UCSD Summer school notes

Quantum multiplayer games, testing and rigidity

1 Multiplayer games

A multiplayer game is a single-round interaction between a *referee* and multiple *players*. The game specifies the actions of the referee: with what distribution she selects the questions to the players, and what tuples of answers are valid for each tuple of questions. The players, traditionally referred to as Alice, Bob, Charlie, etc., are always assumed to attempt to maximize their probability of winning (i.e. providing valid answers) in the game. The probability is over both the referee's and the players' randomness. This quantity, the players' maximum winning probability, is usually called the *value* of the game.

Multiplayer games play an important role in theoretical computer science. Their study ca be motivated from at least two vantage points:

- *Hardness of approximation for constraint satisfaction problems.* The most famous game in this context is the 3SAT "clause-vs-variable" game. In this game there are two players, Alice and Bob. Both players and the referee are given access to the same 3SAT formula φ. Moreover, the players can agree on a strategy after having seen the formula and before the game starts. Once the game starts, the referee selects a clause $C = x \vee y \vee z$ (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. It is not hard to verify that the maximum success probability in this game is directly related to the largest number of clauses of φ that can be simultaneously satisfied by any assignment. Since 3SAT is NP-hard, it follows that the value of a game specified explicitly (in matrix form, i.e. a table specifying explicitly the distribution on questions and the truth table for valid answer tuples) is NP-hard to compute.
  In fact, it follows from the PCP theorem that the value 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.

- *Interactive protocols in cryptography.* In cryptography, games often 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.

Quantum information introduces an exciting twist in the theory of multiplayer games: what if the players are allowed to use entanglement? The game is the same, but the set of allowed strategies has been broadened. While entanglement does not allow the players to communicate, it could in principle allow them to increase their odds of winning, and indeed this is the case: you already saw this in the example of the CHSH game, and we will see many more examples as we go.

Remark 1. One can ask, why stop at quantum? The most general strategies that respect the no-communication assumption are aptly called *non-signaling strategies*. We will not discuss them in this lecture, but aside from being a nice extension of quantum strategies they have recently become very relevant in designing efficient (single-prover) classical delegation protocols; see e.g. [KRR14].

Interestingly, many games have the property that the optimal quantum strategy for the players is essentially unique. This property, called *rigidity*, can be leveraged to devise classical tests that verify that arbitrary quantum devices (the players) perform very 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 a subsequent lecture.

Resources. A great introduction to complexity-theoretic questions about non-local games is the paper by Cleve et al. [CHTW04]. The lecture notes by Ji (see module 5 here) cover CHSH, the Magic Square game, linearity testing (using a slightly different analysis than the one we will give here), and graph-based games.

Notation. In this lecture we use the non-standard notation

$$|\phi_d\rangle \;=\; \frac{1}{\sqrt{d}}\sum_{i=1}^{d}|i\rangle|i\rangle$$

for the maximally entangled state in d dimension. In particular, $|\phi_2\rangle$ denotes an EPR pair.

2 XOR games

The simplest kind of games to think about are XOR games. An XOR game has the following form:

1. The referee selects a pair of questions $(i,j) \in \{1,\dots,m\} \times \{1,\dots,n\}$ 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 payoff to the players is $a_i b_j c_{ij}$, where $c_{ij} \in \{\pm 1\}$. (Note that the payoff is a number in $\{\pm 1\}$. The interpretation is that, if the payoff is $+1$ the players "win", and if it is -1 they "loose". More general real-valued payoffs may be considered.)

Given questions i and j respectively, the goal of the players is to provide answers whose product equals the target value c_{ij}. However, Alice only knows i and Bob only knows j, which is what can make the game challenging. The following example introduces the famous CHSH game as an XOR game.

2.1 Classical strategies

Given an arbitrary XOR game G, introduce a matrix $A \in \mathbb{R}^{n\times m}$ with coefficients $G_{ij} = \pi(i,j)c_{ij}$. Then the maximum expected payoff for the players is

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

It is not hard to verify that this is related to the value $\omega(G)$ of the game, the maximum propbability for the players to be accepted, as $\omega(G) = \frac{1}{2} + \frac{1}{2}\beta(G)$.

Conversely, for any $A \in \mathbb{R}^{m\times n}$, we can define $c_{ij} = \mathrm{sign}(A_{ij})$ and $\pi(i,j) = \frac{A_{ij}}{\sum_{k,l}|A_{kl}|}$ to transform any optimization problem of the form (1) into an XOR game. You may already know that computing (1), or even approximating it to within a small constant factor, is NP-hard in the worst case; if not, verify that there is a simple reduction from MAXCUT. So approximating the value of a 2-player XOR game is NP-hard; this is a slightly stronger result than the hardness for "clause-vs-variable"-type games that we discussed earlier.

Remark 2. In general one may allow the players to use randomized strategies, including both private and shared randomness, to select their answers. It is easy to see that this cannot help in general: if on average (over their random coin tosses) the players achieve a certain payoff, then there must exist some choice of coins that lets them achieve at least the average payoff, and they might as well fix this choice of coins as part of their strategy.

Example 3. We let $m = n = 2$ and $c_{11} = c_{12} = c_{21} = 1$, $c_{22} = -1$. Thus the players "win" if and only if $a_i \oplus b_j = i \wedge j$ (interpreting $i, j \in \{1,2\}$ as Booleans by mapping 1 to 0 and 2 to 1). This game is called the CHSH game, after its inventors, Clause, Horne, Shimony, and Holt. We can evaluate the maximum expected payoff exactly:

$$\beta(G) = \max_{a_i,b_j\in\{\pm 1\}} \sum_{i,j} \pi(i,j)c_{ij}a_ib_j = \max \frac{1}{4}(a_1b_1 + a_1b_2 + a_2b_1 - a_2b_2) = \frac{1}{2}\,.$$[1]

So the classical value of the game is

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

2.2 Quantum strategies

Now let's consider quantum players allowed to use shared entanglement to help determine their answers. The game is the same, and in particular the verifier, and the question and answer sets, remain classical. Using that answers in an XOR game are always binary we can represent each player's strategy as a family of observables.

Definition 4. 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.

[1] The quantity $\beta(G)$ is called the "bias" of the game, which is twice the largest possible advantage over a random strategy.

The laws of quantum mechanics state that if Alice and Bob measure their respective half of an entangled state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ using observables X and Y, then the product of the outcomes they obtain has expectation

$$\begin{aligned}
\mathrm{E}[a \cdot b] &= \Pr\big((a,b) = (0,0)\big) + \Pr\big((a,b) = (1,1)\big) - \Pr\big((a,b) = (0,1)\big) - \Pr\big((a,b) = (1,0)\big) \\
&= \langle\psi|X^0 \otimes Y^0|\psi\rangle + \langle\psi|X^1 \otimes Y^1|\psi\rangle - \langle\psi|X^0 \otimes Y^1|\psi\rangle - \langle\psi|X^1 \otimes Y^0|\psi\rangle \\
&= \langle\psi|X \otimes Y|\psi\rangle \ \in \ [-1,1] \ .
\end{aligned}$$

You can verify that if we restrict to $d = 1$, 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 computation above just returns the product of the observables — this is basically just a classical deterministic strategy.

With this notation in place we can define the optimal expected payoff for quantum players in an XOR game:

$$\beta^*(G) = \sup_{\substack{a_i, b_j \in \mathbb{C}^{d\times d} \\ a_i^2 = b_j^2 = \mathbb{I} \\ a_i = a_i^\dagger \\ b_j = b_j^\dagger}} \sum_{i,j} A_{ij} \cdot \langle\psi|a_i \otimes b_j|\psi\rangle \ . \tag{2}$$

Note that if we restrict the entanglement to $|\phi_d\rangle = d^{-1/2} \sum_{i=1}^d |i\rangle|i\rangle$, a d-dimensional maximally entangled state, we obtain a particularly simple formula, since in this case

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

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

Example 5. 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; in fact, you can recognize some standard matrices from quantum information. 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 5), 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 we proved optimal for classical players.

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 $\frac{1}{2} + \frac{1}{4}$ (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!

2.3 An SDP formulation 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 expected payoff for quantum players is given by (2):

$$\beta^*(G) = \sup_{\substack{A_i,B_j\in\mathbb{C}^{d\times d}\\ A_i^2=B_j^2=\mathbb{I}\\ A_i=A_i^\dagger\\ B_j=B_j^\dagger}} \sum_{i,j} G_{ij}\cdot\langle\psi|A_i\otimes B_j|\psi\rangle \leq \sup_{\substack{\vec{u}_i,\vec{v}_j\in\mathbb{R}^{d^2}\\ \|\vec{u}_i\|=\|\vec{v}_j\|=1}} \sum_{i,j} G_{ij}\vec{u}_i\cdot\vec{v}_j = \mathrm{SDP}(G)\ . \tag{3}$$

This 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

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

The expression on the right-hand side of (3) is nice because it is directly analogous to the expression (1) for the classical value: 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 program will easily recognize that $\mathrm{SDP}(AG)$ is, indeed, an SDP. In particular, the quantity can be approximated to within $\pm\varepsilon$ in time polynomial in n, m, and $\log(1/\varepsilon)$.

Note that (3) 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 earlier is tight.

Theorem 6 (Tsirelson)**.** *For G the CHSH game, it holds that $\beta^*(G) \leq \frac{\sqrt{2}}{2}$.*

Proof. For the CHSH game we can write

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

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, and considered the square of the game value. We will revisit his proof a little later.

Theorem 6 shows that the bound (3) is tight for the CHSH game. What is amazing is that it is *always* tight, for any XOR game! This is another result by Tsirelson.

Exercise 7. Show that given a vector solution to SDP(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 $\langle \phi_D | C(u) \otimes C(v) | \phi_D \rangle$?]*

The fact that (3) is an equality has amazing consequences for the study of XOR games. In particular, it implies that:

- The right-hand size of (3) can be expressed as a semidefinite program, hence the maximum expected payoff of quantum players can be computed efficiently (recall that for classical players it is NP-hard);[2]

- The proof of Tsirelson's theorem in Exercise 7 is explicit, hence an optimal quantum strategy can always be found efficiently;

- Grothendieck's inequality from functional analysis shows that the ratio of SDP$(G)/\beta^*(G)$ is always at most a universal constant, Grothendieck's constant $K_G \leq 1.782\ldots$: this implies that, in XOR games, quantum players can only ever achieve a payoff that is a constant factor larger than the optimal classical payoff.

Remark 8. There are many interesting games that are not XOR games. A good example 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. It is known that XOR games cannot be pseudo-telepathy games.

3 Rigidity for quantum games

Let's look back at the proof of Theorem 6. It has an interesting "rigidity" property. If we try to make all inequalities tight, then we don't have much choice. 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 $v\vec{v}_1$. 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, using that $\vec{v}_0$ and $\vec{v}_1$ are both unit vectors, 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 left 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(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.

[2] This fact holds for XOR games. However, in general it can be a very hard problem to determine the maximum success probability of quantum players in a two-payer game: Slofstra [Slo17] recently showed that the problem is undecidable.

Exercise 9. Suppose $\vec{u}_0, \vec{u}_1, \vec{v}_0, \vec{v}_1$ are unit vectors that achieve a value of $\frac{\sqrt{2}}{2} - \varepsilon$, for some small $\varepsilon > 0$, in SDP(G) for G the CHSH game. 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 above argument formulates "rigidity" of the CHSH game at the level of vectors. Our goal in this section 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 5 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.

3.1 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

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

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 $\|A\|_\sigma = \langle A, A \rangle_\sigma^{1/2}$.

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 10. 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

$$\mathrm{E}_{x,y \in G}\, \Re\big(\big\langle f(x)^* f(y), f(x^{-1}y)\big\rangle_\sigma\big) \;\geq\; 1 - \varepsilon \;, \tag{4}$$

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

Remark 11. The condition (4) in the definition is very closely related to 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, and the condition (4) will arise naturally from our calculations in the next section.

3.2 The Gowers-Hatami theorem

There are many possible notions for approximate group representation. The most often considered one replaces the norm in Definition 10 by the operator norm. An inconvenient of that variant is that in general approximate representations are not always close to exact representations (see, for example, the famous problem on "approximately commuting" versus "nearly commuting" operators). In contrast, Gowers and Hatami [GH15] showed that in the case of Definition 10, approximate group representations 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 12 (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 (the proof is implicit in their argument) that does not seem to allow to recover this estimate.

Note that Theorem 12 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 dimension normalization induced by the norm $\| \cdot \|_\sigma$ will barely notice this, but it will be impossible to "round" the approximate representation obtained to an exact one without modifying the dimension.

Exercise 13. Prove Theorem 12 for the case where G is the single-qubit Weyl-Heisenberg group, which is the 8-element matrix group 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*

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

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

$$\hat{f}(\rho) = \mathrm{E}_{x \in G}\, f(x) \otimes \overline{\rho(x)}. \tag{5}$$

In case G is abelian, we always have $d_\rho = 1$, the tensor product is a product, and (5) reduces to the usual definition of Fourier coefficient. The only properties we will need of irreducible representations is that they satisfy the relation

$$\sum_\rho d_\rho \,\mathrm{Tr}(\rho(x)) = |G| \delta_{xe}\,, \tag{6}$$

for any $x \in G$. Note that plugging in $x = e$ (the identity element in G) yields $\sum_\rho d_\rho^2 = |G|$.

Proof of Theorem 12. 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 (5) of $\hat{f}(\rho)$ and for the third we used (6) 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 $x \in G$,

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

where the last equality uses (6). It then follows that

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

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

3.3 Rigidity for the CHSH game

Now let's see what this all has to do with the CHSH game. We will use the 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 14. *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, succeeds with probability at least $1 - \varepsilon$ 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*

$$\big\| V_A \otimes V_B |\psi\rangle - \frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big) \otimes |\psi'\rangle \big\|^2 = O(\sqrt{\varepsilon})\,, \tag{7}$$

and

$$\big\|(V_A \otimes V_B)(A_0 \otimes \mathrm{Id})|\psi\rangle - (\sigma_X \otimes \mathrm{Id})\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \otimes |\psi'\rangle\big\| \,=\, O\big(\sqrt{\varepsilon}\big)\;, \tag{8}$$

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. (8) 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 see that it can be exploited all the way into a complete "leash" for an arbitrary computation.

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

$$\begin{aligned}(A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1)^2 &= ((A_0 + A_1) \otimes B_0 + (A_0 - A_1) \otimes B_1)^2 \\ &= 4\,\mathrm{Id} \otimes \mathrm{Id} + (A_1A_0 - A_0A_1) \otimes (B_0B_1 - B_1B_0)\;.\end{aligned} \tag{9}$$

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 (9), 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 3.1, this says that

$$\|A_0A_1 + A_1A_0\|_{\rho_A}^2 = O(\varepsilon)\;, \tag{10}$$

i.e. A_0 and A_1 approximately commute.

Now here comes the key observation. Consider the 8-element group H generated by the Pauli matrices σ_X and σ_Z, i.e.

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

Then I claim that A_0 and A_1 induce an approximate representation of H, 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 (10) and the fact that A_0 and A_1 are observables, it is immediate to verify that the conditions of Theorem 12 are satisfied, i.e. f is an $(O(\varepsilon), \rho_A)$-representation of H.

Applying the theorem, there must exist an exact representation of H to which f is close. However, the representation theory of H 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 H, which is precisely given by the Pauli matrices!

We're almost done. 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 then 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.

$$\frac{1}{2}\langle\psi|(V_A \otimes V_B)^\dagger\big((\sigma_X \otimes \sigma_X + \sigma_Z \otimes \sigma_Z) \otimes \mathrm{Id}\big)(V_A \otimes V_B)|\psi\rangle \geq 1 - O(\varepsilon)\,. \qquad (11)$$

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 -1. Therefore, (11) 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. □

Acknowledgments. I thank Richard Cleve and Laura Mančinska for pointing out multiple typos in previous versions of this document. I have no doubt that many remain; I am grateful to anyone for notifying me if they find more.

References

[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.

[CHTW04] Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Computational Complexity, 2004. Proceedings. 19th IEEE Annual Conference on*, pages 236–249. IEEE, 2004.

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

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

[KRR14] Yael Tauman Kalai, Ran Raz, and Ron D Rothblum. How to delegate computations: the power of no-signaling proofs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 485–494. ACM, 2014.

[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.

[Slo17] William Slofstra. The set of quantum correlations is not closed. *arXiv preprint arXiv:1703.08618*, 2017.

[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.