategy to an exact representation of this group, (iii) use that, hopefully, all non-trivial representations give us what we want, i.e. something that looks like the tensor product of n copies of the Pauli group, and (iv) apply Lemma 2.7 to conclude that the shared state is locally isometric to the tensor product of n EPR pairs.

We start in Section 3.1 by studying the group $\mathcal{P}_n$. Then in Section 3.2 we design "tests" for the group product relation.

3.1 The Weyl-Heisenberg group

Denote by $\mathcal{P}_n$ the "n-qubit Weyl-Heisenberg group," i.e. the matrix group generated by n-fold tensor products of single-qubit σ_X and σ_Z matrices. The group $\mathcal{P}_n$ has cardinality $2\cdot 4^n$, and each element of $\mathcal{P}_n$ has

a unique representative of the form $\pm\sigma_X(a)\sigma_Z(b)$ for $a,b \in \{0,1\}^n$. The irreducible representations of $\mathcal{P}_n$ are easily computed from those of $\mathcal{P}$. For us the only thing that matters is that the only irreducible representation g which satisfies $g(-\mathrm{Id}) = -g(\mathrm{Id})$ has dimension 2^n and is given by the defining matrix representation. All other irreducible representations have dimension 1: there are 4^n of them, which are all possible products of the four dimension-1 irreducible representations of $\mathcal{P} = \mathcal{P}_2$.

We now state a version of the Gowers-Hatami theorem tailored to the group $\mathcal{P}_n$ and a specific choice of presentation for the group relations.

Corollary 3.1. *Let n, d be integer, $\varepsilon \geq 0$, $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ a permutation-invariant state, σ the reduced density of $|\psi\rangle$ on either system, and $f : \{X,Z\} \times \{0,1\}^n \to Obs(\mathbb{C}^d)$, the set of observables on $\mathbb{C}^d$. For $a,b \in \{0,1\}^n$ let $X(a) = f(X,a)$, $Z(b) = f(Z,b)$, and assume that $X(a)^2 = Z(b)^2 = I_d$ for all a,b. Suppose that the following inequalities hold: consistency*

$$\mathrm{E}_a \langle\psi|\big(X(a)\otimes X(a)\big)|\psi\rangle \geq 1-\varepsilon\,, \qquad \mathrm{E}_b \langle\psi|\big(Z(b)\otimes Z(b)\big)|\psi\rangle \geq 1-\varepsilon\,, \tag{3.1}$$

linearity

$$\mathrm{E}_{a,a'} \big\|X(a)X(a') - X(a+a')\big\|_\sigma^2 \leq \varepsilon\,, \qquad \mathrm{E}_{b,b'} \big\|Z(b)Z(b') - Z(b+b')\big\|_\sigma^2 \leq \varepsilon\,, \tag{3.2}$$

and anti-commutation

$$\mathrm{E}_{a,b} \big\|X(a)Z(b) - (-1)^{a\cdot b}X(a)Z(b)\big\|_\sigma^2 \leq \varepsilon\,. \tag{3.3}$$

Then there exists a $d' \geq d$, an isometry $V : \mathbb{C}^d \to \mathbb{C}^{d'}$, and a representation $g : \mathcal{P}_n \to U_{d'}(\mathbb{C})$ such that $g(-I) = -I_{d'}$ and

$$\mathrm{E}_{a,b} \big\|X(a)Z(b) - V^* g(\sigma_X(a)\sigma_Z(b))V\big\|_\sigma^2 = O(\varepsilon)\,.$$

Note that the conditions (3.2) and (3.3) in the corollary are very similar to the conditions required of an approximate representation of the group $\mathcal{P}_n$; in fact it is easy to convince oneself that their exact analogue suffices to imply all the group relations. The reason for choosing those specific relations is that they can be checked using games; see the next subsection for this. Condition (3.1) is necessary to derive the conditions for the application of the Gowers-Hatami theorem from (3.2) and (3.3), and is also testable; see the proof.

Nevertheless, note that the number of relations checked in Corollary 3.1 is only $O(4^n)$. This is in contrast to the number of relations that would be checked by applying Theorem 2.3 directly, which is $|\mathcal{P}_n|^2 = O(8^n)$. Thus Corollary 3.1 can be seen as a form of "efficient stability" result, where checking a small subset of all relations suffices to obtain a similar conclusion as checking all of them. We return to this topic at the end of the lecture.

Remark 3.2. Corollary 3.1 can be seen as an extension of the Blum-Luby-Rubinfeld linearity test [BLR93]. The latter makes a similar statement, but for the abelian group $\{\pm\sigma_X(a) \,|\, a \in \{0,1\}^n\} \simeq \mathbb{Z}_2^n$.

Proof. To apply the Gowers-Hatami theorem we need to construct an (ε,σ)-representation f of the group $\mathcal{P}_n$. Using that any element of $\mathcal{P}_n$ has a unique representative of the form $\pm\sigma_X(a)\sigma_Z(b)$ for $a,b \in \{0,1\}^n$, we define $f(\pm\sigma_X(a)\sigma_Z(b)) = \pm X(a)Z(b)$. Next we need to verify that f is an approximate representation. Let $x,y \in \mathcal{P}_n$ be such that $x = \sigma_X(a_x)\sigma_Z(b_x)$ and $y = \sigma_X(a_y)\sigma_Z(b_y)$ for n-bit strings (a_x,b_x) and (a_y,b_y) respectively. Up to phase, we can exploit successive cancellations to decompose $(f(x)f(y)^* - f(xy^{-1}))\otimes$

I as

$$\begin{aligned}
&\big(X(a_x)Z(b_x)X(a_y)Z(b_y) - (-1)^{a_y\cdot b_x}X(a_x+a_y)Z(b_x+b_y)\big)\otimes I \\
&\quad = X(a_x)Z(b_x)X(a_y)\big(Z(b_y)\otimes I - I\otimes Z(b_y)\big) \\
&\qquad +X(a_x)\big(Z(b_x)X(a_y) - (-1)^{a_y\cdot b_x}X(a_y)Z(b_x)\big)\otimes Z(b_y) \\
&\qquad +(-1)^{a_y\cdot b_x}\big(X(a_x)X(a_y)\otimes Z(b_y)\big)\big(Z(b_x)\otimes I - I\otimes Z(b_x)\big) \\
&\qquad +(-1)^{a_y\cdot b_x}\big(X(a_x)X(a_y)\otimes Z(b_y)Z(b_x) - X(a_x+a_y)\otimes Z(b_x+b_y)\big) \\
&\qquad +(-1)^{a_y\cdot b_x}\big(X(a_x+a_y)\otimes I\big)\big(I\otimes Z(b_x+b_y) - Z(b_x+b_y)\otimes I\big).
\end{aligned}$$

(It is worth staring at this equality for a little bit. In particular, note the "player-switching" that takes place in the 2nd, 4th and 6th lines; this is used as a means to "commute" the appropriate unitaries, and is the reason for including (3.1) among the assumptions of the corollary; indeed it is (3.1) that guarantees that those terms are small.) Evaluating each term on the state $|\psi\rangle$, taking the squared Euclidean norm, and then the expectation over uniformly random a_x, a_y, b_x, b_y, the inequality $\|AB|\psi\rangle\| \leq \|A\|\|B|\psi\rangle\|$ and the assumptions of the theorem let us bound the overlap of each term in the resulting summation by $O(\varepsilon)$. Using $\|(A\otimes I)|\psi\rangle\| = \|A\|_\sigma$ by definition and the triangle inequality we have obtained the bound

$$\mathrm{E}_{x,y}\left\|f(x)f(y)^* - f(xy^{-1})\right\|_\sigma^2 = O(\varepsilon).$$

We are now in a position to apply the Gowers-Hatami theorem, which gives an isometry V and exact representation g such that

$$\mathrm{E}_{a,b}\left\|X(a)Z(b) - \frac{1}{2}V^*\big(g(\sigma_X(a)\sigma_Z(b)) - g(-\sigma_X(a)\sigma_Z(b))\big)V\right\|_\sigma^2 = O(\varepsilon). \tag{3.4}$$

Using that g is a representation, $g(-\sigma_X(a)\sigma_Z(b)) = g(-I)g(\sigma_X(a)\sigma_Z(b))$. It follows from (3.4) that $\|g(-I)+I\|_\sigma^2 = O(\varepsilon)$, so we may restrict the range of V to the subspace where $g(-I) = -I$ without introducing much additional error. □

3.2 Testing the Weyl-Heisenberg group relations

Corollary 3.1 makes three assumptions about the observables $X(a)$ and $Z(b)$: that they satisfy approximate consistency (3.1), linearity (3.2), and anti-commutation (3.3). To complete our test we need to show how these relations can be "certified" in a two-player game, i.e. we want a nonlocal game such that achieving a high enough success probability in the game requires that a certain collection of observables, defined from the players' strategy in the game, satisfies the assumptions of the corollary. There are multiple ways that this can be done; we give one. We start by introducing two stand-alone "tests."

Linearity test:

(a) The referee selects $W \in \{X, Z\}$ and $a, a' \in \{0,1\}^n$ uniformly at random. She sends (W, a, a') to one player and (W, a), (W, a'), or $(W, a+a')$ to the other, where the sum is taken modulo 2.[1]

[1] Elements such as (W, a) are labels sent to the players as their question, and carry no other intrinsic meaning.

(b) The first player replies with two bits e_1, e_2, and the second with a single bit f. The referee accepts if and only if the player's answers satisfy the natural relation, e.g. if the third player received $a + a'$ then it should be that $f = e_1 + e_2 \mod 2$.

As always in this section, the test treats both players symmetrically. As a result we can assume that the players' strategy is symmetric, and is specified by a permutation-invariant state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ and a measurement for each question: an observable $W(a)$ associated to questions of the form (W, a), and a four-outcome measurement $\{W_{a,a'}\}$ associated with questions of the form (W, a, a'). The following exercise asks you to verify this fact.

Exercise 3.1. A game is called *symmetric* if $\mathcal{X} = \mathcal{Y}$ and $\mathcal{A} = \mathcal{B}$, the distribution on questions π is invariant under permutation of the two questions, $\pi(x, y) = \pi(y, x)$ for all (x, y), and the verification predicate is symmetric as well, i.e. $V(x, y, a, b) = V(y, x, b, a)$ for all (x, y, a, b). Define a strategy $(|\psi\rangle, \{A_{xa}\}, \{B_{yb}\})$ to be *symmetric* if $\mathcal{H}_\mathsf{A} = \mathcal{H}_\mathsf{B}$, $|\psi\rangle$ is invariant under exchange of the two subsystems, and $A_{xa} = B_{xa}$ for all x, a.

Show that whenever a game $\mathcal{G}$ is symmetric then for any strategy $(|\psi\rangle, \{A_{xa}\}, \{B_{yb}\})$ that succeeds with some probability p in the game there is a symmetric strategy $(|\tilde{\psi}\rangle, \{\tilde{A}_{xa}\})$ that succeeds with the same probability.

The linearity test described above is almost identical to the BLR linearity test, except for the use of the basis label $W \in \{X, Z\}$. The following lemma states conditions that a strategy must satisfy in order to succeed with high probability in the test.

Lemma 3.3. *Suppose that a family of observables $\{W(a)\}$ for $W \in \{X, Z\}$ and $a \in \{0,1\}^n$, generates outcomes that succeed in the linearity test with probability $1 - \varepsilon$, when applied on a symmetric bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ with reduced density matix σ. Then the following hold: approximate consistency*

$$\mathrm{E}_a \langle\psi| \big(X(a) \otimes X(a)\big) |\psi\rangle \;=\; 1 - O(\varepsilon), \qquad \mathrm{E}_b \langle\psi| \big(Z(b) \otimes Z(b)\big) |\psi\rangle \;\geq\; 1 - O(\varepsilon),$$

and linearity

$$\mathrm{E}_{a,a'} \big\| X(a)X(a') - X(a + a') \big\|_\sigma^2 = O(\varepsilon), \qquad \mathrm{E}_{b,b'} \big\| Z(b)Z(b') - Z(b + b') \big\|_\sigma^2 \;=\; O(\varepsilon).$$

Exercise 3.2. Prove the lemma. (In the case of classical strategies, the conditions are an immediate reformulation of the test. The proof for quantum strategies is not much harder.)

Testing anti-commutation can be done using the Magic Square game.

Anti-commutation test:

(a) The referee selects $a, b \in \{0,1\}^n$ uniformly at random under the condition that $a \cdot b = 1$. She plays the Magic Square game with both players, with the following modifications: if the question to the second player is 2 or 4 she sends (X, a) or (Z, b) instead; in all other cases he sends the original label of the question in the Magic Square game together with both strings a and b.

(b) Each player provides answers as in the Magic Square game. The referee accepts if and only if the player's answers would have been accepted in the game.

Using Theorem 2.8 it is straightforward to show the following.

Lemma 3.4. *Suppose a strategy for the players succeeds in the anti-commutation test with probability at least $1-\varepsilon$, when performed on a symmetric bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ with reduced density matrix σ. Then the observables $X(a)$ and $Z(b)$ applied by the player upon receipt of questions (X,a) and (Z,b) respectively satisfy*

$$\mathrm{E}_{a,b:\, a\cdot b=1} \left\| X(a)Z(b) - (-1)^{a\cdot b} Z(b)X(a) \right\|_\sigma^2 \;=\; O\big(\sqrt{\varepsilon}\big). \tag{3.5}$$

3.3 Application: an n-qubit test

We are ready to put all the pieces together and describe a game for testing n EPR pairs. We call this game the "n-qubit Pauli braiding test".

n-qubit Pauli braiding test: With probability $1/3$ each,

(a) Execute the linearity test;

(b) Execute the anti-commutation test;

(c) Execute the following consistency test: Send one player a label $W \in \{X,Z\}$ uniformly at random, and to the other (W,a) for $a \in \{0,1\}^n$ chosen uniformly at random. Expect answers $c \in \{0,1\}^n$ and $c' \in \{0,1\}$ respectively. Accept if and only if $a \cdot c = c'$.

We sketch why this game indeed tests n qubits. Suppose that a family of observables $W(a)$, for $W \in \{X,Z\}$ and $a \in \{0,1\}^n$, together with projective measurements $\{A_{Xa}\}_{a\in\{0,1\}^n}$ and $\{A_{Zb}\}_{b\in\{0,1\}^n}$ and a state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ specify a symmetric strategy that succeeds with probability at least $1-\varepsilon$ in the game.

Using Lemma 3.3 and Lemma 3.4, success with probability $1-\varepsilon$ implies that conditions (3.1), (3.2) and (3.3) in Corollary 3.1 are all satisfied, up to error $O(\sqrt{\varepsilon})$. (In fact, Lemma 3.4 only implies (3.3) for strings a,b such that $a\cdot b = 1$. The condition for string such that $a\cdot b = 0$ follows from the other conditions.) The conclusion of the corollary is that there exists an isometry V such that the observables $X(a)$ and $Z(b)$ satisfy

$$\mathrm{E}_{a,b} \left\| X(a)Z(b) - V^* g(\sigma_X(a)\sigma_Z(b)) V \right\|_\sigma^2 \;=\; O(\sqrt{\varepsilon})\;,$$

for some non-trivial representation g of $\mathcal{P}_n$. Using what we know of non-trivial representations of this group, it follows that there is an isometry which approximately maps $X(a)$ and (b) to $\sigma_X(a)\otimes \mathrm{Id}$ and $\sigma_Z(b)\otimes \mathrm{Id}$ respectively. Finally, we obtain the n EPR pairs using the consistency relations from part (c) of the test and Lemma 2.7.

3.4 Efficient stability

We start with the following definition, that refines Definition 2.1 to the case where an approximate representation is only defined to approximately respect the group multiplication for a pre-specified set of defining relations.

Definition 3.5. Given a finite presentation $G = \langle x_1, \ldots x_m | R_1, \ldots, R_t \rangle$, an integer $d \geq 1$, $\varepsilon \geq 0$, and a d-dimensional positive semidefinite matrix σ with trace 1, an (ε,σ)-representation of G is a collection $X_1, \ldots, X_m \in U_d(\mathbb{C})$ such that

$$\frac{1}{t}\sum_{j=1}^{t} \left\| R_j(\{X_i\}) - \mathrm{Id} \right\|_\sigma^2 \;\leq\; 2\,\varepsilon\,,^2 \tag{3.6}$$

[2]The factor 2 on the right-hand side is unimportant; we include it for consistency with (2.2).

where $R_j(\{X_i\})$ denotes the relation R_j with each occurrence of a generator x_i replaced by X_i.

In the case where the presentation $G = \langle x_1, \dots x_m | R_1, \dots, R_t \rangle$ is the "exhaustive" presentation, with every element of G included in the generators and all relations of the form $x \cdot y \cdot (xy)^{-1} = 1$, then we recover Definition 2.1. If the representation is "near-exhaustive", e.g. every group element (resp. relation of the form $x \cdot y \cdot (xy)^{-1} = 1$) can be obtained by multiplying at most k generators (resp. chaining k defining relations), then it is not hard to see that an (ε, σ) approximate representation in the sense of Definition 3.5 implies an (ε', σ) approximate representation in the sense of Definition 2.1 for some $\varepsilon' = \mathrm{poly}(k)\varepsilon$. What is perhaps much more surprising is that in some cases we can find very "efficient" presentations such that the dependence on k is far milder. The following is shown in [JNV+20].

Theorem 3.6. *There is a presentation $\mathcal{P}_n = \langle X_1, \dots, X_m | R_1, \dots, R_t \rangle$ such that $m, t = \mathrm{poly}(n)$ and any (ε, σ)-representation in the sense of Definition 2.1 is an (ε', σ)-representation in the sense of Definition 2.1, for some $\varepsilon' = \mathrm{poly}\log(n) \cdot \mathrm{poly}(\varepsilon)$.*

Note that since $|\mathcal{P}_n| = O(4^n)$ but $m = \mathrm{poly}(n)$, there are elements of $|\mathcal{P}_n|$ that cannot be written as a product of less than $\mathrm{poly}(n)$ generators $X_1, \dots, X_m$. Naïvely extending a representation in the sense of Definition 2.1 to a map defined over the full group would lead to an (ε', σ)-representation in the sense of Definition 2.1 where $\varepsilon' = \mathrm{poly}(n) \cdot \varepsilon$. The theorem provides an exponential improvement over this naive method.

Based on this theorem it is possible to define a nonlocal game that provides essentially the same guarantees as the Pauli braiding test, but now the total number of questions in the game is polynomial in n, the number of EPR pairs being tested. (See [JNV+21, Section 7] for a complete description of this game.) This exponential improvement turns out to be crucial for the applications to complexity which we discuss in the next lecture.

Lecture 4

The class MIP*

In this lecture we introduce quantum multi-prover interactive proof systems, discuss the recent characterization MIP* = RE, and examine some consequences.

4.1 Multiprover interactive proofs with entangled provers

We start with the main complexity-theoretic definition. Recall that a *promise language*[1] $L = (L_{yes}, L_{no})$ is specified by a pair of disjoint subsets L_{yes}, L_{no} of $\{0,1\}^*$ and that a *complexity class* is a collection of languages.

Definition 4.1. The class MIP* is the class of promise languages $L = (L_{yes}, L_{no})$ such that there is a classical polynomial-time Turing machine M that on input 1^n returns the description of classical circuits for the verifier V_n in a nonlocal game with *two* quantum players (also called *provers* in this context) A and B such that:

- (Completeness:) There is a family of quantum provers $\{A_n, B_n\}_{n\in\mathbb{N}}$ such that for all $x \in L_{yes}$ the interaction of $V_{|x|}$ and $A_{|x|}, B_{|x|}$ on common input x accepts with probability at least $\frac{2}{3}$.
- (Soundness:) For any family of quantum provers $\{A_n, B_n\}_{n\in\mathbb{N}}$, for all $x \in L_{no}$ the interaction of $V_{|x|}$ and $A_{|x|}, B_{|x|}$ on common input x accepts with probability at most $\frac{1}{3}$.

In general one may allow interaction with more than two provers and more than one round; however the two-prover setting is sufficiently interesting for our purposes, and connects most directly with nonlocal games. (Furthermore, it can be shown that in purely complexity-theoretic terms there is no gain to considering more than 2 provers or more than 1 round.)

The goal in complexity theory is to relate different classes of languages . This is especially interesting when the classes are defined in very different terms, as relations between them can provide insights into different models of computation. A pertinent example is the famous equality IP = PSPACE due to [LFKN92, Sha92]. Among the two classes, PSPACE is the simplest to define: this is the class of all languages that can be decided using a polynomial amount of space, and arbitrary time. A complete problem for PSPACE is the *quantified Boolean formula* (QBF) problem, which is to decide if a formula of the form $\exists x_1 \forall x_2 \exists x_3 \cdots (x_1 \wedge x_2 \wedge \neg x_3) \vee (\cdots)$ is satisfiable. Clearly this can be done in polynomial space

[1] Here the *promise* refers to the fact that it is not required that $L_{yes} \cup L_{no} = \{0,1\}^*$

by trying out all possibilities; it is also possible to show that any problem that is solvable in PSPACE can be reduced to this one, and so we say that QBF is *complete* for PSPACE. The class IP is defined very differently: it is the class of languages L such that membership $x \in L_{yes}$ can be decided efficiently by a randomized polynomial-time verifier interacting with a single infinitely powerful prover (so this is the single-prover analogue of MIP*). While it is not too hard to show that IP $\subseteq$ PSPACE, the converse inclusion is not easy at all — to see why, try coming up with a verification protocol for the QBF problem, and keep in mind that the prover is not to be trusted!

Our goal in this lecture is to characterize the complexity of MIP*. The motivation for doing so it to gain insights about computation and entanglement. And possibly more!

Before we do this let's first review what is known about the classical analogue of MIP*, in which the provers are restricted to classical strategies. This restriction affects both the completeness and soundness requirements in Definition 4.1, and so generally any restriction of the set of allowed strategies for the provers will lead to a different complexity class.

4.1.1 Classical multiprover interactive proof systems

The * in MIP* refers to the fact that provers are allowed to use entanglement. If we omit it we get the class MIP of languages that have classical multiprover interactive proof systems. It was shown by Babai, Fortnow and Lund in the early 1990s that MIP = NEXP. This was shown shortly after the aforementioned result IP = PSPACE, which characterizes the unexpectedly large verification power of single-prover interactive proof systems.

Let's recall how MIP = NEXP is shown. The inclusion of MIP $\subseteq$ NEXP is not hard to obtain. To show it we give a non-deterministic exponential time algorithm that exactly computes the maximum acceptance probability of the verifier in an MIP protocol. This algorithm can therefore, given an instance x and a description of the verifier $V_{|x|}$, determine whether $x \in L_{yes}$ (the maximum success probability is $\geq \frac{2}{3}$) or $x \in L_{no}$ (the maximum success probability is $\leq \frac{1}{3}$), promised that one of them is the case, and thus decide any language $L \in$ MIP; thus MIP $\subseteq$ NEXP follows. To devise such an algorithm first observe that in order to do so it suffices to consider the maximum over deterministic strategies, as for any randomized strategy there is a deterministic one that succeeds with at least the same probability. Now note that a deterministic strategy is specified by a list of answers to each possible question for each of the provers. There are at most exponentially many questions because the bit representation of each question must have polynomial length (since the verifier runs in polynomial time) and similarly for answers. Finally, the success probability of a deterministic strategy can be computed exactly in exponential time simply by executing the verification procedure on each possible tuple of questions, weighted by the probability of the question being asked. Therefore, a non-deterministic algorithm can, in exponential time and space, guess an optimal strategy and compute its success probability.

The reverse inclusion, NEXP $\subseteq$ MIP, is harder. To get a hint of how it is shown, consider the problem of verifying that an exponential-size graph is 3-colorable. Formally, an instance x of this problem is specified by a pair $x = (1^n, C)$ where 1^n denotes an integer n written in unary, and C is the description of a classical circuit $C : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. Any x that does not take this form is neither in L_{yes} nor in L_{no}, and need not be considered any further.[2] The circuit C implicitly specifies a graph $G_x = (V_x, E_x)$ with vertex set $V_x = \{0,1\}^n$ and edge set E_x such that $(i,j) \in E_x$ if and only if $C(i,j) = 1$. Then L_{yes} (resp. L_{no}) is the set of all strings x such that G_x is well-defined and is 3-colorable (resp. not 3-colorable). It is known that the language $L = (L_{yes}, L_{no})$ is complete for NEXP; intuitively this is a "scaled-up" version of the result

[2] As usual, we consider that circuits are represented in some given, fixed manner, e.g. as a list of gates and bits that they act on.

that 3-coloring of n-vertex graphs is NP-complete. Now consider the following description for the actions of the verifier in a candidate multiprover interactive proof system for L:

1. The verifier parses its input x as $x = (1^n, C)$.
2. The verifier selects a pair of vertices (i, j) uniformly at random in $\{0,1\}^n \times \{0,1\}^n$. She sends i to Alice and j to Bob.
3. Alice and Bob each reply with a color $a, b \in \{0,1,2\}$.
4. The verifier accepts if and only if any of the following conditions hold: $C(i,j) = 0$ (there is no edge); $i = j$ and $a = b$ (same color for identical vertices); $C(i,j) = 1$ and $a \neq b$ (different colors for neighboring vertices).[3]

It is clear that this protocol has completeness 1: whenever G_x is 3-colorable there is a winning strategy for the provers. Moreover, a moment's thought will reveal that if G_x is not 3-colorable then there is no perfect winning strategy; hence the maximum probability of success in this case is at most $1 - 2^{-\Omega(n)}$ (because any strategy must fail on at least one question). While this is a separation between the two cases, it is not sufficient to establish soundness, which requires that the maximum probability of success for an $x \in L_{no}$ be at most $\frac{1}{3}$.

What the proof of the inclusion NEXP $\subseteq$ MIP shows is that there is in fact a much better verifier, somewhat more involved than the one that we described here, which is such that whenever the graph is not 3-colorable then the maximum success probability is at most $\frac{1}{3}$. Achieving such a protocol essentially entails finding an efficient method that, informally, maps any graph to another graph of polynomially related size such that graphs that are 3-colorable are mapped to graphs that remain 3-colorable, but graphs that are not 3-colorable are mapped to graphs that are *very far* from 3-colorable. Achieving this can be done using advanced tools from the theory of error-correcting codes; we will not be able to say more in this lecture and refer the interested reader to e.g. [AB09].[4]

4.1.2 Interactive proof systems with entangled provers

Our focus is the class MIP*. What does the characterization MIP = NEXP say about it? Not much! The most important point to realize here is that allowing the provers to use entanglement is a double-edged sword:

First, it can affect the soundness property by allowing the provers to "cheat", meaning achieve a higher success probability. We already saw a good example of this with the Magic Square game. While this game doesn't quite look like the 3-coloring protocol we introduced in the previous section, by transforming it it is possible to come up with explicit instances of the latter that are associated with non-colorable graphs but such that there nevertheless exists a quantum strategy which succeeds with probability 1; see for example [Ji13].

As a result we are unable to transfer the lower bound NEXP $\subseteq$ MIP in a "black-box" manner, and the only trivial lower bound on MIP* is PSPACE, as clearly the verifier can ignore all but one of the provers and execute any classical IP protocol with the remaining prover. In fact it is interesting to note

[3] One may modify this protocol by having the verifier only send pairs (i, j) such that either $i = j$ or (i, j) is an edge, since the other case is an automatic "free ride" for the provers; we gloss over this point here.

[4] Technically such a reduction is not obviously necessary, because the definition of MIP allows more complicated protocols than the 3-coloring game described here. Nevertheless, using appropriate manipulations it is possible to show that any proof of NEXP $\subseteq$ MIP does imply such a reduction.

that such a "collapse" to IP does take place when one allows even more power to the provers, in the form of arbitrary non-signaling strategies as defined in Section **??**. Indeed it is not hard to see that the non-signaling constraints are linear, so that it is possible to write the optimal success probability of non-signaling provers in a multiprover interactive proof system as an exponential-size linear program (LP). Using that linear programs can be solved in time polynomial in their size it can be shown that the class of interactive proof systems with non-signaling provers, denoted MIP^{ns}, lies in EXP. Furthermore, if the number of provers is fixed to 2 and the number of rounds of interaction to 1 then the class "collapses" even further to PSPACE, because the associated LP can be solved more efficiently than a general LP; see [Ito10].

Second, entanglement can also affect the completeness property by increasing the power of the provers in the "honest" case. If we start with a classical protocol for a problem in NEXP this is not so interesting, because we already know that the provers have a good strategy without entanglement — we are not making use of the fact that they can do even better with entanglement, and indeed this fact is a new nuisance that we have to deal with in order to establish the soundness property. But what if we start from a more complex problem, that does not necessarily lie in NEXP, and attempt to design a protocol such that completeness *requires* the use of entanglement?

To see how far one might hope to go in this direction we ought to think about *upper bounds* on MIP^*. Recall from the previous section that for MIP we simply enumerated over all possible strategies. In the quantum setting it is not so direct: since we do not place a priori bounds on the complexity of the provers, it is unclear what dimension one should choose in order to find an optimal strategy. If one was able to show an upper bound on the dimension that is sufficient to approach the optimal success probability (as a function of the size of the protocol) then one would automatically get a corresponding upper bound on the complexity of MIP^*. However, no such bound is known! The only upper bound on MIP^* is the following folklore result:

Lemma 4.2. $\text{MIP}^* \subseteq \text{RE}$*, the set of recursively enumerable languages.*

Proof. Recall that a language $L = (L_{yes}, L_{no})$ is recursively enumerable if there exists a Turing machine such that on input x, if $x \in L_{yes}$ then the Turing machine eventually halts and accepts, whereas if $x \in L_{no}$ then the Turing machine may either halt and reject, or it may never halt.

Consider the Turing machine M that on input x specifying a verifier $V_{|x|}$ searches in increasing dimension and with increasing accuracy for a good strategy in the associated protocol. Since we have not introduced a precise formalism for strategies in MIP^* protocols — we will do so for two-prover one-round protocols in Section **??** — we cannot make this too precise. At present it is sufficient to think intuitively that each prover is specified by a dimension of the Hilbert space on which they act, and for each possible question they may receive, in any round, a POVM on their space that is used to determine an answer; these POVM act on an initial quantum state that lies in the tensor product of the prover's Hilbert spaces. (Any unitary actions the provers may take can be incorporated in the POVMs.) For any given dimension d and accuracy ε the space of strategies in dimension at most d can be discretized to a finite set such that the optimum success probability over elements of that set will be within an additive ε of the optimum over all strategies in dimension at most d.

If $x \in L_{yes}$ by definition there must exist a finite dimension d and a strategy in dimension d that succeeds with probability at least (say) $\frac{2}{3} - \frac{1}{100}$; eventually, taking into account discretization errors M will identify a strategy that succeeds with probability at least $\frac{2}{3} - \frac{2}{100}$ and halt with acceptance, having successfully ruled out the case that $x \in L_{no}$. However, in case $x \in L_{no}$ the Turing machine will never find a strategy with success larger than $\frac{1}{3} + \frac{1}{100}$ (where the $\frac{1}{100}$ accounts for possible discretization errors and can be made arbitrarily small), but it will not be able to rule out the existence of such a strategy either; indeed, for all it knows such a strategy may exist in "just one more dimension". □

For a long time it was unclear where the complexity of MIP* lies, between the two "trivial" extremes of IP and RE. In 2012 Ito and the author showed that NEXP $\subseteq$ MIP* by adapting the proof of NEXP $\subseteq$ MIP by Babai et al. In the past few years better lower bounds were obtained. Quite astonishingly, in 2018 Natarajan and Wright [NW19] showed that NEEXP $\subseteq$ MIP*. One reason that this is "astonishing" is because NEEXP is a strictly (unconditionally) larger class that NEXP, and so their result established unconditionally that the presence of entanglement *increases* the verifier's ability to verify languages, even though the latter's complexity has not changed at all (it remains classical polynomial-time)! Building on this result in 2020 Ji et al. [JNV$^+$21] obtained the following characterization.

Theorem 4.3. MIP* = RE.

A complete problem for the class RE is the *halting problem*: given the description of a Turing machine M as input, does M eventually halt? What Theorem 4.3 shows is that this problem, even though it is *not decidable*, can be efficiently *verified* by asking questions to two provers sharing entanglement. In purely complexity-theoretic terms this is an extremely surprising result in and for itself; note that RE contains *any* bounded time or space complexity class — and much more. The following two lectures will be devoted to a sketch of the main arguments that go in the proof of the theorem; these arguments involve the design of tests for multiple qubits as well as delegation protocols and so we will be on familiar terrain. Aside from the complexity theory it turns out that the characterization MIP* = RE has some interesting consequences in the foundations of quantum mechanics as well as in the theory of operator algebras which we discuss next.

4.2 Consequences

Theorem 4.3 is related to a problem in the foundations of quantum non-locality called *Tsirelson's problem*, itself connected to a problem in the theory of von Neumann's algebra usually referred to as *Connes' Embedding Problem* (CEP). To explain the connection we first make a digression and discuss strategies for computing *upper bounds* on the quantum value $\omega^*(\mathcal{G})$.

4.2.1 Computing upper bounds on $\omega^*(\mathcal{G})$

Let's put Theorem 4.3 aside for a moment and aim instead to contradict it by devising an algorithm that approaches the maximum success probability of quantum provers sharing entanglement in a two-prover one-round interactive proof system with verifier V. Equivalently, suppose that given an explicit game $\mathcal{G}$ we aim to approximate the quantum value $\omega^*(\mathcal{G})$. In the proof of Lemma 4.2 we already saw an algorithm, let's call it algorithm A, that returns an increasing sequence of lower bounds

$$v_1 \leq v_2 \leq \cdots \leq v_k \leq \cdots \leq \omega^*(\mathcal{G})$$

by enumerating strategies in increasing dimension and with increasing level of accuracy. Using the definition (**??**) of the entangled value it is clear that $v_k \to_{k\to\infty} \omega^*(\mathcal{G})$. To make algorithm A into an actual approximation algorithm we need to have a sense of when to stop, e.g. when can we guarantee that $|v_k - \omega^*(\mathcal{G})| \leq \frac{1}{100}$?[5] A natural approach is to construct a companion algorithm B that constructs a decreasing sequence of *upper* bounds

$$w_1 \geq w_2 \geq \cdots \geq w_k \geq \cdots \geq \omega^*(\mathcal{G}) \ .$$

[5]The bound $\frac{1}{100}$ is arbitrary; we want it to be small enough to guarantee that the algorithm can eventually distinguish $\omega^*(\mathcal{G}) \geq \frac{2}{3}$ from $\omega^*(\mathcal{G}) \leq \frac{1}{3}$, so any bound $< \frac{1}{6}$ would do.

Given algorithms A and B consider a third algorithm C that given a game $\mathcal{G}$ as input runs both algorithms in an interleaved fashion, computing v_1, w_1, v_2, w_2, etc., halts whenever $|v_k - w_k| \leq \frac{1}{100}$ and returns "YES" if and only if $\frac{1}{2}(v_k + w_k) > \frac{1}{2}$. Now suppose that both (v_k) and (w_k) converge to $\omega^*(\mathcal{G})$. Then C always terminates. Moreover, if $\omega^*(G) \geq \frac{2}{3}$ then $w_k \geq \frac{2}{3}$ for all k and so the value returned is at least $\frac{1}{2}((\frac{2}{3} - \frac{1}{100}) + \frac{2}{3}) = \frac{2}{3} - \frac{1}{50} > \frac{1}{2}$, whereas if $\omega^*(\mathcal{G}) \leq \frac{1}{3}$ it is at most $\frac{1}{2}(\frac{1}{3} + (\frac{1}{3} + \frac{1}{100})) = \frac{1}{3} + \frac{1}{50} < \frac{1}{2}$. Thus C correctly distinguishes between the two cases.

So how do we determine such a sequence of upper bounds (w_k)? A general approach to finding an upper bound on the optimum of some optimization problem is to consider *relaxations* of the problem, i.e. optimization problems whose optimum is easier to find and is guaranteed to be at least as large as the original optimum. For example, consider the following relaxation

$$\begin{aligned}\omega^*(\mathcal{G}) &= \sup_{|\psi\rangle,\{A_a^x\},\{B_b^y\}} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle\psi| A_a^x \otimes B_b^y |\psi\rangle \\ &\leq \sup_{|u_a^x\rangle,|v_b^y\rangle} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle u_a^x | v_b^y \rangle \ , \end{aligned} \tag{4.1}$$

where the supremum on the second line is over all families of vectors $|u_a^x\rangle, |v_b^y\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ such that for every x, the $\{|u_a^x\rangle\}_a$ are orthogonal and $\sum_a \||u_a^x\rangle\|^2 = 1$; similarly for the $|v_b^y\rangle$. The inequality (4.1) is verified by setting $|u_a^x\rangle = A_a^x \otimes \mathrm{Id}\,|\psi\rangle$ and $|v_b^y\rangle = \mathrm{Id} \otimes B_b^y |\psi\rangle$. So (4.1) is a relaxation, very similar to the one we saw for XOR games in Section 1.3 — except that here, it is provably not tight. What did we gain in the process? Crucially, since the objective function in (4.1) only depends on the inner products between the vectors, without loss of generality we can restrict the vectors to lie in a Hilbert space $\mathcal{H}$ such that $\dim(\mathcal{H}) \leq \min(|\mathcal{X}||\mathcal{A}|, |\mathcal{Y}||\mathcal{B}|)$; this is true even if the original $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ were much larger. This means that by exhaustive search we can find arbitrarily good approximations to the optimum (4.1), without having to go beyond a certain fixed dimension that is determined by the size of the game. In fact, (4.1) is an optimization problem that falls in the class of *semidefinite programs* (informally, linear optimization problems over affine sections of the positive semidefinite cone) and can be solved in time polynomial in its size (as opposed to exponential for exhaustive search).

So the optimum (4.1) *can* be determined efficiently. How useful is it, i.e. how good is the inequality (**??**) $\leq$ (4.1)? Unfortunately, in general there can be an arbitrarily large (multiplicative) gap between the two [JP11], and in particular it can be that $\omega^*(\mathcal{G}) \leq \frac{1}{3}$ but (4.1)$\geq \frac{2}{3}$.[6] The relaxation we have devised is thus too coarse for us to obtain a good algorithm right away. But maybe we can do better? What we did so far consists in adding a vector variable to represent $A_a^x \otimes \mathrm{Id}\,|\psi\rangle$ and $\mathrm{Id} \otimes B_b^y |\psi\rangle$. Each of these can be thought of as a degree-1 monomial in the matrix variables $\{A_a^x, B_b^y\}$, evaluated on $|\psi\rangle$. Considering vectors obtained from higher-degree monomials would allow us to impose more constraints, as for example we could require that

$$\big(\langle\psi| A_a^x \otimes \mathrm{Id}\big) \cdot \big(A_a^x \otimes B_b^y |\psi\rangle\big) \;=\; \big(\langle\psi| \,\mathrm{Id} \otimes \mathrm{Id}\big) \cdot \big(A_a^x \otimes B_b^y |\psi\rangle\big) \ ,$$

due to $\{A_a^x\}_a$ being projective. It is not hard to think of other such constraints. For any integer $k \geq 1$ let's define

$$w_k \;=\; \sup_{\Gamma^{(k)} \geq 0} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \Gamma^{(k)}_{xa,yb} \ , \tag{4.2}$$

where the supremum is taken over all positive semidefinite matrices $\Gamma^{(k)}$ of dimension $\binom{|\mathcal{X}||\mathcal{A}|+|\mathcal{Y}||\mathcal{B}|}{k} \times \binom{|\mathcal{X}||\mathcal{A}|+|\mathcal{Y}||\mathcal{B}|}{k}$. Here we think of the entries of $\Gamma^{(k)}$ as being labeled by sequences $(z_1, c_1), \ldots, (z_k, c_k)$

[6] This fact is not obvious, and constructing such "bad examples" is quite difficult. For some restricted types of games, such as XOR games or unique games, the inequality can be shown to be exact or close to exact respectively.

where $z_i \in \mathcal{X} \cup \mathcal{Y}$ and $c_i \in \mathcal{A} \cup \mathcal{B}$, and $\Gamma^{(k)}$ is the Gram matrix of the associated vectors

$$|u_{(z_i,c_i)_{1\leq i\leq k}}\rangle \;=\; C^{z_k}_{c_k} \cdots C^{z_1}_{c_1}|\psi\rangle \;,$$

where $C^z_c = A^z_c \otimes \mathrm{Id}$ if $z \in \mathcal{X}$ and $c \in \mathcal{A}$, $C^z_c = \mathrm{Id} \otimes B^z_c$ if $z \in \mathcal{Y}$ and $c \in \mathcal{B}$, and $C^z_c = 0$ otherwise. In addition, we add any linear constraint on the entries of Γ that follows from the facts that $\{A^x_a\}$ and $\{B^y_b\}$ are projective measurements for all x, y, and that they act on different tensor factors and hence commute.

With this definition we can verify that $w_1 = (4.1)$; this follows since any positive semidefinite matrix Γ has a factorization as a matrix of inner produces. Moreover, $w_1 \geq w_2 \geq \cdots \geq w_k \geq \omega^*(\mathcal{G})$ since each successive level in the "hierarchy" consists in adding additional variables and constraints. Finally, using standard algorithms for semidefinite programs the optimization problem at the k-th level can be solved in time polynomial in its size, i.e. time $(|\mathcal{X}||\mathcal{A}| + |\mathcal{Y}||\mathcal{B}|)^{O(k)}$. Let's call Algorithm B the algorithm that on input k returns w_k.[7]

4.2.2 The commuting value and Tsirelson's problem

Unfortunately it is not the case that $w_k \to_{k\to\infty} \omega^*(\mathcal{G})$. Indeed, if this were the case algorithm C would return arbitrarily good approximations to $\omega^*(\mathcal{G})$ and thus contradict Theorem 4.3. Nevertheless, since (w_k) is non-increasing and larger than $\omega^*(\mathcal{G})$ the sequence must converge to some value. Interestingly, this value is a natural quantity that is referred to as the *commuting value* of the game and defined as

$$\omega^{com}(G) \;=\; \sup_{|\psi\rangle,\{A^x_a\},\{B^y_b\}} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle\psi|A^x_a B^y_b|\psi\rangle \;, \tag{4.3}$$

where the supremum is taken over all states $|\psi\rangle \in \mathcal{H}$ where $\mathcal{H}$ is a (possibly infinite-dimensional) separable Hilbert space and families of projective measurements $\{A^x_a\}$ and $\{B^y_b\}$ on $\mathcal{H}$ such that for all x, y, a, b, A^x_a and B^y_b commute. Since $A \otimes \mathrm{Id}$ and $\mathrm{Id} \otimes B$ always commute it always holds that $\omega^*(G) \leq \omega^{com}(\mathcal{G})$. The hierarchy of values (w_k) is introduced in [NPA08], where they show the following convergence result.

Lemma 4.4. *For any game $\mathcal{G}$ it holds that* $\lim_{k\to\infty} w_k = \omega^{com}(\mathcal{G})$.

Proof. First note that by definition $\omega^{com}(\mathcal{G}) \leq \lim_{k\to\infty} w_k$, since none of the constraints imposed on the definition (4.2) of w_k makes use of the tensor product structure other than to say that $A^x_a \otimes \mathrm{Id}$ and $\mathrm{Id} \otimes B^y_b$ commute.

The remainder of the proof shows the reverse inequality. For any $k \geq 1$ fix a feasible solution $\Gamma^{(k)}$ to the optimization problem (4.2). The entries of $\Gamma^{(k)}$ are indexed by pairs of monomials m in non-commutative variables $\{A^x_a, B^y_b\}$ of degree at most k. Crucially, the constraints on the optimization problem require that (i) $\Gamma^{(k)} \geq 0$, and (ii) this matrix satisfies $\Gamma^{(k)}_{m_1,m_2} = \Gamma^{(k)}_{n_1,n_2}$ whenever both entries are well-defined and $m_1 m_2^* = n_1 n_2^*$ as monomials in $\{A^x_a, B^y_b\}$, because by definition any such constraint is imposed on the optimization problem.

For any monomial m and integer k at least as large as the degree of m let $\tau_k(m) = \Gamma^{(k)}_{m,1}$. Extend τ_k to a linear form on all non-commutative polynomials by setting $\tau_k(m) = 0$ if m has degree larger than k and extending by linearity. Since $|\tau_k| \leq 1$ for each k (this can be verified because the diagonal entries of $\Gamma^{(k)}$ are all constrained to equal 1, so using (i) all entries of $\Gamma^{(k)}$ must have modulus at most 1) by the Banach-Aleoglu theorem the sequence $(\tau_k)_{k\geq 1}$ admits a pointwise convergent subsequence $(\tau_{k_i})_{k_1\leq k_2\leq\cdots}$; let τ be

[7] Technically we need to allow B to return an approximation to w_k. Since well-behaved semidefinite programs such as (4.2) can be solved in time polynomial in their size and in the logarithm of the desired accuracy we could e.g. require that B returns an additive approximation of w_k that is within error at most 2^{-k}; this will suffice for our purposes.

the pointwise limit. Now crucially we observe that τ is a positive linear form. Indeed, for any polynomial $p = \sum_m \alpha_m m$ where m ranges over monomials we have

$$\begin{aligned}\tau(p^*p) &= \lim_i \ \tau_{k_i}(p^*p) \\ &= \lim_i \ \sum_{m,m'} \alpha_m^* \alpha_{m'} \ \tau_{k_i}(m^*m') \\ &= \lim_i \ \alpha^\dagger \Gamma^{(k_i)} \alpha \\ &\geq 0 \ ,\end{aligned}$$

where for the first line we used linearity of τ_{k_i}, for the second line we used the definition of τ_{k_i} (the equality holds for all i such that $k_i \geq \deg(p)$), for the third line we let $\alpha = (\alpha_m)$ and used property (ii), and for the last we used property (i).

At this point we may conclude in a single abstract step by invoking the GNS construction from C^*-algebra theory: for any positive linear functional τ on a C^*-algebra $\mathcal{A}$ there is a $*$-representation π of $\mathcal{A}$ on a Hilbert space $\mathcal{H}$ and a unit vector $|\xi\rangle \in \mathcal{H}$ such that

$$\forall a \in \mathcal{A} \ , \qquad \tau(a) \ = \ \langle \xi | \pi(a) | \xi \rangle \ . \tag{4.4}$$

For us $\mathcal{A}$ is the algebra of non-commutative polynomials in $\{A_a^x, B_b^y\}$ with complex coefficients satisfying the POVM and commutation conditions, and so the image $\tilde{A}_a^x = \pi(A_a^x)$, $\tilde{B}_b^y = \pi(B_b^y)$, together with the state $|\xi\rangle$, immediately gives us a commuting strategy for G with value $\lim_k w_k$:

$$\begin{aligned}\lim_{k\to\infty} w_k = \lim_{i\to\infty} w_{k_i} &= \lim_{i\to\infty} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \Gamma^{(k_i)}_{xa,yb} \\ &= \lim_{i\to\infty} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \Gamma^{(k_i)}_{(xa,yb)} \\ &= \lim_{i\to\infty} \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \tau_{k_i}((xa,yb)) \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \tau((xa,yb)) \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle \xi | \pi(xa,yb) | \xi \rangle \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle \xi | \pi(xa) \pi(yb) | \xi \rangle \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} V(x,y,a,b) \langle \xi | \tilde{A}_a^x \tilde{B}_b^y | \xi \rangle \ ,\end{aligned}$$

where the first line is by definition of w_{k_i}, the second line by the linear constraints (ii), the third by definition of τ_{k_i}, the fourth by definition of τ, the fitfth by (4.4), the sixth because π is a representation and the last by definition of $\tilde{A}_a^x$ and $\tilde{B}_b^y$.

It is also possible to finish the construction more concretely by defining an infinite-dimensional matrix $\Gamma = \lim_i \Gamma^{(k_i)}$, where for the limit to make sense we embed each $\Gamma^{(k_i)}$ as the top left corner of an infinite-dimensional matrix by padding with zeroes. Since all finite minors of Γ are positive semidefinite, it is positive semidefinite and therefore admits a factorization $\Gamma_{m,m'} = \langle m | m' \rangle$ for some $\{|m\rangle\}$ in a Hilbert space $\mathcal{H}$. We can then define $\tilde{A}_a^x$ as the projection on the span of all $|m\rangle$ such that $m = A_a^x m'$ for some m', i.e. the first variable of monomial m is A_a^x. Using the relations satisfied by the inner products between the

vectors $|m\rangle$ (i.e. condition (ii) above) it is possible to verify that the $\tilde{A}_a^x$ together with analogously defined $\tilde{B}_b^y$ and $|\psi\rangle = |1\rangle$ satisfy the required conditions for a commuting strategy, and that the associated value is once again $\lim_k w_k$. □

The two values $\omega^*(\mathcal{G})$ and $\omega^{com}(\mathcal{G})$ were introduced by Tsirelson in a series of papers laying the foundations for the mathematical study of non-locality [Tsi93]. Rather than using the language of games (which at the time was not much in use yet), Tsirelson directly studied the underlying *correlation sets* defined as

$$C^*(n,k) = \Big\{ \big(\langle \psi, A_a^x \otimes B_b^y \psi\rangle\big)_{a,b,x,y} : \mathcal{H}_\mathsf{A}, \mathcal{H}_\mathsf{B} \text{ Hilbert spaces, } \psi \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}, \|\psi\| = 1, \\ \forall (x,y) \in \{1,\dots,n\}^2, \{A_a^x\}_{a\in\{1,\dots,k\}}, \{B_b^y\}_{b\in\{1,\dots,k\}} \text{ POVM on } \mathcal{H}_\mathsf{A}, \mathcal{H}_\mathsf{B} \text{ resp.} \Big\}, \quad (4.5)$$

$$C^{com}(n,k) = \Big\{ \big(\langle \psi, A_a^x B_b^y \psi\rangle\big)_{a,b,x,y} : \mathcal{H} \text{ Hilbert space, } \psi \in \mathcal{H}, \|\psi\| = 1, \\ \forall (x,y) \in \{1,\dots,n\}^2, \{A_a^x\}_{a\in\{1,\dots,k\}}, \{B_b^y\}_{b\in\{1,\dots,k\}} \text{ PVM on } \mathcal{H} \\ \text{s.t. } [A_a^x, B_b^y] = 0 \ \forall (a,b) \in \{1,\dots,k\}^2 \Big\}.^8 \quad (4.6)$$

By taking direct sums of POVMs and scaled vectors it is not hard to see that both sets are convex subsets of $[0,1]^{n^2k^2}$. Note that in the definition of $C^*(n,k)$ we did not restrict the dimension of $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ to be finite. This is to match Tsirelson's presentation; for our purposes the distinction is not important as it is not hard to see that allowing infinite-dimensional strategies in the definition of the entangled value $\omega^*(G)$ does not change the supremum.[9][10] However, in case the Hilbert spaces in *both* definitions are taken to be finite-dimensional then the two sets can be shown to coincide. (This fact essentially follows from von Neumann's Double Commutant Theorem, though it can also be shown directly; we skip the proof.) In his paper Tsirelson states as "fact" the claim that $C^*(n,k) = C^{com}(n,k)$ for arbitrary separable Hilbert spaces and all $n,k \geq 1$. Having realized that a proof of the claim seemed elusive (with the inclusion $C^*(n,k) \subseteq C^{com}(n,k)$ that we already observed being the only obvious one), in a subsequent note[11] Tsirelson reformulates the "fact" as an open problem and, realizing that the answer may be negative, formulates as an "even more important" problem the question of whether the closure $\overline{C^*(n,k)} = C^{com}(n,k)$. (Here the overline designates closure in the usual topology for $\mathbb{R}^{n^2k^2}$. It is not hard to verify that C^{com} is closed.) Two and a half decades after its introduction Tsirelson's first problem was solved by Slofstra [Slo19], who used techniques from the theory of nonlocal games to show the existence of finite n,k such that $C^*(n,k) \neq C^{com}(n,k)$. Until the proof of Theorem 4.3, an apparently purely complexity-theoretic result, Tsirelson's "even more important problem" remained open. However, we can now observe the following corollary to Theorem 4.3.

Corollary 4.5. *There exists finite $n,k \geq 1$ such that $\overline{C^*(n,k)} \subsetneq C^{com}(n,k)$.*

Proof. Suppose for contradiction that $\overline{C^*(n,k)} = C^{com}(n,k)$ for all $n,k \geq 1$. As an immediate consequence, for any game $\mathcal{G}$ it holds that $\omega^*(\mathcal{G}) = \omega^{com}(\mathcal{G})$. Therefore, algorithm C described in Section 4.2.1 always converges in finite time to a correct answer. This contradicts Theorem 4.3, which implies that the problem "Given a game $\mathcal{G}$, is $\omega^*(\mathcal{G}) \geq \frac{2}{3}$ or $\omega^*(\mathcal{G}) \leq \frac{1}{3}$?" is undecidable. □

[9] To show this, observe that any state $|\psi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$, even in infinite dimensions, always has a Schmidt decomposition $|\psi\rangle = \sum_i \lambda_i |u_i\rangle|u_i\rangle$ such that $\sum_i \lambda_i^2 = 1$. $|\psi\rangle$ can be arbitrarily well approximated in finite dimension by truncating the coefficients; using that the restriction of a POVM to a subspace is a POVM we find arbitrarily good approximations to the game value in finite dimension.

[10] It does change the definition of the set however: as shown in [CS18] some elements of $C^*(n,k)$ cannot be represented in finite dimensions.

[11] "Bell inequalities and operator algebras", available at `https://www.tau.ac.il/~tsirel/download/bellopalg.pdf`.

Note how indirect the proof of Lemma 4.5 is! In particular, while it asserts the existence of n, k there is no obvious way to determine what these integers are, or even upper bounds on them, from the proof. In fact it is possible to tweak the argument to get an explicit construction; we refer to [JNV$^+$21] for more.

4.2.3 Connes Embedding Problem

Quantum mechanics and the theory of operator algebras have been intertwined since their origin. In the 1930s [VN32] von Neumann laid the foundations for the theory of (what are now known as) von Neumann algebras with the explicit goal of establishing Heisenberg's matrix mechanics on a rigorous footing (quoting from the preface, in the translation by Beyer: "The object of this book is to present the new quantum mechanics in a unified representation which, so far as it is possible and useful, is mathematically rigorous"). Following the initial explorations of Murray and von Neumann the new theory took on a life of its own, eventually leading to multiple applications unrelated to quantum mechanics, such as to free probability or noncommutative geometry.

In his 1976 paper completing the classification of injective von Neumann algebras [Con76] Connes made a casual remark that has become a central conjecture in the theory of operator algebras. Since we do not have the mathematical language to express it precisely, I will paraphrase Connes' remark as the comment that "any finite von Neumann algebra *ought to* be well-approximated by finite-dimensional matrix algebras." (In more formal terms, CEP states that every von Neumann algebra type II_1 factor embeds into an ultrapower of the hyperfinite II_1 factor.) Although this conjecture may at first seem rather specific (and in fact as far as I know Connes himself did not pursue the question any further than the remark made in his paper), in the two decades that followed the problem rose to prominence thanks to the work of other mathematicians, such as Kirchberg and Voiculescu, who gave equivalent reformulations of the conjecture in operator algebras and free probability. (See e.g. [Cap15] for more reformulations.) Kirchberg's formulation is closest to us: Kirchberg showed that CEP is equivalent to the *QWEP conjecture* about the equivalence of the minimal and maximal tensor products on the full group C^* algebra of a nonabelian free group [Kir93].[12] Informally, the minimal and maximal tensor products of two $C*$ algebras provide two ways to define the closure of the algebraic tensor product, with respect to two different norms, the minimal norm

$$\|x\|_{min} = \sup_{\pi_A, \pi_B} \left\|\pi_A \otimes \pi_B(x)\right\|$$

where the supremum ranges over all pairs of representations $\pi_A : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H}_\mathsf{A})$ and $\pi_B : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H}_\mathsf{B})$, whereas

$$\|x\|_{max} = \sup_{\pi} \|\pi(x)\|$$

where here $\pi : C^*(\mathbb{F}_2) \otimes C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H})$ is any representation that is such that $\pi(a \otimes b) = \pi_A(a)\pi_B(b)$ where $\pi_A, \pi_B : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H})$ are representations with commuting range. Clearly, $\|x\|_{min} \leq \|x\|_{max}$ always, and these two norms can be seen to be the smallest and largest "reasonable" norms that one may put on the tensor product of two C^*-algebras.

With this reformulation it may not be surprising that Kirchberg's QWEP is directly related to Tsirelson's problem, and indeed building on work of Fritz [Fri12] and Junge et al. [JNP$^+$11] Ozawa [Oza13a] showed that Tsirelson's "even more important" problem is equivalent to CEP. This brings us to a second corollary of Theorem 4.3.

[11] The brief discussion in this section is adapted from [Vid19].

[12] Concretely, a C^* algebra can always be represented as a sub-algebra of the algebra of bounded linear operators on a Hilbert space that is closed under taking adjoints, and closed under the norm topology. A von Neumann algebra is further restricted to be closed under the weak toperator topology.

Corollary 4.6. *CEP has a negative answer, i.e. there exists a von Neumann algebra that is not hyperfinite.*

For more background on the relation between Tsirelson's problem and Kirchberg's conjecture, presented in an accessible way, I recommend [Fri12]. For additional results and the connection to CEP, presented in a less accessible way, I recommend [Oza13b].

4.3 An overview of the proof of RE $\subseteq$ MIP*

4.3.1 A cartoon version

At the highest level our proof strategy is as follows. Recall from the previous lecture that for the case of classical protocols one can show the inclusion NEXP $\subseteq$ MIP. While by itself this is already non-trivial, let's take as our starting point the assumption that we are able to show an analogous inclusion for quantum interactive proofs, i.e. NEXP $\subseteq$ MIP$^*(2,1)$.[13] Observe that this inclusion can be recast as a form of delegation. Paraphrasing, the inclusion states that any language $L \in$ NEXP has a multiprover interactive proof systems with quantum provers. Now for any Turing Machine M the language L_M that consists of all n such that there is a string $y \in \{0,1\}^*$ such that M accepts (n,a) in time at most $\exp(n)$ lies in NEXP;[14] moreover for some choices of M it is NEXP-complete.[15] Thus there is an efficient reduction from Turing machines M to verifiers V such that for all integer n, on input z such that $|z| = n$,

$$\exists a \ : \quad M \text{ accepts } (z,a) \text{ in time} \leq \exp(n) \tag{4.7}$$

then $\omega^*(V_n(z)) \geq \frac{2}{3}$, and if no such a exists then $\omega^*(V_n(z)) \leq \frac{1}{3}$. Now suppose that we're able to achieve a somewhat stronger reduction, where for the starting point we replace the condition (4.7) by

$$\textit{On average over } x \sim \mathcal{U}_N\,, \quad \Pr_x\Big(\exists a \ : \ M \text{ accepts } (z,x,a) \text{ in time} \leq \exp(n)\Big) \geq \frac{2}{3}\,, \tag{4.8}$$

where $N = 2^n$ and for every n, $\mathcal{U}_N$ is the uniform distribution on $\{0,1\}^N$. (Suppose also that a symmetric condition holds for soundness.) This would be a form of delegation for (exponential-time) AM ("Arthur-Merlin") protocols, where an AM protocol is one in which the verifier can send a uniformly random string as question to the prover before receiving the proof. Note that the step we just made is highly non-trivial because of the introduction of a distribution on x; delegating randomized computations like this is hard because there is no easy means to verify that the computation is being performed with the "right choice" of the random string x—indeed, we need to make sure to detect cases where it might be that there exists (x,a) such that M accepts (n,x,a), but it is still very unlikely to be the case when x is chosen at random. As we will see later the use of quantum provers and entanglement will be useful to achieve this.

Let's do one last leap of faith and suppose that we have an even stronger reduction, that applies directly to exponential-size multiprover interactive proofs. Precisely, we'd replace the condition (4.8) by

$$\textit{On avg over } (x,y) \sim \mathcal{U}_N \times \mathcal{U}_N\,, \ \Pr_{(x,y)}\Big(\exists (a,b) \ : \ M \text{ accepts } (z,x,y,a,b) \text{ in time} \leq \exp(n)\Big) \geq \frac{2}{3}\,, \tag{4.9}$$

[13] This inclusion is shown in [IV12] for 5 provers. The 2-prover version follows from the work in [JNV$^+$20].

[14] This formulation is a bit unusual due to the use of the letter n to represent the input, which is usually called x; this is for later convenience. Here n is written in binary. Note that the time bound implies that without loss of generality $|y| \leq \exp(n)$.

[15] An example would be to take M the Turing machine that parses n as an implicitly represented graph $n = (1^{n'}, C)$ and expects y to be an explicit coloring for the $2^{n'}$ vertices of the graph; see Section 4.1.1.

where in addition we'd require that (a, b) are generated locally by quantum provers sharing entanglement, such that the provers are given x and y respectively; formally, given (x, y) the pair (a, b) should be distributed as $\langle\psi|A_a^x \otimes B_b^y|\psi\rangle$ for some state $|\psi\rangle$ (independent of (x, y)) and POVM $\{A_a^x\}$ and $\{B_b^y\}$. Once again we'd also require a symmetric condition with probabilities $\leq \frac{1}{3}$ for soundness.

Let's call the resulting reduction a "compression" procedure: it takes as input an exponential-time verifier V and returns a polynomial-time verifier V^{COMPR} that has the same completeness and soundness properties: if there is a good strategy for V_n there is also one for $V\text{COMPR}_n$ and vice-versa. Then I claim that by iterating this compression procedure we could obtain progressively stronger inclusions, from EXP $\subseteq$ MIP* to EEXP $\subseteq$ MIP* to .... any time complexity that is a finite tower of exponentials.[16] Recall that for well-chosen M, the problem of given an integer n, does M halt in at most 2^n steps is EXP-complete. Now suppose that e.g. we have a family of verifiers $\{V_n\}$, implicitly depending on M, such that $\omega^*(V_n) \geq \frac{2}{3}$ if M halts in $\leq 2^n$ steps, and $\omega^*(V_n) \leq \frac{1}{3}$ otherwise; such a family follows from EXP $\subseteq$ MIP$^*(2, 1)$. Now define $\{V_n^{\text{COMPR}}\} = \text{COMPR}(\{V_{2^n}\})$. Then by definition $\omega^*(V_n^{\text{COMPR}}) \geq \frac{2}{3}$ if $\omega^*(V_{2^n}) \geq \frac{2}{3}$ if M halts in $\leq 2^{2^n}$ steps, and similarly $\omega^*(V_n^{\text{COMPR}}) \leq \frac{1}{3}$ otherwise. Thus EEXP $\subseteq$ MIP*. Iterating this procedure and stretching things a little bit, this would give us the inclusion $\mathsf{TIME}(T(n)) \subseteq$ MIP* for any computable function T. And then taking the "limit", we'd get RE $\subseteq$ MIP*...?

Obviously there's a lot of moving pieces in this description. The goal in this lecture is to make them sufficiently precise as to be believable, and eventually arrive at a core "nugget" that encapsulates the key step that needs to be proven—which we'll do in the next lecture. For now we focus on, first, setting things up so that the above sketch can be made more precise, and second, discussing in more detail the "compression" procedure, which is the key part where the use of quantum provers is essential.

4.3.2 The Halting problem

Recall from the last lecture that the two main consequences of RE $\subseteq$ MIP* that we discussed, negative answers to Tsirelson's problem and to Connes' Embedding Conjecture, both follow from the fact that the problem "Given a two-player one-round game G such that $\omega^*(G) \geq \frac{2}{3}$ or $\omega^*(G) \leq \frac{1}{3}$, which is the case?" is undecidable. In this lecture we will show that there is a computable map $\mathcal{F}$ from Turing Machines M to games $G = G_M$ such that if M halts then $\omega^*(G) \geq \frac{2}{3}$, whereas if M does not halt then $\omega^*(G) \leq \frac{1}{3}$. To argue that this indeed shows that the aforementioned problem is undecidable, we recall the proof that the Halting problem is undecidable.

Definition 4.7. The language L_{HALT} is the set of all $x \in \{0, 1\}^*$ such that x is the description of a Turing Machine M such that M, when it is executed on an empty tape, eventually halts.

For convenience we use the notation $\overline{M} \in \{0, 1\}^*$ to denote the description of a Turing Machine M, using some canonical representation. Recall that there exists a "universal" Turing machine $\mathcal{U}$ that on input $\overline{M}$ and x simulates the execution of M on input x.

Lemma 4.8. *The language L_{HALT} is undecidable.*

Proof. Suppose for contradiction that there exists a Turing Machine A such that given as input $\overline{M}$, A halts with "YES" in case M halts on the empty tape, and A halts with "NO" otherwise. Now consider the following Turing Machine B. When run on an empty tape, B first executes A on $\overline{B}$. If A halts with "YES" then B enters an infinite loop. If A halts with "NO" then B halts with "YES". Does B halt? We have reached a contradiction, therefore A does not exist. □

[16] We could do the same argument for non-deterministic time complexities, but it is easier to present in the deterministic case. The place where we do need non-determinism is for the compression procedure.

Note that in the proof of Lemma 4.8 we designed a Turing Machine B that at some point performs an instruction that depends on its own "source code" $\overline{B}$. That this is allowed is a consequence of Kleene's recursion theorem, which is basically a generalization of the standard diagonalization argument. We will use this possibility again later.

4.3.3 Compression

We make more precise what we need of the magical "compression procedure" discussed in Section 4.3.1. First we introduce a restricted class of verifiers.

Definition 4.9. A *normal form verifier* is a Turing Machine V that on input n returns the description of a Turing Machine R_n (the "referee," or "decision procedure") that on input $(x, y, a, b) \in \{0,1\}^{4n}$ returns a value $d \in \{0,1\}$. To R_n we associate a two-player one-round game G_n whose question and answer sets are $X = Y = A = B = \{0,1\}^n$ and such that the question distribution π is uniform on $\{0,1\}^n \times \{0,1\}^n$ and the referee predicate is given by R_n.

We let $\mathsf{TIME}_V(n)$ be the worst-case running time of R_n over all inputs (x, y, a, b). If $\mathsf{TIME}_V(n) \leq (\lambda n)^\lambda$ for some integer $\lambda \geq 1$ and all $n \geq 1$ then we say that V is *λ-bounded*.

Note that in the definition we fixed the question distribution used for the game G_n to the uniform distribution. At this stage this is mostly for convenience. Later we will realize that this is too restrictive, and so one should bear in mind that the definition can be generalized to allow various classes of distributions, where the key point is that the distribution should be fixed and independent of V.

We need one last definition.

Definition 4.10. For a nonlocal game $\mathcal{G}$ and a probability $p \in [0,1]$ let $\mathcal{E}(\mathcal{G}, p)$ denote the smallest integer $d \geq 1$ such that there exists a strategy $(|\psi\rangle, \{A_a^x\}, \{B_b^y\})$ for the players in $\mathcal{G}$ that has success probability at least p and such that $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$. If no such strategy exists, then $\mathcal{E}(\mathcal{G}, p) = \infty$.

For example, for the Magic Square game it is possible to show that $\mathcal{E}(\mathcal{G}, 1) = 4$ (two qubits per player), and in fact there is a $c < 1$ such that $\mathcal{E}(\mathcal{G}, p) = 4$ for all $p \in [c, 1]$. For a game that has a perfect classical strategy we have $\mathcal{E}(\mathcal{G}, p) = 1$ for all $p \in [0,1]$.

Let's make the following specification for a compression procedure.

Claim 4.11. *There is a polynomial-time computable mapping* COMPR *that takes as input a Turing machine description $\overline{V}$ and an integer λ written in unary and returns a Turing machine description $\overline{V^{\text{COMPR}}} = \text{COMPR}(\overline{V}, \lambda)$ such that the following conditions hold:*

(a) V^{COMPR} is always a normal form verifier such that $\mathsf{TIME}_{V^{\text{COMPR}}}(n) \leq p_{\text{COMPR}}(\lambda + n)$, for some universal polynomial p_{COMPR} independent of V.

(b) If V is a normal form λ-bounded verifier then for every $n \geq 1$ letting $N = 2^n$ the following hold:

(b.i) If $\omega^*(R_N) = 1$ then $\omega^*(R_n^{\text{COMPR}}) = 1$.

(b.ii) $\mathcal{E}(G_n^{\text{COMPR}}, \frac{1}{2}) \geq \max\left\{\mathcal{E}(G_N, \frac{1}{2}), N\right\}$.

The key point about Claim 4.11 is that the running time of R_n^{COMPR} can be much smaller than that of R_N, yet it preserves essential properties of it, stated in (b.i) and (b.ii).

We make a few comments on the requirements stated in the claim. First of all, even though eventually we only need to create a *computable* mapping $\mathcal{F}$ from Turing Machines to games, it will be important that

here COMPR is required to run in polynomial time. Second, it will also be essential that for any input $(\overline{V}, \lambda)$ to COMPR the output $\overline{V^{\text{COMPR}}}$ is the description of a time-bounded verifier. Note that this is not hard to enforce in practice by hard-coding some kind of time-out mechanism in the definition of V^{COMPR}. Finally, observe that condition (b.ii) states something a little stronger (strictly speaking, incomparable) than the "soundness preservation" condition we considered in Section 4.3.1. Indeed the fact that we are able to make a statement about entanglement will play an important role in the final argument. (On the other hand, that the conditions apply to $N = 2^n$ as opposed to e.g. $N = n + 1$ is not important; since it is what comes out of the proof we keep it here—what matters is that V_n^{COMPR} reproduces properties of V_N for some $N > n$ while having complexity comparable to V_n, not V_N.) Finally, note that due to condition (b) running COMPR on a trivial input that always accepts already yields an interesting family of games: due to (b.i) we will have $\omega^*(V_n^{\text{COMPR}}) = 1$ for all n, and due to (b.ii) achieving any value larger than $\frac{1}{2}$ will necessarily require a quantum state of local dimension at least $N = 2^n$.

These observations show that designing a procedure COMPR that fulfills all conditions will likely not be an easy task. Nevertheless, let's put that task aside for the time being and see how the desired reduction can be completed assuming the validity of Claim 4.11.

4.3.4 A self-referential verifier

Fix a Turing Machine M and an integer $\lambda \geq 1$ and consider the following normal form verifier $V = V_{M,\lambda}$, that implicitly depends on M and λ (and in fact is efficiently computable from the pair $(\overline{M}, \lambda)$). For any $n \geq 1$, we describe the decision procedure R_n using high-level "pseudocode". Given $(x, y, a, b) \in \{0,1\}^{4n}$, R_n does the following:

1. R_n simulates M on the empty tape for n steps. If M halts then R_n accepts (i.e. it returns the value '1', irrespective of its inputs (x, y, a, b). Otherwise, if M has not halted in n steps then R_n proceeds to the next item.
2. R_n computes $\overline{V^{\text{COMPR}}} = \text{COMPR}(\overline{V}, \lambda)$.
3. R_n returns the decision $(R^{\text{COMPR}})_n(x, y, a, b) \in \{0,1\}$.

Note that in giving this high-level description of a Turing Machine V, that on input 1^n returns the description $\overline{R_n}$, we have referred to the description $\overline{V}$ itself. That this is possible, i.e. V is a well-defined Turing Machine, is a consequence of Kleene's recursion theorem—this is similar to the self-referential call we made for the definition of algorithm B in the proof of Lemma 4.8.

The following three claims establish the key properties of this construction.

Claim 4.12. *For any Turing Machine M there is an integer $\lambda \geq 1$ which is computable from $|M|$ and such that V is λ-bounded.*

Proof. By definition V on input 1^n returns a decision procedure R_n that takes four inputs of length n each, so it is a normal form verifier. It remains to estimate its running time. First we estimate $|\overline{V}|$. Clearly, the actions to be performed in each of the three steps can be described using $\text{poly}(|\overline{M}|, \lambda)$ bits. Note that the description $\overline{\text{COMPR}}$ does not depend on anything, so its size is a constant.

Next we estimate the running time of R_n. The first step, the simulation of M for n steps, takes time $p_1(n, \overline{M})$ for some universal polynomial p_1. The second step, the computation of $\overline{V^{\text{COMPR}}}$, takes time $p_2(\overline{V}, \lambda)$, for some universal polynomial p_2 that bounds the running time of COMPR. The last step, the evaluation of $R_n^{\text{COMPR}}(x, y, a, b)$, takes time $p_{\text{COMPR}}(\lambda + n)$ by property (a) in Claim 4.11.

Overall the running time is $\mathrm{poly}(n, \overline{M}, \lambda)$ for some universal polynomial. This can be bounded above by the expression $(\lambda n)^\lambda$ for all $n \geq 1$ provided λ is large enough compared to $\overline{M}$. □

For the remaining two claims we fix λ to the value promised in Claim 4.12 and let $\{R_n\}$ and $\{G_n\}$ be the family of decision procedures and games respectively implied by the verifier V specified from M and λ.

Claim 4.13. *Suppose that M halts on an empty input tape. Then $\omega^*(R_n) = 1$ for all n.*

Proof. Let T be the number of steps taken by M to halt. Then for all $n \geq T$ the decision procedure R_n always accepts its inputs at step 1. Therefore $\omega^*(R_n) = 1$ for all $n \geq T$. Now we show by (strong) downwards induction from $m = T$ to 1 that $\omega^*(R_m) = 1$. We showed the induction hypothesis for $m = T$ already. Suppose it true up to some value $m > 1$. Then since M does not halt in $(m-1)$ steps, the decision procedure R_{m-1} proceeds to step 2. and executes $(R^{\text{COMPR}})_{m-1}$. Since $2^{m-1} > m-1$, it follows from the induction hypothesis that $\omega^*(R_{2^{m-1}}) = 1$. Using property (b.i) in Claim 4.11 we have that $\omega^*(R_{m-1}) = 1$, as desired. □

Claim 4.14. *Suppose that M does not halt. Then $\omega^*(R_n) \leq \frac{1}{2}$ for all $n \geq 1$.*

Proof. We show that $\mathcal{E}(G_n, \frac{1}{2}) = \infty$ for all $n \geq 1$. This shows that no finite strategy can achieve a success probability larger than $\frac{1}{2}$, and taking the limit that $\omega^*(R_n) \leq \frac{1}{2}$, as desired. Since M does not halt, for any n, R_n proceeds to step 2. and returns the desision of R_N^{COMPR} where $N = 2^n$. By property (b.ii) in Claim 4.11 it follows that for all $n \geq 1$,

$$\mathcal{E}(G_n, \frac{1}{2}) \geq \max\left\{\mathcal{E}\big(G_N, \frac{1}{2}\big), N\right\}.$$

By straightforward induction, $\mathcal{E}(G_n, \frac{1}{2}) \geq T$ for any integer T, so it must be ∞. □

4.3.5 A game for the halting problem

We now describe the reduction $\mathcal{F}$. On input $\overline{M}$, $\mathcal{F}$ first computes the integer λ whose existence is promised in Claim 4.12. Then $\mathcal{F}$ computes a description of the decision procedure R_1 from the start of Section 4.3.4, based on $\overline{M}$ and λ. Finally, $\mathcal{F}$ returns a description of the associated game $G = G_1$.[17]

Suppose first that M halts. Then by Claim 4.13 it holds that $\omega^*(G) = 1$. Suppose now that M does not halt. It follows from Claim 4.14 that $\omega^*(G) \leq \frac{1}{2}$. This completes the reduction.[18]

Remark 4.15. It is worth pausing to appreciate the significance of this reduction. Beyond the stated inclusion of complexity classes, it makes quite a striking statement about the complexity that may lurk behind simple, finite, observable phenomena in quantum mechanics. What the existence of $\mathcal{F}$ states is that for *any* problem that can be encoded in the halting of a Turing machine there is a game, that moreover is easily computable from the Turing machine, that "witnesses" this fact. Consider for example the Riemann Hypothesis (RH). There is a simple Turing Machine M that halts if and only if RH is provable in ZFC (Zermelo-Fraenkel set theory with the axiom of choice included). Indeed M simply enumerates over all possible proofs in ZFC and checks if they are (i) valid proofs and (ii) prove RH. Moreover, the Turing machine M is large, but not absurdly so; probably a few millions of characters are more than enough. This means that we can, in

[17] It is interesting that ultimately we only need G_1, but to arrive at constructing it we had to consider infinite families of games.

[18] While we promised to obtain a separation between values $\frac{2}{3}$ and $\frac{1}{3}$, we only obtained one between 1 and $\frac{1}{2}$. There is nothing in the line of argument that is special about $\frac{1}{2}$ and we could have done the same replacing it by $\frac{1}{3}$, giving us an even stronger reduction than desired. In general, the values $\frac{2}{3}$ and $\frac{1}{3}$ can be amplified towards 1 and 0 respectively by applying techniques from parallel repetition; see e.g. [Yue16].

principle but also in practice, write a simple computer program that will return the rules for a moderately-sized nonlocal game $\mathcal{G} = \mathcal{G}_{RH}$ such that $\omega^*(\mathcal{G}) = 1$ if and only if RH is provable in ZFC. Isn't this amazing?

Remark 4.16. A natural question is what is the *commuting value* ω_{com} of the games $\mathcal{G} = \mathcal{G}_{M,\lambda}$. Naturally this value is always at least $\omega^*(M)$, and moreover for *some* infinite family of M it must be the case that $\omega_{com}(\mathcal{G}) = 1$ even if M does not halt (and hence $\omega^*(\mathcal{G}) \leq \frac{1}{2}$), as otherwise using algorithm C from the previous lecture we would be able to solve the Halting problem. We do not know if $\omega_{com}(\mathcal{G}) = 1$ for all games $\mathcal{G}$ in the range of the reduction $\mathcal{F}$.

Bibliography

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of computer and system sciences*, 47(3):549–595, 1993.

[BOGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 113–131. ACM, 1988.

[Cap15] Valerio Capraro. *Connes' Embedding Conjecture*, pages 73–107. Springer International Publishing, Cham, 2015.

[CLS17] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1):012202, 2017.

[CM14] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming*, pages 320–331. Springer, 2014.

[CN16] Matthew Coudron and Anand Natarajan. The parallel-repeated magic square game is rigid. *arXiv preprint arXiv:1609.06306*, 2016.

[Con76] Alain Connes. Classification of injective factors cases II_1, II_∞, III_λ, $\lambda \neq 1$. *Annals of Mathematics*, pages 73–115, 1976.

[CS17] Andrea Coladangelo and Jalex Stark. Robust self-testing for linear constraint system games. *arXiv preprint arXiv:1709.09267*, 2017.

[CS18] Andrea Coladangelo and Jalex Stark. Unconditional separation of finite and infinite-dimensional quantum correlations. *arXiv preprint arXiv:1804.05116*, 2018.

[DCOT17] Marcus De Chiffre, Narutaka Ozawa, and Andreas Thom. Operator algebraic approach to inverse and stability theorems for amenable groups. *arXiv preprint arXiv:1706.04544*, 2017.

[Din07] Irit Dinur. The pcp theorem by gap amplification. *Journal of the ACM (JACM)*, 54(3):12, 2007.

[Fri12] Tobias Fritz. Tsirelson's problem and Kirchberg's conjecture. *Reviews in Mathematical Physics*, 24(05):1250012, 2012.

[GH15] William Timothy Gowers and Omid Hatami. Inverse and stability theorems for approximate representations of finite groups. *arXiv preprint arXiv:1510.04085*, 2015.

[IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 217–228. IEEE, 2009.

[Ito10] Tsuyoshi Ito. Polynomial-space approximation of no-signaling provers. *Automata, Languages and Programming*, pages 140–151, 2010.

[IV12] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for nexp sound against entangled provers. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 243–252. IEEE, 2012.

[Ji13] Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint arXiv:1310.3794*, 2013.

[JNP+11] Marius Junge, Miguel Navascues, Carlos Palazuelos, David Perez-Garcia, Volkher B Scholz, and Reinhard F Werner. Connes' embedding problem and Tsirelson's problem. *Journal of Mathematical Physics*, 52(1):012102, 2011.

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test. *arXiv preprint arXiv:2009.12982*, 2020.

[JNV+21] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. *Communications of the ACM*, 64(11):131–138, 2021.

[JP11] Marius Junge and Carlos Palazuelos. Large violation of bell inequalities with low entanglement. *Communications in Mathematical Physics*, 306(3):695, 2011.

[Kaz82] David Kazhdan. On ε-representations. *Israel Journal of Mathematics*, 43(4):315–323, 1982.

[Kir93] Eberhard Kirchberg. On non-semisplit extensions, tensor products and exactness of group C^*-algebras. *Inventiones mathematicae*, 112(1):449–489, 1993.

[KKM+11] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate. *SIAM Journal on Computing*, 40(3):848–877, 2011.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM (JACM)*, 39(4):859–868, 1992.

[MYS12] Matthew McKague, Tzyh Haur Yang, and Valerio Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(45):455304, 2012.

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10(7):073013, 2008.

[NW19] Anand Natarajan and John Wright. Neexp is contained in mip. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518. IEEE, 2019.

[Oza13a] Narutaka Ozawa. About the Connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[Oza13b] Narutaka Ozawa. About the connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[Sha92] Adi Shamir. IP= PSPACE. *Journal of the ACM (JACM)*, 39(4):869–877, 1992.

[Slo19] William Slofstra. The set of quantum correlations is not closed. In *Forum of Mathematics, Pi*, volume 7. Cambridge University Press, 2019.

[SW88] Stephen J Summers and Reinhard Werner. Maximal violation of bell's inequalities for algebras of observables in tangent spacetime regions. In *Annales de l'Institut Henri Poincare Physique Theorique*, volume 49, pages 215–243, 1988.

[Tsi93] Boris S Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement*, 8(4):329–345, 1993.

[Vid19] Thomas Vidick. From operator algebras to complexity theory and back. *Notices of the American Mathematical Society*, 66(10), 2019.

[VN32] J Von Neumann. Mathematische grundlagen der quantenmechanik. 1932.

[WBMS16] Xingyao Wu, Jean-Daniel Bancal, Matthew McKague, and Valerio Scarani. Device-independent parallel self-testing of two singlets. *Physical Review A*, 93(6):062121, 2016.

[Yue16] Henry Yuen. A parallel repetition theorem for all entangled games. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.