

Games in the matrix

Ilijas Farah^[1]

Do all the possible computations on very large matrices look alike? Nobody knows, but let's play a game and see how far we get!

1. Games

Let's start with one of the simplest interesting mathematical structures. A *simple graph* is a set of “vertices” some of which are connected by an “edge”. This can be illustrated conveniently as in the examples in Figure 1. We will consider only simple graphs and refer to them simply as “graphs”. We write



Figure 1: The left figure shows a basic example for a graph on 5 vertices. As the name suggests, edges are drawn as lines which connect the vertices. On the right we have the “cycle graph” C_6 .

$G = (V, E)$, where V is the set of vertices and E is the set of edges. If two vertices $v, w \in V$ are connected by an edge, we write $v E w$. To compare two graphs, we will play a game. Given graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ and a natural number $n \geq 1$, the game $\Gamma(G_1, G_2, n)$ has two players called “Challenger” and “Duplicator” and is played in n innings according to the following rules. In each inning Challenger chooses a vertex in either one of the

^[1] Partially supported by NSERC.

graphs. Duplicator responds by choosing a vertex in the other graph. Thus, in the j -th inning a vertex a_j in V_1 and a vertex b_j in V_2 are chosen. We add the requirement that each vertex may be chosen at most once. After n innings, Duplicator wins if all i and j satisfy $a_i E_1 a_j$ if and only if $b_i E_2 b_j$; otherwise Challenger wins. In essence, Duplicator is trying to show the two graphs are structurally similar, while Challenger is trying to expose a structural difference between them.

Let's get some easy observations out of the way. If we choose G_1 and G_2 to describe the same graph, then Duplicator can win the game by “duplicating” Challenger's moves. If both graphs have n vertices, then Duplicator can win the game only if the graphs coincide.

Each game considered in this note is a finite game of perfect information, meaning that each player has complete knowledge of the history of the game up to their current turn. For example, both chess and Go are finite games of perfect information, but each of these games can end in a draw. In this note, we do not consider games in which a draw is possible. This implies that one of the players has a winning strategy. A *winning strategy* in $\Gamma(G_1, G_2, n)$ for Duplicator is a description of which move exactly Duplicator should play in each position of the game such that by following this description Duplicator wins the game regardless of how Challenger plays. A winning strategy for Challenger is defined analogously.

We move on to a (somewhat) nontrivial example of a winning strategy for Duplicator. For $m \geq 3$ let C_m be the *cycle graph*, whose vertices are $\{1, \dots, m\}$ and two vertices i, j are connected by an edge if and only if $|i - j| = 1$ or $\{i, j\} = \{1, m\}$. Figure 1 shows the cycle graph on 6 vertices. Then, Duplicator wins the game $\Gamma(C_m, C_k, n)$ if m and k are sufficiently large compared to n – we want the smaller value of m and k to be greater than 2^{n+1} – by using the following strategy.

To simplify matters, we can assume that Challenger and Duplicator each play in only one of the graphs for the entire game. Dropping this assumption would only need Duplicator to be more conscious about their moves but the strategy would in principle still be the same. In the first inning, Challenger plays any vertex a_1 and Duplicator plays any vertex b_1 . After the second inning, the vertices a_1, a_2 and b_1, b_2 each split their graphs into two arcs. Duplicator ensures that either the two arcs are “long” (there are at least $2^n = \frac{2^{n+1}}{2}$ intermediate vertices), or the corresponding short arcs have the same length. From this point on, Duplicator maintains the following idea: small distances are matched exactly, and large distances are kept large. To be more precise, in the k -th inning there are three possible situations. First, if Challenger plays on a short arc, Duplicator chooses the corresponding vertex on the matching short arc. Second, if Challenger plays on a long arc but close (at most 2^{n-k} intermediate




Figure 2: The position after five innings. Since m and k are very large, we zoomed out so much that the vertices blended into two continuous circles. The corresponding short (red) arcs are of the same length, while the long (blue) arcs are simply sufficiently long. See [4] for this and other details.

vertices, half the previous turn’s bound for largeness) to an already chosen vertex, Duplicator matches that distance. Last, if Challenger plays far from all chosen vertices, Duplicator also chooses a distant vertex – at least 2^k steps to the nearest already-chosen vertex. By applying this strategy, Duplicator keeps the distances between the chosen vertices great enough and after n innings all played vertices satisfy $a_i E_1 a_j$ if and only if $b_i E_2 b_j$, which means Duplicator wins. The strategy is illustrated in Figure 2.

This tells us that short games cannot distinguish long cycles. It requires a syntactical definition to make this precise.^[2] A *first-order statement* about graphs is expressed by “quantification” over the vertices,^[3] and it involves the adjacency relation E and the equality relation $=$. For this purpose we have to introduce some mathematical (or rather logical) notation: First of all, there is the *universal quantifier* $\forall$ which reads as “for all” and the *existential quantifier* $\exists$ which reads as “there exists”. Since we are interested in first-order statements about graphs quantified over the vertices, $\forall x$ reads as “for all vertices x ” and similarly $\exists x$ reads as “there exists a vertex x ”. Furthermore, we need the logical connectives $\wedge$ (“and”), $\vee$ (“or”), $\neg$ (“not”), and $\rightarrow$ (“implies”). For example,

$$\varphi_1: (\exists x)(\forall y)y \neq x \rightarrow y E x$$

reads as “there is a vertex x such that for all vertices y , $y \neq x$ implies that y and x

^[2] A mathematical logician is a mathematician who takes the syntax seriously.

^[3] A second-order statement is one in which quantification over arbitrary subsets of the structure is allowed. It is considerably more expressive than the first-order logic, which lies in the “Goldilocks zone” in terms of expressiveness. In the final section we will see some evidence for this.

are connected by an edge” and hence asserts that some vertex is adjacent to all other vertices. Now try it yourself! Can you translate the following statement?

$$\varphi_2: (\forall x)(\forall y)x \neq y \rightarrow (\exists z)z E x \wedge \neg(z E y) \quad [4]$$

But what does this have to do with games? Well, everything. But before we go into the details, have a look at the two graphs in Appendix A. Play the game on the two graphs in three innings. Can you find a winning strategy for either Challenger or Duplicator?

Are you back? Here is the solution. For any two graphs G_1 and G_2 , if G_1 satisfies φ_2 and G_2 does not (check that this is true for the two graphs in Appendix A), then Challenger wins $\Gamma(G_1, G_2, 3)$ as follows. In the first two innings, Challenger chooses a_1 and a_2 in G_2 so that no z satisfies $z E a_1$ and $\neg(z E a_2)$. No matter how b_1 and b_2 are chosen, in the third inning we can choose b_3 so that $b_3 E b_1$ and $\neg(b_3 E b_2)$, and Duplicator is doomed.

Now have a look at the graphs in Appendix B. Play the game on the two graphs with varying numbers of innings. How many innings are necessary for a winning strategy for Challenger to exist? [5]

For every first-order statement ψ there is n large enough so that there is a winning strategy for Challenger in the game $\Gamma(G_1, G_2, n)$ if G_1 and G_2 disagree on the truth of ψ . The least such n is the *quantifier depth* of ψ ; let’s just say that the name is well chosen instead of giving a proper definition. Conversely, if G_1 and G_2 agree on the truth of all first-order statements of quantifier depth n , then there is a winning strategy for Duplicator in the game $\Gamma(G_1, G_2, n)$.

These are the Ehrenfeucht–Fraïssé games, affectionately known as the EF-games (see [7, §3.3]). Their variations can be used to compare mathematical structures of all sorts: linear orderings, groups, rings, and many more. In each case, EF-games provide efficient means for verifying whether two given structures satisfy the same first-order statements. Before moving on, we state another fun fact about games on graphs: every sequence of graphs has a subsequence such that for every fixed number of innings n , any two graphs in it that are “far enough down the road” are indistinguishable by the n -inning EF-game – that is, there is a winning strategy for Duplicator.

2. Matrices

The same idea – comparing structures by playing games – now applies to a very different kind of object: matrices. We start with the necessary background.

[4] Solution: “For all vertices x and for all vertices y , $x \neq y$ implies that there exists a vertex z such that an edge connects z and x and no edge connects z and y .” This means that for any two distinct vertices x, y some vertex is connected to x but not to y .

[5] The answer is 4.

During a legendary 1925 vacation in Heligoland, Werner Heisenberg (1901–1976) noticed that because of the noncommutativity of matrix multiplication, matrix algebras provide a natural setting for quantum mechanics. A few years later, John von Neumann (1903–1957) pointed out that the infinite-dimensional analog of matrix algebras is more suitable for this purpose, and suggested certain algebras of operators as a model. We will see an example of such an algebra in the final section. Let’s first take a closer look at what matrix algebras are all about.

Consider the (orthogonal) projection p to the line $x = y$, or the rotation ρ_α of the plane by an angle α around the origin (Figure 3). These transformations




Figure 3: The projection to the line $x = y$ and the rotation by $\frac{\pi}{2}$ around the origin as applied to points in the plane.

are implemented by matrices, via

$$p\left(\begin{pmatrix} a \\ b \end{pmatrix}\right) = \begin{pmatrix} a/2 + b/2 \\ a/2 + b/2 \end{pmatrix} = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix},$$

$$\rho_\alpha\left(\begin{pmatrix} a \\ b \end{pmatrix}\right) = \begin{pmatrix} \cos(\alpha)a - \sin(\alpha)b \\ \sin(\alpha)a + \cos(\alpha)b \end{pmatrix} = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}.$$

In case you haven’t seen it yet, the product of a 2×2 matrix and a 2-vector is defined as follows: the first entry of the product of a matrix and a vector is the dot product of the first row of the matrix with the vector, and its second entry is the dot product of the second row of the matrix with the vector. This works the same in the higher-dimensional context, for example when $n = 3$:

$$\begin{pmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^3 c_{1i}a_i \\ \sum_{i=1}^3 c_{2i}a_i \\ \sum_{i=1}^3 c_{3i}a_i \end{pmatrix}$$

To multiply two $m \times m$ matrices, a and b , one applies this formula to each column of b and produces the corresponding column of the result, ab .

Fix $m \geq 2$ and let $\mathbb{C}^m$ be the space of m -tuples (called m -vectors) of complex numbers. This is a *vector space* over $\mathbb{C}$; what this means is that we can add vectors pointwise and multiply them by a complex number (called a *scalar*). Complex matrices act on complex vectors by matrix multiplication – those are the transformations of vectors we are interested in. For the upcoming, we will also need a third operation, the *adjoint* a^* of a matrix a : To compute it, transpose the matrix (flip the entries of the matrix over its diagonal) and take complex conjugates of its entries. Thus, $((a_{ij})_{i,j \leq m})^* = (\overline{a_{ji}})_{i,j \leq m}$. For example, the adjoint of the matrix that corresponds to a rotation by an angle α is the matrix that implements the rotation by the angle $-\alpha$.

By $\mathbb{M}_m$ we denote the *algebra of all complex $m \times m$ matrices* $a = (a_{ij})_{i,j \leq m}$, that is, we equip the matrices with multiplication, addition, adjoint, and multiplication by scalars. By 1_m we denote the *unit matrix* that has 1's down the main diagonal and 0's elsewhere; in other words, $a_{ij} = 1$ if $i = j$ and $a_{ij} = 0$ otherwise. It has the property that $1_m a = a 1_m = a$ for all a in $\mathbb{M}_m$.

2.1. Let the games begin

For l, m , and n one could consider the following game: Our two players choose elements from $\mathbb{M}_l$ and $\mathbb{M}_m$ as they did for graphs, and Duplicator wins if the chosen matrices $a_1, \dots, a_n$ and $b_1, \dots, b_n$ “look the same”. One could insist on exact equality, but this turns out to be too strict. Strict equality is fine for graphs and other discrete structures, but when in the matrix there is no need for exactness.^[6] Let's make this precise. The length of an m -vector ξ with coordinates $x_1, \dots, x_m$ is

$$\|\xi\| = \sqrt{\sum_{i=1}^m |x_i|^2}.$$

This is justified by the Pythagorean theorem if $m = 2$ and generalizes to higher dimensions. There are at least two common ways to define the length (called “norm”) of a matrix. First, every $a = (a_{ij})_{i,j \leq m}$ in $\mathbb{M}_m$ can be treated like an m^2 -vector, and its *normalized Hilbert–Schmidt* norm is

$$\|a\|_{\text{HS}} = \sqrt{\frac{1}{m} \sum_{i,j=1}^m |a_{ij}|^2}.$$

It is “normalized” because of the factor $\frac{1}{m}$, assuring that $\|1_m\|_{\text{HS}} = 1$.

The *operator norm* may be slightly less intuitive and it is notoriously difficult to compute for an arbitrary matrix, but it is in some sense even more natural than $\|\cdot\|_{\text{HS}}$. It is defined as the “maximal stretching coefficient”: $\|a\|_{\text{OP}} = r$ means that $\|a\xi\| = r\|\xi\|$ for some vector ξ and no other vector is stretched further by a .

[6] Please don't read much into this sentence.

Let the correct games begin. For l, m, n , and $\varepsilon > 0$, we will consider the game $\Gamma_{\text{HS}}(\mathbb{M}_l, \mathbb{M}_m, n, \varepsilon)$ which is played like the strict version, with two differences:

- (P1) Players can choose only matrices x that satisfy $\|x\|_{\text{OP}} \leq 1$.^[7]
(P2) Duplicator wins if and only if for all i, j, k and all scalars y, z such that $\max(|y|, |z|) \leq 1$, each of

$$\begin{aligned} & \|a_i\|_{\text{HS}} - \|b_i\|_{\text{HS}}, \\ & \|a_i a_j - a_k\|_{\text{HS}} - \|b_i b_j - b_k\|_{\text{HS}}, \\ & \|y a_i + z a_j - a_k\|_{\text{HS}} - \|y b_i + z b_j - b_k\|_{\text{HS}}, \\ & \|a_i^* - a_j\|_{\text{HS}} - \|b_i^* - b_j\|_{\text{HS}} \end{aligned}$$

is no greater than ε .

This list of conditions means that while the matrices are not strictly equal, they look almost the same – up to a small “error” of size ε . The game $\Gamma_{\text{OP}}(\mathbb{M}_l, \mathbb{M}_m, n, \varepsilon)$ is defined analogously, with the payoff in (P2) computed using $\|\cdot\|_{\text{OP}}$ instead of $\|\cdot\|_{\text{HS}}$.

For fixed m and l , as n increases or ε decreases, the game becomes more difficult for Duplicator. For smaller ε , there are fewer matrices satisfying (P2) for Duplicator to choose from. As in the case of graphs, Duplicator wins $\Gamma_{\text{HS}}(\mathbb{M}_l, \mathbb{M}_m, n, \varepsilon)$ and $\Gamma_{\text{OP}}(\mathbb{M}_l, \mathbb{M}_m, n, \varepsilon)$ if $\mathbb{M}_l$ and $\mathbb{M}_m$ agree (up to ε) on the values of “continuous first-order statements” – evaluation of continuous first-order statements results in real numbers instead of truth values; see [2] and [6] – of quantifier depth n .

A first result for winning strategies for matrix games is that for all sufficiently small $\varepsilon > 0$ and all even m and odd l , there is a winning strategy for Challenger in the game $\Gamma_{\text{OP}}(\mathbb{M}_l, \mathbb{M}_m, 6, \varepsilon)$. Just like in the discrete case, the existence of a winning strategy can be reformulated as asserting that some first-order statement is evaluated differently at $\mathbb{M}_m$ and $\mathbb{M}_l$. The first-order statement associated with the winning strategy for Challenger in the game $\Gamma_{\text{OP}}(\mathbb{M}_l, \mathbb{M}_m, 6, \varepsilon)$ is

$$\psi: \inf_{\|x\|_{\text{OP}} \leq 1} \max(\|x^* x - (x^* x)^2\|_{\text{OP}}, \|x^* x + x x^* - 1\|_{\text{OP}}).$$

This statement distinguishes whether or not there are matrices in $\mathbb{M}_l$ and $\mathbb{M}_m$ which are a certain kind of “partial isometry”. Roughly, a partial isometry is a transformation that preserves distances on part of the space and collapses the rest to zero, and ψ asserts that in addition the dimensions of these two parts are equal. Such matrices only exist in $\mathbb{M}_m$ if m is even. This implies that the value of ψ in $\mathbb{M}_n$ is 0 if n is even and $> 10^{-10^{10}}$ if n is odd.^[8]

^[7] The subscript OP is not a typo. The correct ambient for our purposes is the operator unit ball of $\mathbb{M}_m$, regardless of which one of the norms we consider.

^[8] There are better estimates but I am just a set theorist.

In other words, the collection of true first-order statements of $\mathbb{M}_m$ “knows” the parity of m . Actually, the true first-order statements of $\mathbb{M}_m$ even “know” the divisors and the value of m . The associated winning strategies for Challenger rely on a modification of ψ .

How different is Γ_{HS} from Γ_{OP} ? Can the Challenger distinguish the parities of the dimensions in this game? Well, no. For all n and $\varepsilon > 0$ and all sufficiently large m , there is a winning strategy for Duplicator in the game $\Gamma_{\text{HS}}(\mathbb{M}_m, \mathbb{M}_{m+1}, n, \varepsilon)$. In words: with respect to the HS-norm, matrix algebras of even and odd dimension look almost the same. A similar observation applies to divisibility of the dimension by any given $k \geq 3$. What, if any, properties of large matrix algebras can be detected by their $\|\cdot\|_{\text{HS}}$ -first order theories? Nobody knows.

Question 1. *Given n and $\varepsilon > 0$, is there a winning strategy for Duplicator in the game $\Gamma_{\text{HS}}(\mathbb{M}_l, \mathbb{M}_m, n, \varepsilon)$ whenever m and l are large enough?*

A positive answer to this question would imply that sufficiently large matrices equipped with the HS-norm cannot be distinguished by first-order statements. We will return to this point in Section 3

2.2. Normalizers

If you can't solve a problem, then there is an easier problem you can solve: find it. George Pólya (1887–1985)

Let's change the game by requiring the players to play only permutation matrices. A *permutation matrix* is a matrix all of whose entries are equal to 0 or to 1 that has exactly one 1 in each row and each column. We denote the set of all $m \times m$ permutation matrices by $\mathcal{S}_m$.

Let $\Gamma(\mathcal{S}_m, \mathcal{S}_l, n, \varepsilon)$ be defined as follows.

- (N1) The players alternate choosing elements of $\mathcal{S}_m$ and $\mathcal{S}_l$, as before.
- (N2) Duplicator wins if and only if the maximal value of $\|a_i a_j - a_k\|_{\text{HS}} - \|b_i b_j - b_k\|_{\text{HS}}$ for all i, j, k is smaller than ε , that is, with respect to matrix multiplication, the chosen matrices in $\mathcal{S}_m$ and $\mathcal{S}_l$ look almost the same.

This game is more similar to $\Gamma_{\text{HS}}(\mathbb{M}_m, \mathbb{M}_l, n, \varepsilon)$ than it may appear: the restriction to permutation matrices allows us to play an interesting game, even if we reduce the conditions on looking similar only to similarity with respect to matrix multiplication. We can now answer the analog of Question 1.

Theorem. *There are n and $\varepsilon > 0$ such that there are arbitrarily large m and l for which there is a winning strategy for Challenger in the game $\Gamma(\mathcal{S}_m, \mathcal{S}_l, n, \varepsilon)$.*

This is a consequence of [1, Theorem 5.6], although this is not at all obvious from what we have seen so far. The proof involves some fascinating ideas that we have no time to discuss such as sofic groups, expander graphs, Ulam-stability of approximate homomorphisms, comparison of first-order theories of alternating groups of order $p^4 - 1$ for primes $p \geq 7$, and something that awaits us in the final section. It is time to take the red pill,^[9] step out of the matrix, and take a bird's-eye view.

3. The really, really, really big picture

I have lied to you. Question 1 is just a shadow of Question 2, which was posed by Sorin Popa. A negative answer to Question 1 implies a negative answer to Question 2.^[10]

Question 2. *Are all tracial ultraproducts of matrix algebras isomorphic?*

It seems that I have a lot of explaining to do (pun intended). An (ultra)filter on $\mathbb{N}$ provides a systematic way to distinguish between subsets of $\mathbb{N}$ that are considered “large” and those that are considered “small”. More precisely, a *nonprincipal ultrafilter on $\mathbb{N}$* (all ultrafilters in this snapshot are nonprincipal and on $\mathbb{N}$) is a family of subsets $\mathcal{U}$ of $\mathbb{N}$, whose elements we think of as the *large subsets of $\mathbb{N}$* , that satisfies the following conditions.

- (U1) For all A, B in $\mathcal{U}$, also their intersection $A \cap B$ is in $\mathcal{U}$.
- (U2) For all $A \in \mathcal{U}$ and $A \subseteq C \subseteq \mathbb{N}$, also C is in $\mathcal{U}$.

The first condition says that the intersection of two large subsets of $\mathbb{N}$ is still large. The second condition says that any set containing a large set is itself large. Together, these conditions define a *filter on $\mathbb{N}$* .

- (U3) For every $A \subseteq \mathbb{N}$ exactly one of A or $\mathbb{N} \setminus A$ belongs to $\mathcal{U}$.

This condition strengthens a filter on $\mathbb{N}$ to an *ultrafilter on $\mathbb{N}$* : every subset of $\mathbb{N}$ is classified as either large or small (then its complement is large) – there is no other option.

- (U4) No finite set belongs to $\mathcal{U}$.

This final condition says that all finite sets are considered small and makes the ultrafilter *nonprincipal*.

^[9] After reading this far you may have the impression that it was written by a fan of the Matrix movies. This is not true at all; I just had to grab your attention somehow. Sorry.

^[10] Is the converse true? Who knows. Read on.

Ultrafilters provide a way to describe the long-term behavior of a sequence of numbers (and also other mathematical objects). Suppose we have a sequence of numbers $x_1, x_2, \dots$. Informally, we could ask if the values of the sequence eventually stay close to some certain number x . An ultrafilter allows us to make this question precise: We think of an ultrafilter as a gadget that tells us which sets of indices are large and which are small (and thus negligible). We then say that the sequence *converges* to x along the ultrafilter, if for every $\varepsilon > 0$ the set of indices $\{n: |x - x_n| < \varepsilon\} \subset \mathbb{N}$ is large, that is, an element of the ultrafilter $\mathcal{U}$. In plain words: according to the ultrafilter, the values of the sequence are very often close to x . A remarkable property of ultrafilters is that every bounded sequence converges to a unique number for a given ultrafilter. This idea extends beyond numbers. For example, if you have a sequence of matrices in $\mathbb{M}_m$ which are bounded (referring to their size with respect to a norm as described in the previous section), then there is also a unique matrix to which the matrices converge along an ultrafilter.

Before we move on with an explanation of Question 2 and how it is related to nonprincipal ultrafilters, let's return to graphs to get a feel for “ultraproducts” and their exciting properties. Given a sequence $G_1, G_2, \dots$ of graphs and an ultrafilter $\mathcal{U}$ on $\mathbb{N}$, their *ultraproduct* with respect to $\mathcal{U}$, denoted by $\prod_{\mathcal{U}} G$ is the graph which consists of the following data: We start by picking a vertex a_n from each graph G_n . This results in a sequence of vertices $(a_1, a_2, \dots)$. Now we compare two sequences $(a_1, a_2, \dots)$ and $(b_1, b_2, \dots)$ and say that they are the same, if they agree in “most” of the components – this means, we want $\{n: a_n = b_n\} \in \mathcal{U}$. As a vertex of $\prod_{\mathcal{U}} G$ we hence choose sequences of vertices of the respective graphs and consider two such sequences the same if they are the same in most of the components. Analogously we define the edges of $\prod_{\mathcal{U}} G$. Two vertices a and b (that is, sequences of vertices of the respective graphs) are connected by an edge, if most of their components are connected by an edge, meaning $\{n: a_n E_n b_n\} \in \mathcal{U}$, where E_n denotes the edge relation of G_n .

For every first-order statement φ about graphs, φ holds in $\prod_{\mathcal{U}} G$ if and only if the set $\{n: \varphi \text{ holds in } G_n\}$ belongs to $\mathcal{U}$ – that is, it holds in most of the respective graphs. This is Łoś’s Theorem (named after Jerzy Łoś (1920–1998)), rightfully known as the Fundamental Theorem of Ultraproducts. As we saw in Section 1, in a sequence of graphs, the true first-order statements of most graphs in the sequence are indistinguishable and therefore in particular the ultraproduct $\prod_{\mathcal{U}} G$ does not depend on the choice of $\mathcal{U}$. All of this applies to other discrete mathematical structures such as linear orderings, groups, rings, and many more, but I owe you an explanation of Question 2.

Back on track, let $\mathbb{M} = \{(a_1, a_2, \dots): a_m \in \mathbb{M}_m \text{ and } \|a_m\|_{\text{OP}} < \infty \text{ for all } m\}$. We equip $\mathbb{M}$ with coordinatewise multiplication, addition, adjoint and multiplication by scalars. As for graphs, however, we need to find a weaker notion

of equality of the sequences in $\mathbb{M}$. We say that two sequences of matrices are the same, if “most” of their components are close to one another. Closeness is measured with respect to $\|\cdot\|_{\text{HS}}$ and “most” is captured by the ultrafilter $\mathcal{U}$. This results in the *tracial ultraproduct* $\prod^{\mathcal{U}} \mathbb{M}_m$. This is an infinite-dimensional analog of the matrix algebras M_m . The continuous version of Łoś’s Theorem implies that the answer to Question 1 is positive if and only if the true first-order statements of $\prod^{\mathcal{U}} \mathbb{M}_m$ do not depend on the choice of $\mathcal{U}$.

By [5], the negation of Cantor’s Continuum Hypothesis (CH) implies a negative answer to Question 2: different choices of ultrafilters, that is, different ways of deciding which matrix dimensions are considered as large, can produce structurally different tracial ultraproducts. CH is a famous statement about the sizes of infinite sets, shown to be neither provable nor disprovable from the standard axioms of mathematics. Instead of offering a detour on CH this late in the snapshot, I recommend that the next time you meet a set theorist you ask whether CH is true. You may hear answers that are about as opinionated as they are different.^[11] The answer to Question 1 does not depend on CH or other standard additional set-theoretic axioms; we just don’t know what it is.

Acknowledgements

I am indebted to Dima Shlyakhtenko for suggesting that I write something like this and for helpful remarks on the first draft of this paper, to Diego Bejarano and Andrea Vaccaro for noticing a few typos, and to Vadim Alekseev for two insightful suggestions that greatly improved the exposition.



^[11] Also see the Paul Erdős (1913–1996) anecdote in [3, p. 251]. This book also happens to be my favourite reference for ultraproducts, countable saturation, and whatnot ☺.

References

- [1] V. Alekseev and A. Thom, *On non-isomorphic universal sofic groups*, arXiv:2406.06741v3, 2026.
- [2] I. Ben Yaacov, A. Berenstein, C.W. Henson, and A. Usvyatsov, *Model theory for metric structures*, Model Theory with Applications to Algebra and Analysis (Z. Chatzidakis, D. Macpherson, A. Pillay, and A. Wilkie, eds.), vol. 2, London Math. Soc. Lecture Notes Series, no. 350, Cambridge University Press, 2008, pp. 315–427.
- [3] I. Farah, *Combinatorial set theory and C^* -algebras*, 2nd ed., Springer Monographs in Mathematics, Springer, 2026.
- [4] ———, *Games in the matrix: Director’s cut*, arXiv:2511.07456v2, 2026.
- [5] I. Farah, B. Hart, and D. Sherman, *Model theory of operator algebras II: Model theory*, Israel J. Math. **201** (2014), 477–505, <https://doi.org/10.1007/s11856-014-1046-7>.
- [6] B. Hart, *An introduction to continuous model theory*, Model theory of operator algebras (I. Goldbring, ed.), De Gruyter, 2023, pp. 83–132.
- [7] W. Hodges, *Model theory*, Encyclopedia of Mathematics and its Applications, vol. 42, Cambridge university press, 1993.



A. Graphs for a game

Two graphs for a three-inning game. Can you find a winning strategy for Challenger or Duplicator?



B. Graphs for another game

Two graphs to play games on. How many innings are necessary for a winning strategy for Challenger to exist?



Ilijas Farah is a research chair on
foundations of operator algebras at York
University.

Mathematical subjects
Algebra and Number Theory, Analysis,
Discrete Mathematics and Foundations

License
Creative Commons BY-SA 4.0

DOI
10.14760/SNAP-2026-011-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Junior Editor
Victor Bloch
junior-editors@mfo.de

Senior Editor
Anja Randecker
senior-editor@mfo.de

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9–11
77709 Oberwolfach
Germany

Director
Gerhard Huisken



Mathematisches
Forschungsinstitut
Oberwolfach



IMAGINARY
open mathematics