From Operator Algebras to Complexity Theory and Back

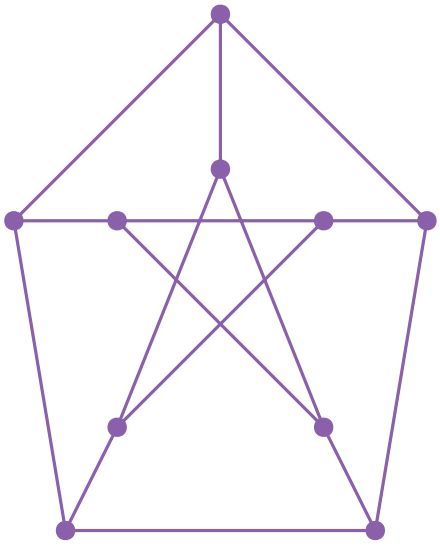

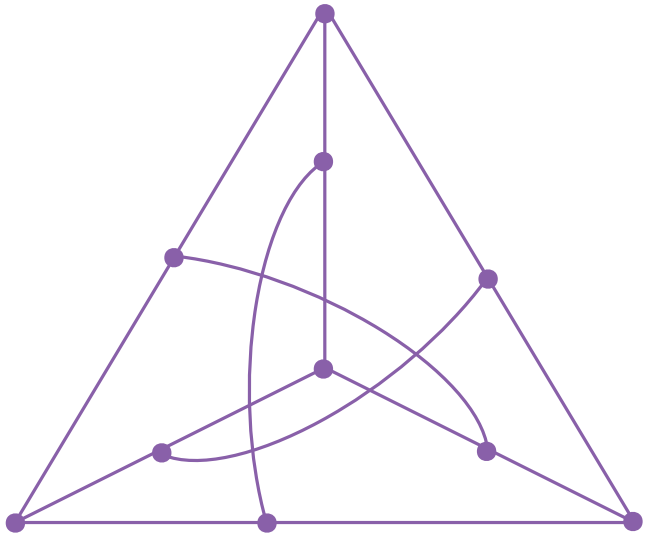

Thomas Vidick

Quantum mechanics and the theory of operator algebras have been intertwined since their origin. In the 1930s [20] 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 [6] Connes made a casual remark that has become a central conjecture in the theory of operator algebras. Paraphrasing, Connes' remark was that any finite von Neumann algebra, i.e., one that has a finite trace, "ought to" be well approximated by finite-dimensional matrix algebras. Thanks to the work of other mathematicians, such as Kirchberg and Voiculescu, the remark, now known as *Connes' embedding conjecture* (CEC), has become one of the most important open problems in operator algebras. (Formally, the CEC states that every von Neumann algebra type II_1 factor embeds into an ultrapower of the hyperfinite II_1 factor.) Kirchberg showed that CEC 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 [12]. Voiculescu gave a reformulation in terms of the existence of matrix microstates in free probability [19]. Rădulescu showed that a group is hyperlinear if and only if its group von Neumann algebra satisfies CEC [16]. Goldbring and Hart showed that CEC holds if and only if every type II_1 tracial von Neumann algebra has a computable universal theory [8].[1] Many more equivalent formulations are known (see e.g. [4] for a survey).

In this note we are concerned with an equivalent formulation of CEC known as "Tsirelson's problem." The

Thomas Vidick is a professor of computing and mathematical sciences at California Institute of Technology. His email address is vidick@caltech.edu.

The author gratefully acknowledges support from the IQIM, an NSF Physics Frontiers Center (NSF Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-12500028), NSF CAREER Grant CCF-1553477, AFOSR YIP award number FA9550-16-1-0495, and a CIFAR Azrieli Global Scholar award.

DOI: https://doi.org/10.1090/noti1980

[1] *This connection with a question in logic is, to the best of our knowledge, not related in any direct way with the connection discussed in the present article.*

problem arose from Tsirelson's study of the nonlocal properties of entanglement, a puzzling phenomenon first brought to light by Einstein, Podolsky, and Rosen. Loosely speaking, Tsirelson's problem asks about the appropriate way to model locality in quantum mechanics: it asks if the Hilbert space associated with spacelike isolated regions always factors as a tensor product of Hilbert spaces on which observables associated with each region can be localized.

With the development of quantum computing, entanglement and the nonlocality of quantum mechanics have come to be seen as resources for computation and communication. Thus the CEC, through its equivalence with Tsirelson's problem, is tied back to questions in quantum information whose study has been pursued largely independently of the developments in operator algebras that led to Connes' original remark. The purpose of this note is to show how work in quantum computing theory leads to a complexity-theoretic conjecture whose proof would imply a negative answer to Tsirelson's problem. (Interestingly, a *refutation* of the conjecture is not known to have any implications for CEC.) Informally, the complexity-theoretic conjecture states that the class MIP* of problems that can be decided by a polynomial-time verifier interacting with quantum provers sharing entanglement contains undecidable languages; we explain these terms as the article progresses.

We start with a formulation of Tsirelson's problem and tie it to Bell's work in the foundations of quantum mechanics. We then dive into the key notion of *interactive proof*, which has played a major role in the development of complexity theory over the past three decades. This allows us to introduce the theory of nonlocal games in quantum information and relate it, through the framework introduced by Bell, to Tsirelson's problem. Finally, we close the loop by formulating a conjecture on the power of quantum interactive proof systems whose proof would lead to a refutation of CEC.

Tsirelson's Problem

In the early 1980s Boris Tsirelson wrote a series of papers laying out the mathematical formalism for the systematic study of the nonlocal properties of quantum mechanics. In one of the papers [18] Tsirelson introduces two sets that capture certain kinds of distributions that arise from measurements on entangled states. For a (separable) Hilbert space $\mathcal{H}$ a *projection valued measure* (PVM) on $\mathcal{H}$ is a finite collection $\{P_1, \dots, P_m\}$ of projections on $\mathcal{H}$ such that $\sum P_i = \mathrm{Id}$. For arbitrary finite indexing sets X, Y, A, B Tsirelson considers the convex subsets Q^c_{ABXY} and Q^s_{ABXY} of $[0,1]^{A\times B\times X\times Y}$, where the superscripts c and s stand for *commuting* and *spatial*, respectively:

$$\begin{aligned} Q^c_{ABXY} = \{ \, & (\langle \psi, A^x_a B^y_b \psi \rangle)_{a,b,x,y} \, : \\ & \mathcal{H} \text{ Hilbert space}, \ \psi \in \mathcal{H}, \ \|\psi\| = 1, \\ & \forall (x,y) \in X \times Y, \ \{A^x_a\}_{a\in A}, \{B^y_b\}_{b\in B} \text{ PVM on } \mathcal{H} \\ & \text{s.t. } [A^x_a, B^y_b] = 0 \ \forall (a,b) \in A \times B \} \, , \end{aligned} \tag{1}$$

$$\begin{aligned} Q^s_{ABXY} = \{ \, & (\langle \psi, A^x_a \otimes B^y_b \psi \rangle)_{a,b,x,y} \, : \\ & \mathcal{H}_A, \mathcal{H}_B \text{ Hilbert spaces}, \ \psi \in \mathcal{H}_A \otimes \mathcal{H}_B, \\ & \|\psi\| = 1, \ \forall (x,y) \in X \times Y, \\ & \{A^x_a\}_{a\in A}, \{B^y_b\}_{b\in B} \text{ PVM on } \mathcal{H}_A, \mathcal{H}_B, \text{ resp.} \} \, . \end{aligned} \tag{2}$$

By taking direct sums of PVMs and scaled vectors it is not hard to see that both sets are convex. Moreover, in case the Hilbert spaces in both definitions are taken to be finite-dimensional, the two sets can be shown to coincide. In his paper Tsirelson states as "fact" the claim that $Q^s_{ABXY} = Q^c_{ABXY}$ for arbitrary separable Hilbert spaces and all finite A, B, X, Y. Having realized that a proof of the claim seemed elusive (with the inclusion $Q^s_{ABXY} \subseteq Q^c_{ABXY}$ being the only obvious one), in a subsequent note[2] 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 $\overline{Q^s_{ABXY}} = Q^c_{ABXY}$. Two and a half decades after its introduction Tsirelson's first problem was solved by Slofstra [17], who used techniques from the theory of nonlocal games to show the existence of finite indexing sets A, B, X, Y such that $Q^s_{ABXY} \neq Q^c_{ABXY}$. But Tsirelson's "even more important problem" remains open:

Tsirelson's "even more important" problem:
Does $\overline{Q^s_{ABXY}} = Q^c_{ABXY}$ for all finite sets A, B, X, Y?

Building on work of many others, including Fritz [7] and Junge et al. [11], Ozawa [15] showed that Tsirelson's "even more important" problem (hereafter referred to as "Tsirelson's problem") is equivalent to CEC, thereby lifting the problem from a question in the foundations of quantum mechanics to a central conjecture in operator algebras. In the remainder of this article we explain the relation of Tsirelson's problem with quantum nonlocality, as seen through the lens of complexity theory. We start by explaining the origins of Tsirelson's problem in the foundations of quantum mechanics.

Bell Experiments

On to quantum information. Tsirelson's problems are rooted in the quantum phenomenon of *entanglement*. We will not attempt to give a precise mathematical definition of the term here or explain its physical underpinnings:

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

roughly, a state of multiple particles is said to be entangled when some observable degrees of freedom (e.g. angular momentum) of the particles are highly correlated.

Almost three decades after the publication of the famous paper by Einstein, Podolsky, and Rosen (EPR) using the "spookyness" of entanglement to argue the incompleteness of quantum mechanics, John Bell in 1964 was the first to introduce a concrete (though impractical) thought experiment that sharply delineates the predictions of quantum mechanics from those of classical theory. Bell considered the following scenario, today referred to as a "Bell experiment." Suppose that two distant physical systems (e.g. spatially localized collections of particles) are initialized in an arbitrary state; the systems may be as correlated as is allowed by the physical theory.[3] Suppose further that the first (resp., second) system can be measured using any one of a finite collection of possible procedures A^x (resp., B^y) indexed by $x \in X$ (resp., $y \in Y$). Suppose finally that performing the measurements A^x on the first system and B^y on the second yields a pair of outcomes $(a, b) \in A \times B$, where A, B are finite sets. Define the resulting "correlation set" as the convex set $K_{ABXY} \subseteq [0,1]^{A\times B\times X\times Y}$ that contains all tuples (p_{abxy}) such that there is an initial quantum state for the systems and measurements on them that lead to outcomes (a, b) with probability p_{abxy} whenever measurements A_x and B_y are performed. (The convexity of K follows as soon as the theory counts any probabilistic mixture of allowed states as an allowed state.)

Interestingly, giving a precise mathematical definition of the correlation set Q_{ABXY} associated with measurements on quantum systems requires us to make a nontrivial design choice. In quantum mechanics the state of a physical system is represented by a positive linear functional ω of norm 1 on $\mathcal{B}(\mathcal{H})$, the bounded linear operators acting on a separable Hilbert space $\mathcal{H}$. It turns out to be sufficient to restrict our attention to vector states, which are those ω such that there exists a unit vector $\psi \in \mathcal{H}$ such that $\omega(A) = \langle \psi, A\psi \rangle$. To each measurable quantity, such as the location of a photon or the spin of an electron, is associated an *observable*, which is a bounded self-adjoint operator on $\mathcal{H}$. In general each of the two systems can be measured using a certain set of allowed observables, $\mathcal{O}_A \subseteq \mathcal{B}(\mathcal{H})$ for the first and $\mathcal{O}_B \subseteq \mathcal{B}(\mathcal{H})$ for the second.

In quantum mechanics a measurement in general perturbs the state that it is performed on. As a result, two different observables cannot always be measured simultaneously: the order in which the measurements are performed may matter. In a Bell experiment it is assumed that the two systems are "distant" and in particular do not interact. Thus in any reasonable physical realization of a Bell experiment it ought to be possible to perform the measurements on both systems in any order and obtain the same distribution on outcomes. Von Neumann [20] showed that this *joint measurability* condition is equivalent to the algebraic requirement that any $A^x \in \mathcal{O}_A$ and $B^y \in \mathcal{O}_B$ commute. Given a pair of observables (A^x, B^y) that satisfy this condition, we can specify the probability that their joint measurement on a state ψ returns a pair of outcomes (a, b) as follows. Using that observables are self-adjoint we can write the spectral decompositions as $A^x = \sum_a \lambda_a A^x_a$ and $B^y = \sum_b \mu_b B^y_b$ with a finite set of real eigenvalues λ_a, μ_b (the assumption that there is a finite number of distinct eigenvalues is unimportant but generally follows from the interpretation that each eigenvalue of an observable is associated with a measurement outcome). The probability of observing (a, b) is then given by $\langle \psi,\ A^x_a B^y_b\ \psi \rangle$. Using that A^x_a, B^y_b are positive semidefinite and commute, and that $\sum_a A^x_a = \sum_b B^y_b = \mathrm{Id}$, the formula specifies a well-defined family of distributions. The resulting correlation set is precisely the "quantum commuting set" Q^c_{ABXY} introduced in (1).

In his study of quantum correlations Bell was not particularly interested in the distinction between Q^s and Q^c—as already mentioned, he considered only finite-dimensional systems for which the sets coincide. Rather, Bell was interested in how either set relates to the set of classical correlations C_{ABXY}, defined analogously except that the initial state of the systems is required to have a classical description (equivalently, $\mathcal{O}_A$ and $\mathcal{O}_B$ both generate commuting algebras). Bell's work was motivated by a desire to give a logically clear statement that would place the EPR thought experiment on a firm mathematical footing, a task in which he largely succeeded!

Concisely stated, Bell's result is the identification of an explicit point in Q^s_{ABXY} that does not lie in C_{ABXY} for some large enough X, Y. (Bell considers the case of infinite X, Y, but his construction can be easily discretized. A few years later Clauser et al. gave an explicit separation for $|X| = |Y| = |A| = |B| = 2$.) An insightful reformulation of Bell's theory can be given using the language of *nonlocal games*. We motivate these games by showing how they arise in an a priori entirely distinct line of work, the theory of interactive proofs in complexity theory.

In nonrelativistic quantum mechanics it is generally assumed that the number of degrees of freedom of the physical systems considered is finite, in which case the underlying Hilbert space $\mathcal{H}$ is finite-dimensional. (In some cases, such as the position and momentum observables, $\mathcal{H}$ is infinite, but the algebra of physically

[3] *For example, consider two "distant" coins that are both in state "heads" with probability $\frac{1}{2}$ and in state "tails" with probability $\frac{1}{2}$. This is an "allowed state" in classical physics.*

relevant observables is a type I von Neumann algebra, for which the ensuing discussion applies as well.) In finite dimensions it is a simple consequence of Schur's lemma that for any two mutually commuting collections of observables $\mathcal{O}_A$ and $\mathcal{O}_B$ acting on $\mathcal{H}$ there is an isomorphism $\mathcal{H} \simeq \bigoplus_i (\mathcal{H}_{A,i} \otimes \mathcal{H}_{B,i})$ such that $\mathcal{O}_A = \bigoplus_i \mathcal{O}_{A,i}$ and $\mathcal{O}_B = \bigoplus_i \mathcal{O}_{B,i}$, where for each i, $\mathcal{O}_{A,i}$ (resp., $\mathcal{O}_{B,i}$) acts nontrivially only on $\mathcal{H}_{A,i}$ (resp., $\mathcal{H}_{B,i}$). In such cases the sets Q^c_{ABXY} and Q^s_{ABXY} coincide, and indeed Bell formulated his theory using tensor products. More generally, in relativistic quantum mechanics the local algebra of observables associated with a bounded region of space-time is generally believed to be of type III. Under additional physical assumptions such as strict spacelike separation between regions or some bounded energy-density assumptions, it is possible to show rigorously that the algebras have the *split property*; i.e., there exists a type I factor $\mathcal{F}$ such that $\mathcal{A} \subset \mathcal{F} \subset \mathcal{B}'$ [3]. In these cases it is also known that Tsirelson's problem has an affirmative answer. Yet not all algebras have the property! For example, it is known not to hold for certain unbounded spacelike separated regions [3].

Interactive Proofs

We begin with an instructive example. Suppose we are given an explicit description of two graphs G and H on the same vertex set $V = \{1, \dots, n\}$ (see Figure 1). Can you tell if the two graphs are isomorphic? On the example you probably can, but in general this is a hard problem, by which we mean that there is no known algorithm that runs in time at most some fixed polynomial in n and provides the correct answer on all pairs of graphs on n vertices for all integers $n \geq 1$. But if someone—an all-powerful *prover*—conveniently hands you a *proof* of isomorphism in the form of an explicit map between the graph's respective vertex sets, it is easy to verify that the map is a bijection that preserves the adjacency relation of the graphs and hence is a valid graph isomorphism.

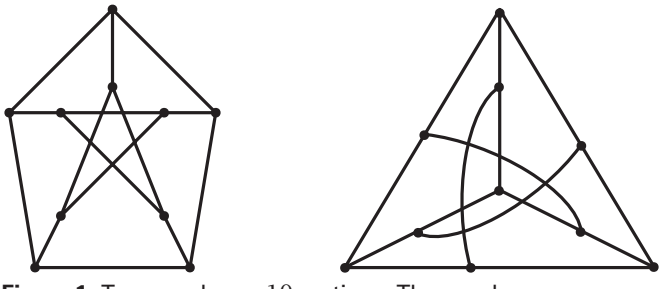

Figure 1. Two graphs on 10 vertices. The graphs are isomorphic.

Thus isomorphic graphs always have valid proofs of isomorphism that can be efficiently verified. In complexity-theoretic terms we say that the graph isomorphism problem lies in the class NP of problems that have efficiently verifiable proofs.[4] What about the complement problem, where the verification procedure aims to verify a proof supporting the claim that the two graphs are *not* isomorphic? This seems more difficult, and indeed there is no known efficiently verifiable proof for graph nonisomorphism. Yet consider the following *interactive* proof, executed between the *verifier* (who wishes to verify that the graphs are nonisomorphic) and the *prover*. At the first step the verifier privately selects a uniformly random bit $c \in \{0, 1\}$ and a uniformly random permutation π on $\{1, \dots, n\}$. If $c = 0$ the verifier sends the permuted graph $\pi(G)$ obtained by relabeling the vertices of G according to π to the prover. If $c = 1$ the verifier sends $\pi(H)$ to the prover. The prover is asked to respond with the label of one of the two graphs, G or H. The verifier accepts the prover's answer if and only if the returned label matches that of the chosen graph (i.e., G in case $c = 0$ and H in case $c = 1$).

To check that this is a valid proof system for graph nonisomorphism we need to consider two cases. First, in case the graphs G and H are not isomorphic there exists a strategy for the prover that is always accepted by the verifier. Indeed, by investing sufficient computational effort (e.g. iterating over all possible relabelings) it is possible for the prover to uniquely determine which of G or H it was sent a relabeling of. Second, in case the graphs G and H are isomorphic then it is easy to see that no prover, however powerful, can succeed with probability strictly larger than $\frac{1}{2}$. This is because in this case the prover receives a uniformly random relabeling of G, or equivalently of H, and has no way to determine from which graph it was obtained.

Summarizing, in case the graphs are not isomorphic there is a prover that is always accepted, and in case they are not, any prover is accepted with probability at most $\frac{1}{2}$. Repeating the interaction a few times in sequence and accepting only if all interactions accept amplifies the gap between 1 and $\frac{1}{2}$ exponentially fast. In complexity-theoretic terms, we have shown the graph nonisomorphism problem lies in the class IP of problems that have efficiently verifiable randomized interactive proofs.[5]

[4] *Formally, the graph isomorphism problem is modeled as a* language, *i.e., the set* $L_{iso} \subseteq \{0,1\}^*$ *of all binary representations of pairs of graphs* (G_0, G_1) *such that* G_0 *and* G_1 *are isomorphic. The statement* $L_{iso} \in$ NP *means that there is an efficient verification procedure such that given an arbitrary* $x \in \{0,1\}^*$, *if* $x \in L_{iso}$, *then there is a proof* π *that the verification procedure accepts, whereas if* $x \notin L_{iso}$, *then no proof is accepted.*

[5] *The consideration of* randomized *proofs comes with an inevitable caveat: for any statement needing verification, the verifier may have a small chance of making the wrong decision. This fact is an integral part of the definition of a "randomized verification procedure." It is generally required that the chance of making an error should be bounded below a fixed constant, say 1/3. In our example this is achieved by performing two repetitions of the protocol and accepting if*

This example demonstrates the power of randomization and interaction in verifying proofs. The discovery of this gain in expressive power initiated major lines of work in complexity theory and cryptography along which research very actively continues today. One of the most important milestones, which is relevant for this article, is the characterization by Babai, Fortnow, and Lund [2] of the class of problems that can be decided by an efficient verifier who may interact with *two* (or more) cooperating but *mutually noninteracting* provers. Here by "noninteracting" we mean that the provers are not allowed to communicate directly between themselves while they are being interrogated by the verifier. (The provers may still discuss, and agree on, a common strategy ahead of time.)

From the provers' point of view, two is not more powerful than one, as the prover is always considered to have infinite computational power. Instead, the addition of a second prover gives more power to the verifier, who now has at her disposition the entire policeman's arsenal in her interrogation of (possibly) colluding spies. A concrete way in which this power can be used is by sending an entire list of questions to one prover and a single question from the list to the other. By checking that the provers return consistent answers, the verifier can ascertain that the first prover's answer to each of the questions in the list is made independently of the other questions (since it has to match the second prover's answer, which by definition depends on a single one of the questions). The same effect could not be achieved with a single prover, even through sequential interaction, as there would be no way to guarantee that the prover's answers in a certain round are independent from his questions in earlier rounds.

This observation allows us to think of the information held by the second prover, the list of all possible answers that could be given to each of the verifier's individual questions, as an exponentially long proof that the verifier has the ability to query at polynomially many locations (by asking the first prover for all entries (s)he cares about and checking consistency with the second at a randomly chosen entry). The aforementioned result of Babai et al. is that any language that can be verified by an *exponential-time* verifier given a static exponentially long proof can be verified by a *polynomial-time* verifier in this model. In complexity-theoretic terms, NEXP $\subseteq$ MIP, where NEXP is the exponential-time version of NP and MIP denotes the class of problems that can be decided by a randomized polynomial-time verifier interacting with multiple provers.

Tracing back through Babai et al.'s construction, after additional refinements complexity theorists have arrived at the striking statement that any polynomial-size proof that can be verified in polynomial time by reading the entirety of the proof can systematically be (efficiently) encoded in a format such that the proof can be verified by reading only a very small, in fact constant, number of entries in the proof! This result is known as the PCP theorem, for "probabilistically checkable proofs," a landmark result in complexity theory. For our purposes we need not go further in this direction. Instead we investigate a variant of multiprover interactive proofs in which the provers may implement a quantum strategy that makes use of entanglement.

and only if both repetitions accept.

Nonlocal Games

Starting with the work of Cleve et al. [5] in 2004, computer scientists started investigating the consequences of entanglement for the theory of multiprover interactive proofs. To simplify the presentation and make the connection with Bell experiments, it is convenient to focus on interactive proofs that involve a single round of interaction. Such interactive proofs can be described using the language of *nonlocal games* that we now introduce.

A nonlocal game involves a *referee* (a.k.a. verifier) and two cooperating but noncommunicating *players* (a.k.a. provers), generally referred to as "Alice" and "Bob." The game proceeds in a single round: the referee selects a pair of questions (x, y) according to a distribution π on $X \times Y$. They send x to Alice and y to Bob. The players each reply with an answer, $a \in A$ for Alice and $b \in B$ for Bob. Finally, the referee evaluates a decision predicate $V(a, b|x, y) \in \{0, 1\}$. If they obtain 1 we say that the players win, and if not the players lose. The rules of the game (in the form of the sets X, Y, A, B), the distribution π, and the predicate V are publicly known. The players may agree on a strategy ahead of time, but once the game has started they are no longer allowed to communicate. The "value" of a game is the maximum probability, over the verifier's choice of questions and the players' strategy, that the players win in the game. To define this precisely we distinguish between classical strategies, in which each player is restricted to evaluating a function $f_A : X \times \Omega \to A$ and $f_B : Y \times \Omega \to B$, respectively, where Ω is an arbitrary probability space,[6] and quantum strategies, in which each player may perform a measurement on a quantum state. This naturally leads to three possible "values" of a nonlocal game: the *classical value*

$$\omega(G) = \sup_{f_A, f_B} \sum_{x,y} \pi(x, y) \sum_{a,b} V(a, b|x, y) \times \int_\Omega 1_{f_A(x,\omega)=a} 1_{f_B(y,\omega)=b} \, d\omega \,, \tag{3}$$

where the supremum ranges over all probability spaces Ω and measurable $f_A : X \times \Omega \to A$ and $f_B : Y \times \Omega \to B$;

[6] *The space Ω models the use of "shared randomness" between the players, i.e., any question-independent prior information they may wish to share.*

the *spatial value*

$$\omega^s(G) = \sup_{A^x, B^y} \sum_{x,y} \pi(x,y) \sum_{a,b} V(a,b|x,y) \times \langle \psi, (A_a^x \otimes B_b^y)\, \psi \rangle , \tag{4}$$

where the supremum is taken over all separable Hilbert spaces $\mathcal{H}_A$, $\mathcal{H}_B$ and finite collections of observables $\mathcal{O}_A = \{A^x\}_x \subseteq \mathcal{B}(\mathcal{H}_A)$ and $\mathcal{O}_B = \{B^y\}_y \subseteq \mathcal{B}(\mathcal{H}_B)$; and the *commuting value*

$$\omega^c(G) = \sup_{A^x, B^y} \sum_{x,y} \pi(x,y) \sum_{a,b} V(a,b|x,y) \times \langle \psi, A_a^x B_b^y\, \psi \rangle , \tag{5}$$

where the supremum is taken over all separable Hilbert spaces $\mathcal{H}$ and finite collections of observables $\mathcal{O}_A = \{A^x\}_x$ and $\mathcal{O}_B = \{B^y\}_y$ on $\mathcal{B}(\mathcal{H})$ such that $[A^x, B^y] = 0$ for all x, y.

In general, $\omega(G) \leq \omega^s(G) \leq \omega^c(G)$. By employing suitable scaling arguments to relate arbitrary linear functionals to games it can be shown that Tsirelson's problem is equivalent to the assertion that $\omega^s(G) = \omega^c(G)$ for all G. Before tackling this question we give an example that demonstrates $\omega(G) < \omega^s(G)$. Since a nonlocal game can naturally be translated into a Bell experiment by associating a measurement to each possible question to a player in the game (with as many outcomes as there are possible answers to that question), such a game establishes a separation between C_{ABXY} and Q^s_{ABXY}, proving Bell's theorem.

The example is called the "Magic Square game" (MS) and is due to Aravind [1], building on work of Mermin and Peres. In this game the players are asked to fill in entries of a "magic square" to which there is no solution (see Figure 2 for an explanation of the rules of the game; there is no relation between the game and the more usual kind of "magic square" studied in combinatorics). Aravind showed that for any strategy that classical players could employ, there is always a question on which the strategy must sometimes return a wrong answer, whereas there exists a quantum strategy that involves performing measurements on a specific quantum state using which the players can succeed 100% of the time. (The quantum strategy is explained in the next section.) In other words, the classical value of the Magic Square game satisfies $\omega(MS) < 1$ (in fact, $\omega(MS) = 17/18$), but the quantum spatial value is $\omega^s(MS) = 1$.

After the publication of the EPR paper most physicists brushed aside the "weirdness" of entanglement, preferring to focus on those aspects of quantum mechanics that "work" and yield useful scientific predictions. Although the work of Bell and others established the nonlocality of entanglement on a firm footing, both theoretical (no "spooky action at a distance") and experimental (culminating in the recent "loophole-free" tests [9]), for a long time

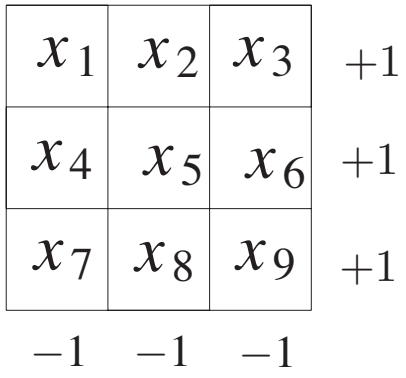

Figure 2. The Magic Square game. The first player, Alice, is sent the label of a row or column chosen uniformly at random. The second player, Bob, is sent the label of a uniformly random entry in Alice's row or column. Alice should return a $\{+1, -1\}$-valued assignment to the three variables in her row or column such that the product of the entries is as indicated on the figure. Bob should return a $\{+1, -1\}$-valued assignment to the variable he is asked about that matches Alice's assignment to the same variable. The players' winning probability is the probability, over the referee's choice of questions as well as randomness in their strategy, that the players' answers satisfy the constraints imposed. The fact that there is no assignment to the variables satisfying all constraints implies that the players do not have a perfect consistent strategy.

the special correlations afforded by quantum mechanics remained an oddity, of interest to researchers in foundations but of little practical relevance. The situation changed drastically in the early 1990s with the discovery that entanglement could act as a resource for quantum information tasks. For example, the technique of "superdense coding" allows transmission of two bits of information using a single qubit (quantum bit) aided by one pair of entangled particles. In 1991 Ekert introduced a protocol for quantum key distribution, a task in quantum cryptography that constitutes one of the most promising applications of quantum information to communication networks, based on entanglement: Ekert argued that by verifying that two distant, cooperating parties share a quantum state that allows them to succeed in a nonlocal game, the parties are able to certify that the quantum state they share is entirely uncorrelated with any third party and in particular from a malicious eavesdropper (as such sharing would necessarily weaken the entanglement, a phenomenon known as the *monogamy* of entanglement). In the next section we investigate consequences of entanglement and nonlocal games in another direction: complexity theory.

Consequences of Nonlocality for Complexity Theory

Consider again the Magic Square game. Observe that the game can be repurposed as a multiprover interactive proof system, as introduced in the section "Interactive Proofs."

Specifically, in the game the referee always sends a triple of questions to the first player, Alice, and one question chosen uniformly at random from the triple to the second player, Bob. Thus a classical deterministic strategy for Bob can be described using nine variables $x_i \in \{\pm 1\}$ representing Bob's answer to each of his nine possible questions, as illustrated in Figure 2. For the players' strategy to succeed in all cases (for all possible question pairs that could be chosen by the referee) these nine values together must form a satisfying assignment to the magic square constraints: $x_1x_2x_3 = 1$, $x_1x_4x_7 = -1$, etc. Indeed, consistency with Alice's strategy and the requirement that Alice's answers satisfy the row and column parity constraints imply that in a strategy that wins with probability 1 all equations must be simultaneously satisfied by the assignment specified by Bob's strategy. Yet the astute mathematician will have no difficulty devising her own proof that there does not exist an assignment to the nine variables that simultaneously satisfies all six constraints. Hence there does not exist a classical strategy in the Magic Square game that succeeds with certainty on all possible pairs of questions.

Why does this argument not rule out the existence of a perfect quantum strategy? One may think of the additional flexibility given to a quantum strategy as follows. Although the system of six equations in nine variables implied by the square has no solution in terms of variables in $\{-1, 1\}$, it has a *noncommutative* solution in the following sense: there exist nine 4×4 Hermitian matrices $X_1, \ldots, X_9$ that each square to identity and such that moreover the three matrices in any row (resp., column) (i) commute and (ii) multiply to $+\mathrm{Id}$ (resp., $-\mathrm{Id}$). (The solution is not too hard to find; as a hint, it suffices to consider matrices with coefficients in $\{0, \pm 1, \pm i\}$.) These matrices are the quantum analogue of an assignment (instead of being fixed values in $\{-1, 1\}$ they have eigenvalues in that range). Physically, the players' quantum strategy takes the following form. Initially, the players each have two spin-$\frac{1}{2}$ particles in their possession. The joint state of the four particles is described by a unit vector $\psi \in \mathbb{C}^4 \otimes \mathbb{C}^4$. Upon reception of its question, a player performs a measurement on its two particles. If the player is Bob, the measurement is the one associated with the observable X_i, where i is the index of his question. The outcome of the measurement is a value in $\{-1, 1\}$ that Bob returns as his answer. If the player is Alice, the measurement is the one associated with the three observables $X_{i_1}, X_{i_2}, X_{i_3}$ in her row or column. The fact that these three operators always commute guarantees that they can be jointly measured. The outcome is a triple of values in $\{-1, 1\}$ that Alice returns as her answer. It can then be verified, using property (ii) above and the laws of quantum mechanics, that the joint measurement of Alice's and Bob's observables on a well-chosen state ψ (independent of their respective questions) *always* results in answers that satisfy the verifier's checks in the game.

The existence of games such as the Magic Square game shatters the complexity theorist's world view regarding multiprover interactive proofs. Recall that the power of proof systems such as MIP rests on the use of "cross-checking" of the prover's answers; it is such cross-checking that allowed us to think of a multiprover proof system as a device through which the verifier is able to obtain random access to an exponentially long proof. What the example shows is that the connection between strategy and proof does not extend to the quantum case: there exist formulas that have no satisfying assignment, such as the formula that underlies the row and column constraints of the "magic square," yet the provers are able to provide valid answers to the verifier's checks 100 percent of the time. Indeed, in the case of a quantum strategy the extraction performed earlier—writing down Bob's answer to each of his possible questions—is nonsensical, because it entirely ignores the fact that in general Bob's answers may be strongly correlated, due to entanglement, with Alice's.

Nevertheless, a sequence of results in complexity theory has established techniques to design proof systems that are "resistant" to the provers' malicious use of entanglement: it is now known that any proof system, including the one of Babai et al., can be encoded in such a way that entanglement is no longer useful to malicious provers. This establishes the inclusion NEXP = MIP $\subseteq$ MIP* [10], where the * refers to the provers' allowed use of entanglement. (For historical reasons the model used for entangled-player strategies in MIP* is the "finite-dimensional" model corresponding to (4); the complexity class obtained from (5) is denoted MIP^{co}.) But we are interested in a much more exciting possibility, which is that entanglement between the provers may be used by the verifier to their advantage, yielding proof systems that allow deciding problems beyond NEXP. A priori there is no limit to how complex of a problem may be decided in this model—there is no limit to the complexity of the provers' strategy; only the verifier is restricted to being efficient. Indeed, our current state of knowledge is consistent with, and in fact points to, the possibility that MIP* contains undecidable problems. This, however, would have surprising consequences for Tsirelson's problem and Connes' embedding conjecture, as we explain in the next section.

A Complexity-Theoretic Approach to Tsirelson's Problem

Recall the characterization MIP = NEXP that follows from the work of Babai et al. on classical multiprover interactive proof systems. As discussed in the previous section, allowing the provers to share entanglement breaks the soundness of certain proof systems (such as a proof system that relies on the soundness of the Magic Square game), but

not all. Recently an important additional step has been taken by Natarajan and Wright [13], who show that MIP* is *strictly larger* than NEXP and in particular strictly larger than its classical, entanglement-free counterpart MIP. In symbols, the new result is that NEEXP $\subseteq$ MIP*, where NEEXP stands for nondeterministic *doubly* exponential time. In words, by querying two provers sharing entanglement, a polynomial-time verifier has the ability to verify the validity of a proof of doubly exponential length!

Our purpose is not to dwell too much on what to the noncomplexity theorist may seem like rather exotic alphabet soup, but rather to explore the potential significance of results such as Natarajan and Wright's and possible improvements thereof for Tsirelson's problem. For this we take a closer look at the optimization problems (4) and (5) used to define the spatial and commuting value of a game, respectively. Suppose we are given a set of rational coefficients $(\pi(x,y)V(a,b|x,y))_{a,b,x,y}$ that specify a game, i.e., an input to either (4) or (5). How would one go about algorithmically estimating the supremum to within some accuracy ε?

We sketch two possible approaches, the first "from below" and the second "from above." The first approach directly attempts to exhaustively search over all possible strategies for the players. In the simpler case of classical deterministic players, a strategy is a map from questions to answers. Such a map can be written using space $|X|\log|A|$ for Alice and $|Y|\log|B|$ for Bob simply by listing the answer that the player would give to each possible question. Guessing an optimal strategy at random within that space yields an algorithm that runs in nondeterministic time $O(|X|\log|A| + |Y|\log|B|)$. In an interactive proof system the verifier is restricted to run in polynomial time, so it can only write and read questions and answers of polynomial length, implying that $|X|, |Y|, |A|, |B|$ are each at most of exponential size. This algorithm thus implies the inclusion MIP $\subseteq$ NEXP, which as we saw applies to classical interactive proofs—but not quantum.

Indeed, the case of quantum players sharing entanglement is more difficult for two reasons. First, quantum strategies are specified by "continuous" objects, a quantum vector state $\psi \in \mathcal{H}$, and collections of observables $\mathcal{O}_A, \mathcal{O}_B \subseteq \mathcal{B}(\mathcal{H})$. Second, even if we restrict our attention to finite-dimensional strategies there is no a priori obvious bound on the dimension of a Hilbert space that is sufficient to support a state and observables that achieve a value close to the optimum. The first difficulty is easily handled: in fixed dimension d, it is possible to consider a nested sequence of finite nets $N_1 \subseteq \cdots \subseteq N_k \subseteq \cdots$ over all states and observables in that dimension such that N_k has size $k^{O(d^2)}$ and for any strategy in dimension d, there is a strategy in N_k that performs almost as well, up to an additive $(1/k)$ loss in the objective value (4). Let $\omega^s_{\leq n}$ denote the success probability of the best strategy encountered within the net N_k in dimension d over all $d, k \leq n$. Then $\{\omega^s_{\leq n}\}_{n\geq 1}$ is a bounded nondecreasing sequence of values that converges to (4) from below.[7] Note that a similar approach does not seem to apply for (5), as the operators considered in the supremum in (5) may not have finite-dimensional approximations.

It seems much more difficult to give any numerical approximation to the supremum (5). Perhaps surprisingly, there exists a "dual" approach that considers the problem of optimizing over a larger region than the feasible region of (5). The most naïve relaxation replaces each term $\langle\psi, A^x_a B^y_b \psi\rangle$ by a coefficient α_{abxy} and considers the supremum over all α_{abxy} in $[0,1]$. In general this leads to a wild overestimate which can be refined by introducing additional constraints on the α_{abxy}. For example, it should be that for any x, y the coefficients sum to 1 over all a, b, because $\sum_a A^x_a = \sum_b B^y_b = \mathrm{Id}$, and ψ is normalized. To go further the authors of [14] introduce an increasing hierarchy of constraints: informally, the idea is to introduce additional coefficients to represent quantities such as $\langle\psi, (A^{x_1}_{a_1} B^{y_1}_{b_1} A^{x_2}_{a_2} \cdots B^{y_k}_{b_k})\psi\rangle$ and use relations such as $(A^x_a)^2 = A^x_a$ and the commutation relations to place constraints that relate different coefficients together. Considering all relations that hold on coefficients obtained from products of length at most n yields a sequence of finite optimization problems with supremum $\omega^c_{\leq n}$ such that $(\omega^c_{\leq n})_{n\geq 1}$ is a nonincreasing sequence that can be shown to converge to the optimum of (5) from above.

We have devised two procedures: the first returns a sequence of values converging to (4) from below, and the second converges to (5) from above. Although we have not discussed the runtime of the procedures, as it does not matter for the general argument, it can be shown that both run in time $\exp(\mathrm{poly}(n))$. The "from above" procedure can be further optimized by using techniques from semidefinite programming, and it is used in practice to obtain numerically tight upper bounds on small nonlocal games that arise in experiments or in cryptographic applications. Irrespective of the runtime, the existence of these procedures has an important consequence. Suppose the values (5) and (4) happen to be equal. Then we claim that there exists an algorithm that provided as input coefficients $\{\pi(x,y)V(a,b|x,y)\}$ always halts and correctly decides between the case when (5) is at least $\frac{2}{3}$, or at most $\frac{1}{3}$, provided that one of the two cases is promised to hold. Indeed, such an algorithm can be obtained by executing both procedures in parallel and stopping as soon as either the "from below" procedure returns a value that strictly

[7] *To show this rigorously, write $\psi \in \mathcal{H}_A \otimes \mathcal{H}_B$ using the "Schmidt decomposition" as $\psi = \sum_i d_i\, u_i \otimes v_i$ with $d_1 \geq d_2 \geq \cdots$ and consider projections P_j and Q_j on $\{u_i : 1 \leq i \leq j\}$ and $\{v_i : 1 \leq i \leq j\}$, respectively.*

exceeds $\frac{1}{3}$ or the "from above" procedure returns a value strictly less than $\frac{2}{3}$.

Now consider a problem L that lies in the class MIP*. By definition, this means that there is a procedure that transforms instances of the problem (such as, in the case of the graph nonisomorphism problem, pairs of graphs) into coefficients $G = \{\pi(x,y)V(a,b|x,y)\}$ such that if the instance is a positive instance (the graphs are not isomorphic), then the associated value $\omega^s(G)$ is 1 or close to it, say at least $\frac{2}{3}$; and if it is a negative instance, then $\omega^s(G)$ is much smaller, say less than $\frac{1}{3}$. Executing the algorithm described above on G it follows that membership in L can be decided by an algorithm that always halts. In the language of complexity theory, this means that every problem in MIP* is decidable. It is worth explicitly formulating this finding:

> If $\overline{Q^s_{ABXY}} = Q^c_{ABXY}$, then the class MIP* contains only decidable problems.

How likely is this consequence to hold? In a major breakthrough two years ago, Slofstra [17] showed that if one were to remove the $\frac{2}{3}/\frac{1}{3}$ promise, then there would exist games such that the question "$\omega^s(G) = 1$ or $\omega^s(G) < 1$" is undecidable. (As mentioned earlier, leading up to this result Slofstra was able to prove false Tsirelson's "fact" from [18].) In such a case, however, the algorithm described above does not work, because there is no finite gap $\delta > 0$ such that the algorithm can safely stop in case the "from below" procedure obtains a value that exceeds $1 - \delta$. The best result known to date on the class MIP* is the aforementioned inclusion of problems in NEEXP. While this remains a far cry from undecidable languages, NEEXP is such a gargantuan class that extending the result to NEEEXP, and more, all the way to undecidable languages, is an enticing research program that could lead to a complexity-theoretic refutation of Tsirelson's problem and by extension of Connes' conjecture. We ambitiously formulate it as a conjecture:

> *Conjectured complexity-theoretic counterpoint to Tsirelson's problem:*
> The class MIP* contains undecidable problems.

Interestingly, a proof of the complexity-theoretic conjecture could be obtained without exhibiting any of the objects that CEC or Tsirelson's problem posit to not exist, such as a type II_1 algebra with no finite-dimensional approximations. In any case we may still be far from such a refutation; indeed, CEC may still hold its ground, so that MIP* would be decidable. In complexity-theoretic land this arguably would be a much more expected and comfortable situation. Isn't it rather "spooky" when undecidability crops up in a "reasonable" model of computation?

ACKNOWLEDGMENTS. I am indebted to Volkher Scholz, William Slofstra, Henry Yuen, and the *Notices* referees for multiple comments that greatly improved the presentation of this article.

References

[1] Aravind PK. Quantum mysteries revisited again, *Amer. J. Phys.*, no. 10 (72):1303–1307, 2004, DOI 10.1119/1.1773173. MR2086837

[2] Babai L, Fortnow L, Lund C. Nondeterministic exponential time has two-prover interactive protocols, *Comput. Complexity*, no. 1 (1):3–40, 1991, DOI 10.1007/BF01200056. MR1113533

[3] Buchholz D. Product states for local algebras. *Comm. Math. Phys.*, no. 4, (36):287–304, 1974. MR0345546

[4] Capraro V, Lupini M. *Introduction to sofic and hyperlinear groups and Connes' embedding conjecture*, with an appendix by Vladimir Pestov, Lecture Notes in Mathematics, vol. 2136, Springer, Cham, 2015. MR3408561

[5] Cleve R, Hoyer P, Toner B, Watrous J. Consequences and limits of nonlocal strategies, *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004*, 2004, 236–249.

[6] Connes A. Classification of injective factors. Cases II_1, II_∞, III_λ, $\lambda \neq 1$, *Ann. of Math. (2)*, no. 1 (104):73–115, 1976, DOI 10.2307/1971057. MR0454659

[7] Fritz T. Tsirelson's problem and Kirchberg's conjecture, *Rev. Math. Phys.*, no. 5 (24):1250012, 2012, DOI 10.1142/S0129055X12500122. MR2928100

[8] Goldbring I, Hart B. Computability and the Connes embedding problem, *Bull. Symb. Log.*, no. 2 (22):238–248, 2016, DOI 10.1017/bsl.2016.5. MR3532694

[9] Hensen B et al. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres, *Nature*, no. 7575 (526):682, 2015

[10] Ito T, Vidick T. A multi-prover interactive proof for NEXP sound against entangled provers, 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012, 2012, 243–252. MR3186611

[11] Junge M, Navascues M, Palazuelos C, Perez-Garcia D, Scholz VB, Werner RF. Connes embedding problem and Tsirelson's problem, *J. Math. Phys.*, no. 1 (52):012102, 2011, DOI 10.1063/1.3514538. MR2790067

[12] Kirchberg E. On nonsemisplit extensions, tensor products and exactness of group C^*-algebras, *Invent. Math.*, no. 3 (112):449–489, 1993, DOI 10.1007/BF01232444. MR1218321

[13] Natarajan A, Wright J. *NEEXP in MIP**. Technical report, arXiv:1904.05870, 2019.

[14] Navascués M, Pironio S, Acín A. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations, *New J. Phys.*, no. 7 (10):073013, 2008

[15] Ozawa N. About the Connes embedding conjecture: algebraic approaches, *Jpn. J. Math.*, no. 1 (8):147–183, 2013, DOI 10.1007/s11537-013-1280-5. MR3067294

[16] Rădulescu F. A non-commutative, analytic version of Hilbert's 17th problem in type II_1 von Neumann algebras,

Von Neumann algebras in Sibiu; Theta Ser. Adv. Math., 10, Theta, Bucharest, 2008:93–101. MR2512326
[17] Slofstra W. The set of quantum correlations is not closed, *Forum Math. Pi* (7):e1, 41 pp., 2019, DOI 10.1017/fmp.2018.3. MR3898717
[18] Tsirelson B S. Some results and problems on quantum Bell-type inequalities, *Hadronic J. Suppl.*, no. 4 (8):329–345, 1993. MR1254597
[19] Voiculescu D. The analogues of entropy and of Fisher's information measure in free probability theory. II, *Invent. Math.*, no. 3 (118):411–440, 1994, DOI 10.1007/BF01231539. MR1296352
[20] von Neumann J. *Mathematische Grundlagen der Quantenmechanik*, Unveränderter Nachdruck der ersten Auflage von 1932. Die Grundlehren der mathematischen Wissenschaften, Band 38, Springer-Verlag, Berlin-New York, 1968 (German). MR0223138

Thomas Vidick

Credits

Figures 1 and 2 are courtesy of the author.
Photo of the author is courtesy of Matt Scobel.