MIP* = RE

David Wakeham

March 13, 2020

1 Correlations

1.1 Introduction

- We'll be talking about the recent paper $\mathsf{MIP}^* = \mathsf{RE}$ (2001.04383), by Ji, Natarajan, Vidick, Wright and Yuen.
- This means that the yes-no questions you can reliably answer by quickly chatting to entangled provers are the same questions which a Turing machine, if the answer is 'yes', will eventually say 'yes' to. This is dramatically more powerful than anyone expected, and according to people who know about these things, it is probably the important result in computational complexity theory this century (so far).
- But I'm not a complexity theorist, and I guess neither are you. Why should we care? It turns out that this result tells us deep facts about entanglement, quantum mechanics, and the theory of operator algebras, so I will focus on these.

1.2 CHSH inequality

- Our path starts with a simple calculation. Say Alice and Bob each have a qubit, and a set of operators A_0, A_1, B_0, B_1 with outcomes taking values ± 1. Alice and Bob are going to do their measurements separately and without communicating, so we require that the A and B commute.
- We now ask the same question John Bell asked: can I use the pattern of bipartite correlation, which I can think as a joint probability distribtion $P(ab)$ over measurement outcomes a and b, to rule out local hidden variable theories? In other words, can I show the world isn't classical?
- As Bell found, the answer is yes: he computed some bounds on classical correlations which Alice and Bob can violate with quantum mechanics. In the simple two-qubit case, the bound is called the *CHSH inequality*. If we define

$$\mathcal{B} = A_0 B_0 + A_0 B_1 + A_1 B_0 - A_1 B_1,$$

then squaring gives

$$\mathcal{B}^2 = 4I - [A_0, A_1][B_0, B_1].$$

Classical measurements themselves commute, i.e. $[A_0, A_1] = [B_0, B_1] = 0$, so we immediately get a bound on the expectation of $\mathcal{B}$:

$$|\langle \mathcal{B} \rangle| \leq \sqrt{\langle \mathcal{B}^2 \rangle} = 2.$$

- We can think of any pattern of correlations $P(ab)$ as a point on the hypercube $[0,1]^{16}$ (technically the 15-dimensional hyperplane of probability distributions), since ab can take 16 values. The classical correlations are easy to describe. We start with deterministic assignments to A_i, B_j. By giving Alice and Bob access to shared classical randomness, you enlarge this set of points to its convex hull. The faces of this polytope saturate the CHSH inequality.

- This situation immediately generalises to a finite number of observables and outcomes: the classical correlations are given by the convex hull of deterministic assignments, with Bell inequalities on the faces. We'll call the set of classical bipartite correlations $\mathcal{C}_{nk}$, where both Alice and Bob can choose n measurements to perform on their systems, and each measurement has k outcomes. (We can allow for both more observers, and a more complicated pattern of measurements, but it won't actually affect any of our conclusions.)

1.3 Tsirelson bounds

- So, that's the classical story. But let's return to our original two-qubit scenario. If Alice and Bob share an entangled state, we can see from repeated measurements that they violate CHSH. So this experimentally rules out local hidden variable theories, or at least, the simplest ones.

- But by how much can we violate the inequality? In other words, is there a quantum Bell inequality constraining quantum correlations? It turns out there is, and we can find it by reusing our expression for $\mathcal{B}^2$ and applying Cauchy-Schwarz:

$$|\mathcal{B}^2| = |4I - [A_0, A_1][B_0, B_1]| \leq 4 + |[A_0, A_1]| \cdot |[B_0, B_1]| \leq 4 + 2|A_0||A_1||B_0||B_1| = 8.$$

  We learn that quantum mechanics obeys $|\langle \mathcal{B} \rangle| \leq 2\sqrt{2}$.

- This is called a *Tsirelson bound*. It's sharp, in the sense that you can achieve it. For instance, you can pick a Bell state $(|01\rangle - |10\rangle)/\sqrt{2}$, and the operators

$$A_0 = Z_1, \quad A_1 = X_1, \quad B_0 = -\frac{Z_2 + X_2}{\sqrt{2}}, \quad B_1 = \frac{Z_2 - X_2}{\sqrt{2}}.$$

  Proceeding in this, you can get a bunch of special points and then generate their convex hull.

- Let's consider the same generalisation as before, where Alice and Bob have n observables, each with k outcomes. We want to find the analogue of Bell inequalities, so these Tsirelson bounds, and the set of bipartite correlations allowed by quantum mechanics.

- Perhaps surprisingly, this is much, much harder, and ambiguous in a deep way. But loosely speaking, this is because, in the quantum-mechanical case, we need to talk about states as well as measurements, and talking about states requires a Hilbert space.

- The problem is hard because we're only given numbers n and k, so we need to figure the right Hilbert space or spaces for the measurements to act on. We need to think about representations. The problem is ambiguous because we have a choice about how to represent the fact that Alice and Bob make separate measurements and don't communicate.

1.4 Tensor products vs commuting operators

- Here are the two choices. First, we could act on tensor product Hilbert spaces $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_b$. This is the natural thing to do in finite-dimensional, non-relativistic quantum mechanics, for instance if Alice holds some qubits and Bob holds some qubits. We'll call the allowed set of tensor-product correlations $\mathcal{Q}_{nk}^{\mathrm{T}}$. Formally, we can define this as

$$\mathcal{Q}_{nk}^{\mathrm{T}} = \{p(ab|xy) : p(ab|xy) = \langle\psi|A_a^x \otimes B_b^y|\psi\rangle\},$$

  where x and y labels observables, a and b label outcomes, the operators $\{A_a^x : \mathcal{H}_A \to \mathcal{H}_A\}$ and $\{B_b^y : \mathcal{H}_B \to \mathcal{H}_B\}$ are projective measurements on Alice and Bob's factors, and $|\psi\rangle \in \mathcal{H}_B \otimes \mathcal{H}_B$ is a state in a tensor product Hilbert space.

- You might say this is unncessarily strong: all we really need is for Alice and Bob's measurements to commute. We'll call this set of correlations from commuting operators $\mathcal{Q}_{nk}^{\mathrm{C}}$, and we have

$$\mathcal{Q}_{nk}^{\mathrm{C}} = \{p(ab|xy) : p(ab|xy) = \langle\psi|A_a^x B_b^y|\psi\rangle\},$$

  where now the operators $\{A_a^x, B_b^y : \mathcal{H} \to \mathcal{H}\}$ are POVMs on the full Hilbert space which commute $[A_a^x, B_B^y] = 0$, and $|\psi\rangle \in \mathcal{H}$ for some Hilbert space.

- Tensor product automatically implies commuting, and quantum correlations certainly include the classical ones, so we have

$$\mathcal{C}_{nk} \subset \mathcal{Q}_{nk}^{\mathrm{T}} \subseteq \mathcal{Q}_{nk}^{\mathrm{C}}.$$

  A natural question is: are the quantum sets equal?

- For a finite-dimensional Hilbert space, it turns out that the two choices are equivalent. If Alice and Bob's measurements commute, I can always factorise the Hilbert space to respect this. But it's easy to come up with infinite-dimensional examples where Alice and Bob's operators commute, but we can't factorise the Hilbert space. The canonical example is quantum field theory, where you can't factorise the Hilbert space spatially, essentially because of short range entanglement. If I try and cut out a region around Alice, I have to sever a bunch of spatial Bell pairs.

1.5 Connes' embedding conjecture

- And in fact, you can using functional analysis to show that for infinite-dimensional Hilbert spaces, $\mathcal{Q}_{nk}^{\mathrm{C}}$ form a closed set, while the tensor product correlations $\mathcal{Q}_{nk}^{\mathrm{T}}$ don't.
- Problem solved, right? Not exactly. Tsirelson asked the subtler question: is the *closure* of tensor product correlations $\overline{\mathcal{Q}_{nk}^{\mathrm{T}}}$ equal to commuting correlations $\mathcal{Q}_{nk}^{\mathrm{C}}$? In other words, can I approximate a pattern of commuting quantum correlations arbitrarily well by a tensor product correlation?
- This may seem pedantic, but it's actually a deep question about entanglement. Take some correlations on an infinite-dimensional tensor product. You can use standard arguments about convergence to show that this can be arbitrarily well-approximated by a finite-dimensional tensor product. So the question becomes: can I approximate the patterns of correlation possible with infinite entanglement (commuting case) by finite, but arbitrarily large, amounts of entanglement (tensor product case)?
- I will just mention in passing that we can also translate this question into the language of operator algebras. The observables in commuting quantum correlations are related to an object called the Type II_1 factor von Neumann algebra. The question then becomes: can I approximate Type II_1 factors by matrices? (In operator algebra-ese, "approximate by matrices" is more properly "embed in an ultrapower of the hyperfinite Type II_1 algebra".) Fields medallist Alain Connes asked if this was the case in a footnote ~ 40 years ago. Called Connes' Embedding Conjecture, it has been one of the central unsolved problems in operator algebra theory.

1.6 NPA hierarchy

- Ok, after all this song and dance, let's say what is known about these quantum correlations. The set of commuting correlations is described by something called the *NPA hierarchy*. Consider the strings of n POVMs A_a^x, B_b^y, and their linear span $\mathcal{S}_n$. There is a set of identities $\mathcal{E}_n$ which hold for these strings simply by virtue of the fact that the A_a^x, B_b^y are projective measurements and $[A_a^x, B_b^y] = 0$.
- We can check if a pattern $p(ab)$ satisfies the $\mathcal{E}_n$ using semidefinite programming (SDP), and if it does, find a positive semidefinite matrix Γ_n called a *certificate*, whose elements are expectation values of overlaps of elements in $\mathcal{S}_n$. The positivity property is essentially a Tsirelson bound at level n, and progressively constrains the allowed commuting quantum correlations.
- We can check that p is not allowed by simply testing for each n. If it's not an allowed correlation, at some finite n we will see it violates an inequality. If it *is* allowed, the hierarchy is sufficient, but in general you need to test an infinite number of things! The process may not halt.
- The exception is when the operator has a finite-dimensional representation. When this happens, there is a way for us to read this off the certificate Γ_n, and in fact construct the

representation directly! This is the case, for instance, in our original two-qubit CHSH example.

- The main result of this NPA hierarchy is an approximation to $\mathcal{Q}_{nk}^C$ from above. As we increase n, we make the set of allowed correlations smaller, the constraints on certificates Γ_n get tighter, and the Tsirelson bounds get smaller. We can bound $\mathcal{Q}_{nk}^T$ (and hence $\mathcal{Q}_{nk}^C$) from below, using finite-dimensional approximations, but (Tsirelson's problem) we don't know if the lower bounds on $\mathcal{Q}_{nk}^C$ are sharp.

2 Complexity

2.1 Nonlocal games

- So, we are led by consideration from the consideration of correlation and entanglement to fairly deep questions about operator algebras. This suggests something interesting is going on here. To make things even more interesting, we're going to start moving in the direction of computational complexity by turning correlation into a *nonlocal game*. Rather than considering the whole joint probability distribution of Alice and Bob measurements, a referee picks x, y according to some probability distribution π. Alice and Bob perform these measurements and return outcomes a and b.

- This is a game, so to determine whether the players win or lose, the referee applies some rules in the form of a binary predicate $V(ab|xy) \in \{0, 1\}$, with 0 for failure and 1 for success. In order to maximise their chances of success, the players get to choose their measurements $\{A_a^x, B_b^y\}$ and the shared randomness or entanglement (the state $|\psi\rangle$) they have access to. We think of these choices as a *strategy*.

- The *value* of the game $G = (\pi, V, n, k)$ is the probability of success. But we can define this relative to the strategies we allow! There is the classical value, where we take the supremum of V over classical strategies $p \in \mathcal{C}_{nk}$

$$\text{val}(G) = \sup_{p \in \mathcal{C}_{nk}} \sum_{xyab} \pi(xy) V(ab|xy) p(ab|xy).$$

  Similarly, we can define the tensor product value $\text{val}^{\text{T}}(G)$ and commuting value $\text{val}^{\text{C}}(G)$.

- Since classical strategies are quantum, and tensor strategies commute, it follows immediately that

$$\text{val}(G) \leq \text{val}^{\text{T}}(G) \leq \text{val}^{\text{C}}(G).$$

  Tsirelson's problem becomes whether $\text{val}^{\text{T}}(G) = \text{val}^{\text{C}}(G)$ for all G.

2.2 Multiprover interactive proofs

- We can finally make contact with computational complexity by turning this nonlocal game into a protocol for verifying proofs. We have to change hats, and think of the referee as

the *verifier*, trying to check if some claim is valid, for instance, that two graphs are non-isomorphic. The players are now *provers*, whose goal is to convince the verifier of the statement.

- In more detail, the verifier will encode a yes-no problem into the game. Alice and Bob meet beforehand and agree on a strategy and share some randomness or entanglement, with no constraints on their computational power. Then the game starts, and the verifier randomly asks questions, and based on the answers, makes a decision using the function V. It's a "yes" if $V = 1$ and "no" if $V = 0$. Finally, we assume V is computed by some polynomial time algorithm. The *value* of the game that we discussed earlier is just the maximum probability that the provers can succeed in convincing the verifier of a "yes" answer.

- Since the verifier is asking random questions, this is a randomised proof, and since they interact with multiple provers, it's called a *multiprover interactive proof system*. The complexity class $\mathsf{MIP}(2,1)$ (where "MIP" stands for *multiprover interactive proof*, and $(2,1)$ for "two provers with one round of interaction") consists of all yes-no question that can reliably answered using a nonlocal game with classical provers.

- More precisely, I can take a problem (e.g. "are these graphs non-isomorphic"), describe the inputs via bit strings, and create *language* L consisting of all problem instances for which the answer is yes. Then $L \in \mathsf{MIP}(2,1)$ if there is an efficient mapping from problem instances z to games G_z, such that:

  - if $z \in L$, the value $\text{val}(G_z) > 2/3$ (completeness);
  - if $z \notin L$, the value $\text{val}(G_z) < 1/3$ (soundness).

- In other words, if the answer is "yes", then there is a high probability of convincing the verifier, and if the answer is no, there is a low probability.

- There are a couple of neat results I'll mention. First of all, it turns out that restricting to two provers and one round of interaction is no restriction at all:

$$\mathsf{MIP}(2,1) = \mathsf{MIP}.$$

  Second, you can think of a multiprover interactive system as a verifier taking a polynomial-length path through an exponentially tree of possible answer. You can think of the tree as a static, exponentially long proof, so if there are no provers to interact with, the verifier can read the whole in exponential time. So

$$\mathsf{MIP} = \mathsf{NEXP},$$

  the class of the exponential-time checkable proofs. This is the exponential version of NP. This tells us that multiprover interactive proofs are extremely powerful!

2.3 MIP*

- So, MIP was defined using classical provers. The class MIP* is exactly the same, except that now the provers are quantum: they can share entanglement and not just classical randomness. One of the intermediate results of the paper is that, without loss of generality, we can assume we have two provers and one round of interaction, so $\mathsf{MIP}^*(2,1) = \mathsf{MIP}^*$.

- We have to make a choice about what sort of quantum correlations we're allowing, so we choose tensor products. (The corresponding complexity class for commuting value is called $\mathsf{MIP}^{\mathrm{co}}$, and I might comment on it later.) So, MIP* is the class of languages which can be reliably verified by a classical, polynomial-time verifier interacting with multiple provers in different tensor factors of a Hilbert space. Same definition as before, but we replace $\mathrm{val}(G_z)$ with the "entangled value" $\mathrm{val}^{\mathrm{T}}(G_z)$.

- The relationship between MIP and MIP* is actually unclear. Consider a language $L \in \mathsf{MIP}$. Since $\mathrm{val}^{\mathrm{T}}(G_z) \geq \mathrm{val}(G_z)$, completeness wrt MIP* is unaffected, but soundness may fail. Similarly, for $L \in \mathsf{MIP}^*$, soundness in MIP holds, but completeness may fail.

- It turns out, for the languages reliably verified with classical provers, the entangled value is approximately equal to the classical value, so these languages remain sound wrt quantum provers. The question then is: how big is MIP*? How hard are the problems I can solve by interacting with entangled quantum provers? Before this paper, the best result was $\mathsf{NEEXP} \subseteq \mathsf{MIP}^*$, showing that MIP* is considerably more powerful than MIP: it lets a verifier efficiently check proofs which are *doubly* exponentially long. But maybe it can do other things too.

2.4 Tsirelson's conjecture and decidability

- Let's return to our original problem of understanding correlations. Solving problems in MIP* is the same as approximating the entangled value of some nonlocal game, $\mathrm{val}^{\mathrm{T}}(G)$. In turn, this is related to optimising over the set of tensor product quantum correlations.

- Thinking back to what we know about these correlations, let's think about how we might approximate the value. We made the observation earlier that we can sharply approximate the set of correlations, and hence the entangled value, from below, using finite-dimensional Hilbert spaces. You can just exhaustively search through these space, check all the consistent assignments, and see what you get. That will be a lower bound on the entangled value.

- Let's assume that the answer to Tsirelson's problem, or equivalently Connes' embedding conjecture, is affirmative, and the entangled value is equal to the commuting value. Then we also have a sharp approximation from above, by virtue of the NPA hierarchy.

- Both of these take a long time to run, but the point is that we now have a procedure which is guaranteed to halt and tell us if the entangled value is $\leq 1/3$, or more than $2/3$,

given a "promise" that one of these situations holds. I just dovetail, and run them at the same time.

- This is important. It's telling us that if Tsirelson's conjecture is true, then all the languages MIP^* are decidable: you can decide if a problem instance $z \in L$ or $z \notin L$ using an algorithm which is guaranteed to halt and produce an output.

2.5 Undecidability

- This finally leads us to the paper itself, which I'm only going to briefly comment on. Ji, Natarajan, Vidick, Wright and Yuen claim that MIP^* contains undecidable languages. In fact, they show that

$$\mathsf{MIP}^* = \mathsf{RE},$$

  where the complexity class RE is the set of *recursively enumerable* languages, i.e. those accepted by some Turing machine, however long it takes to run. More precisely, $L \in \mathsf{RE}$ if there is a Turing machine $\mathcal{M}$ and $z \in L$ iff $\mathcal{M}$ halts and accepts when given z as an input.

- This is a huge class, containing not only every decidable language, but also undecidable languages like the Halting Problem. Here, the question is just: will a Turing machine halt on some input? This is in RE because I can just run the Turing machine on the input and see if it halts; if it does, I give a big thumbs up and say "yes, it halts!"

- Since this algorithm is just to run a Turing machine and wait for it to stop, and RE is really the class of all such problems, it suggests correctly that the Halting Problem is *complete* for the class. If you contain the Halting Problem, you contain RE!

- But there is no general procedure $\mathsf{halts}(\mathcal{M}, z)$ for telling when a Turing machine doesn't halt. If there was, you could define an evil recursive function

$$\mathsf{badHalts}() = \begin{cases} \text{if } \mathsf{halts}\,(\mathsf{badHalts}()) & \text{then loop} \\ \text{else} & \text{halt.} \end{cases}$$

  This halts iff it doesn't halt. The only way out of this contradiction is to assume the predicate is *partial*, i.e. doesn't always know the answer.

- Let's pause for a moment and absorb just how powerful multiple interactive provers are. They can reliably persuade a polynomial verifier that a given Turing machine does or does not halt, even though there is no deterministic procedure for doing this. It follows that the answer to Tsirelson's problem is negative, and Connes' conjecture is false.

2.6 $\mathsf{MIP}^* = \mathsf{RE}$

- Let's finish with a few brief details about the proof. It proceeds by designing an entangled interactive proof for the Halting Problem. In more detail, they find an efficient way to

take any Turing machine $\mathcal{M}$ (implicitly including the input), and find a game $G_{\mathcal{M}}$ such that, if the Turing machine halts, then

$$\text{val}^{\text{T}}(G_{\mathcal{M}}) = 1,$$

and if it doesn't, then

$$\text{val}^{\text{T}}(G_{\mathcal{M}}) \leq 1 - \epsilon.$$

I can get the usual soundness parameter by choosing $\epsilon = 2/3$.

- It follows immediately that deciding whether a pattern of correlations $p(ab)$ is in the set of tensor correlations $\mathcal{Q}^{\text{T}}$, or even ϵ-far from it in ℓ_1 distance in the set of probability distributions (promised one or the other is the case), is equivalent to the Halting Problem and therefore undecidable. There are some inputs for which you can literally never answer this question! The set of tensor correlations is dramatically more complicated than anyone thought.

- Since the Halting Problem is complete for RE, it follows that RE $\subseteq$ MIP*, and the reverse inclusion is trivial. Thus,

$$\text{MIP}^* = \text{RE}.$$

We thought that MIP was powerful, but consulting entangled provers make you a god.

2.7 MIPco

- You might wonder how field theory fits into the picture. Local algebras in AQFT are Type III algebras, which are like the commuting Type II factors but without a trace. So they are not described by MIP*. I suspect the corresponding complexity class is MIP$^{\text{co}}$, the set of multiprover interactive proofs with commuting provers.

- That raises the question: what the heck is MIP$^{\text{co}}$? It's known that MIP$^{\text{co}}$ is contained in the complemenet of recursively enumerable languges coRE, i.e. the ones where a Turing machine halts for $z \notin L$ rather than $z \in L$. This is different from RE, because if they were the same, then the Halting Problem would be decidable. I could recursively enumerate both the things which halt and the things which don't halt. Because they're not the same, MIP$^{\text{co}} \neq$ MIP*.

- That's what we would expect, given that there seems to be this huge difference between commuting and tensor models. But what is MIP$^{\text{co}}$? No one knows, but as the paper points out, a pleasing "dual" result would be

$$\text{MIP}^{\text{co}} \stackrel{?}{=} \text{coRe}.$$

They even give a one-paragraph outline of how that proof might go. So I'm not sure if it's a conjecture or another 165-page exercise for the diligent reader.

2.8 Appendix: proof seen from very far away

- The main idea seems to taking a (uniformly generated) family of nonlocal games in some specific *normal form*, say $\{G_n\}$, and *compressing* it exponentially to form a new family of games in normal form $\{G'_n\}$ so that soundness and completeness are preserved:

$$\begin{aligned} \mathrm{val}^{\mathrm{T}}(G'_n) = 1 \quad &\Longrightarrow \quad \mathrm{val}^{\mathrm{T}}(G'_{2^n}) = 1 \\ \mathrm{val}^{\mathrm{T}}(G_n) \leq \frac{1}{2} \quad &\Longrightarrow \quad \mathrm{val}^{\mathrm{T}}(G'_{2^n}) \leq \frac{1}{2}. \end{aligned}$$

- If $\mathcal{E}(G, q)$ is the minimum local dimension of the entangled state needed for players to succeed with probability q, then the compression procedure also has the property

$$\mathcal{E}\left(G'_n, \frac{1}{2}\right) \geq \max\left\{\mathcal{E}\left(G_{2^n}, \frac{1}{2}\right), 2^{2^{\Omega(n)}}\right\}.$$

  So, roughly, compression makes the dimension of the entangled state bigger.

- The second step is to iterately compress to find an entangled interactive proof of the Halting Problem. To begin with, consider a Turing machine $\mathcal{M}$, and a family of nonlocal games $\{G^{(0)}_{\mathcal{M},n}\}$ with the property that if $\mathcal{M}$ halts in at most n steps, then $\mathrm{val}^{\mathrm{T}}(G^{(0)}_{\mathcal{M},n}) = 1$, and otherwise, $\mathrm{val}^{\mathrm{T}}(G^{(0)}_{\mathcal{M},n}) \leq 1/2$.

- Apparently this family is easy to construct in normal form. So you can compress it to get a new family $\{G^{(1)}_{\mathcal{M},n}\}$, such that if $\mathcal{M}$ halts in 2^n steps, then $\mathrm{val}^{\mathrm{T}}(G^{(1)}_{\mathcal{M},n}) = 1$, and if it doesn't halt, the value is $\leq 1/2$. Any strategy which achieves success with probability $q \geq 1/2$ requires an entangled state of dimension at least $2^{2^{\Omega(n)}}$.

- Morally speaking, you iterate the compression procedure until the first game in the family solves the Halting Problem. Technically, you don't do this an infinite number of times, but instead, look for a fixed point of the compression procedure, $\{G^{(\infty)}_{\mathcal{M},n}\}$. The nonlocal game encoding the Halting Problem for $\mathcal{M}$ is then

$$G_{\mathcal{M}} = G^{(\infty)}_{\mathcal{M},1}.$$

- The details are really hard and I don't understand. It's 165 pages for a reason! But this is the proof seen from a great distance.

3 References

- $\mathsf{MIP}^* = \mathsf{RE}$ (2020), JNVWY.
- From Operator Algebras to Complexity Theory and Back (2019), Thomas Vidick.
- A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations (2008), NPA.
- $\mathsf{MIP}^* = \mathsf{RE}$ (2020), Scott Aaronson.