# Research Statement – Adam Bene Watts

## 1 Overview

The laws of quantum mechanics describe a universe in which objects can exist, evolve, and interact in ways entirely unlike the classical world around us. The mathematics describing this behavior is clear, yet its consequences continue to surprise us, despite over a century of research and experiment. My research focuses on understanding quantum mechanics operationally: on asking what is and is not possible in a quantum mechanical world. Answering these questions can lead to important practical applications, as we push the limits of our ability to engineer and control quantum systems to obtain real-world benefits. At the same time, these answers can also lead to important advances in theory, as they sharpen our intuition for the deeply counter-intuitive world of quantum mechanics and the mathematics that surround it.

Much of my research involves considering restricted physical scenarios and then asking what behavior quantum mechanics allows in them. Here the focus is on *information*, rather than specific physical phenomena. In what ways can the outcomes of measurements made on physically distinct, but entangled, quantum systems be correlated? How can classical data be processed by simple quantum circuits? And what can be learned from a few copies of a quantum state? Restricting quantum mechanics in this way often reveals sophisticated mathematical structure which can, in turn, be used to reveal new truths about quantum mechanics.

Another recurring theme in my research involves investigations of simplified models or classical analogues of some aspect of quantum mechanics. How much of quantum behavior can be described by these models? And what behaviors do these models fail to capture? A key example here comes from low-degree non-commutative sum of squares (ncSoS) approximations, which store partial information about quantum states via their low order moments.[1] Other examples come from my work concerning quantum circuits and learning from quantum states, which can be compared to classical circuits and learning from samples drawn from classical probability distributions, respectively.

By definition, the realm of quantum mechanics is one in which there is a gap between what is actually possible and what our classical intuition suggests should be true. My research aims to bridge that gap wherever possible.

## 2 Past Work

In the following I discuss some details of my past research in a way that I hope will be accessible to non-experts in the field. This discussion is organized broadly by themes, some of which encompass many papers and some only a few. For reasons of brevity some papers are not discussed at all: notable omissions include projects concerning probabilistic graph colouring [9] and extremal combinatorics [36], which I worked on before beginning my PhD. A complete list of my publications can be found in my CV.

### 2.1 Nonlocal Games and Quantum Correlations

Consider a collection of non-communicating observers who are sent messages, then subsequently make measurements on a shared quantum state. It has been known since the work of Bell that the measurement outcomes obtained by these observers can be correlated with the messages they are sent in ways that are impossible to reproduce classically. But exactly what correlations are possible? Recent series of breakthroughs [27, 11, 21, 17] have shown that fully describing this set of correlations in general is an incredibly (or technically, impossibly) complex task, and that the exact set of correlations that can be produced depends critically on the underlying structure of the Hilbert space in which the quantum state lives. In particular, [17] showed that measurements made on a state living in an infinite dimensional Hilbert space could obtain correlations not contained in the closure of the set of correlations achievable by measurements made on finite dimensional Hilbert spaces. Through a chain of previous established results [22] this result also disproved

[1] These approximations are discussed in more detail in Section 2.3.

Connes' longstanding embedding conjecture, giving a striking example of how advances in our understanding of quantum mechanics can also lead to advances in the surrounding mathematics.

An important tool involved in the study of these correlations are nonlocal games: situations in which the observers (in this context often called *players*) chose a shared state and measurement with the goal of maximizing a linear functional defined on the space of correlations. Bell's original test can be phrased as a two-player nonlocal game called an XOR game. In these games all measurements are restricted to binary outcomes and the linear functional depends only on the parity of the measurement outcomes. In a now-classic 1987 result [28], Tsirelson showed that optimal measurement strategies for all two-player XOR games could be computed in polynomial time, and that these only required measurements made on finite dimensional Hilbert spaces. Despite this, the complexity of determining optimal strategies for XOR games with 3 or more players remains open, as does the closely related question of what resources (i.e. what dimension of Hilbert space) are required to play these games optimally.[2]

In [32] coauthors and I gave an algebraic characterization of when an XOR game with any number of players could be played *perfectly*, meaning that there existed a measurement strategy for the game achieving the maximum possible value on the associated linear functional. Via a theorem of alternatives, we then showed that perfect strategies for these games only required that the players share a simple finite dimensional state called a GHZ state. These results also showed that it was possible to decide in polynomial time if a given game had a perfect strategy, provided that game came from a subclass of XOR games called symmetric XOR games. In the same paper we also used first moment method techniques to analyze randomly generated XOR games and found a threshold above which these games were unlikely to be perfect, along with exponential lower bounds on the time complexity required to discover this fact using a semi-definite programming hierarchy known as the non-commutative sum of squares hierarchy.[3]

In [6] Bill Helton and I gave a careful analysis of 3 player XOR games using the criterion developed in [32] and showed that it was possible to determine if any 3 player XOR game had a perfect strategy in polynomial time. This proof was centered around an involved algebraic argument – essentially showing that a specific family of instances of the subgroup membership problem on a right-angled Coxeter group were efficiently solvable. Notably, this argument was developed "from scratch" for this problem and, to the best of my knowledge, does not match any other standard techniques for showing solvabilitity of the subgroup membership problem. In a follow-up work [34] Bill and I, along with a USCD graduate student, used the previous results to investigate 3XOR games numerically and gave evidence for a sharp phase transition between an almost-always-perfect and almost-always-imperfect phase for these games.

XOR games are not the only nonlocal games with an algebraic characterization of when they are perfect. A different algebraic criterion characterizes the existence of perfect strategies for closely related family of games called Binary Constrain System (BCS) games [10], and yet another criterion applies to a rich class of nonlocal games called synchronous games [15]. These criterion appear distinct, as do the methods used to prove them. In [33] coauthors and I showed all these criteria could be viewed as consequences of a particular non-commutative Nullstellensatz. This result adds another link between non-commutative algebraic geometry and nonlocal games and further clarifies the algebraic machinery underlying previous work on nonlocal games.

## 2.2 Shallow Quantum Circuits

Near-term quantum computers are likely to be small, noisy machines incapable of maintaining coherence for long periods of time. Yet even these simple machines are expected to have super-classical abilities [4]. The study of short depth circuits asks what computations can be performed by circuits with depth (i.e. number of layers of computation) much shorter than their input length. This question has practical significance, as these short depth circuits are exactly the circuits suitable for implementation on near term machines. And it offers theoretical advances: as will be discussed in a moment, studying short depth circuits allows for unconditional separations between the power of quantum circuits and their classical counterparts – separations which are far beyond our theoretical capabilities otherwise.

In a recent breakthrough result [7], Bravyi Gosset and Koenig compared the ability of constant depth quantum and classical circuits when solving search problems: that is problems for which the goal is to map

[2] Though some bounds are known on the dimension needed to achieve a certain advantage over classical strategies, see [25, 8].
[3] This hierarchy will be discussed in more detail in Section 2.3.

an input bitstring to a "valid" classical output. For this class of problems they showed an unconditional separation between the ability of constant depth quantum circuits and constant depth classical circuits without fanout gates. Equivalently, in the language of complexity theory, they showed an unconditional separation between the classes $\mathsf{QNC}^0$ and $\mathsf{NC}^0$ with respect to search problems. In [35] coauthors and I strengthened this result to give an unconditional separation between constant depth quantum circuits and classical circuits *with* fanout gates (but without parity gates). In the language of complexity theory, this showed a separation between the classes $\mathsf{QNC}^0$ and $\mathsf{AC}^0$. The proof of this result required bringing together a variety of ideas, including previous work concerning the power of shallow quantum circuits with fanout gates [16], ideas from measurement based computation [26], and the proof of a novel switching lemma.

In [37] I worked with Natalie Parham, a masters student at the University of Waterloo (now a graduate student at Columbia), to prove an unconditional separation between quantum and classical circuits for a different type of computational task. We described a distribution which could be sampled from approximately by quantum circuits, but which could not be approximately sampled from by any constant depth, fanout free, classical circuit given access to (a bounded number of) random bits.[4] In the language of complexity theory, we showed an unconditional sampling separation between the complexity classes $\mathsf{QNC}^0$ and $\mathsf{NC}^0$. While this separation only holds in the fanout-free setting, sampling separations are in general harder to show than search separations. In fact, this result gave the first unconditional quantum-classical sampling separation known for shallow circuits. Proving this result required adapting some techniques used to prove hardness of sampling in the classical setting [29] and carefully engineering a quantum circuit using a variety of old [16] and new tricks.

## 2.3 The Non-Commutative Sum of Squares Hierarchy and Hamiltonian Complexity

An important tool in much of my research has been a hierarchy of semidefinite programs known as the non-commutative sum of squares (or ncSoS) hierarchy. When applied to quantum mechanics, the key idea of this hierarchy is to observe that quantum states can be viewed as positive linear maps acting on polynomials formed from some set of observables, and then approximate these states by maps that are restricted to be positive only on low degree polynomials in the observables.[5] This gives an outer approximation the the set of quantum states, which the Helton-McCollough Positivstellensatz [14] then says converges as the degree of the polynomials considered goes to infinity. In the nonlocal game setting, these observables represent the measurements that can be made by players and the ncSoS hierarchy (also called the NPA hierarchy when applied to nonlocal games) gives an upper bound on the correlations achievable with (commuting) measurements made on an infinite dimensional state.

In the local Hamiltonian problem a large matrix (the Hamiltonian) is described via a sum of local observables, and one is asked to compute its minimum eigenvalue. Usually these local observables are described as tensor products of Pauli matrices. The ncSoS algorithm described in the previous paragraph can then be applied in this setting with the Pauli matrices taking the role of the set of observables. The resulting hierarchy of semidefinite programs, called the Quantum Lasserre hierarchy, has proven to be an important tool in our study of the local Hamiltonian problem. Of particular interest is the application of this hierarchy to a family of local Hamiltonians called Quantum Max Cut Hamiltonians [12, 23]. As suggested by the name, computing the minimum eigenvalue of these Hamiltionans can be viewed as a quantum analog of the well-studied (classical) max cut problem, and the use of the quantum Lasserre hierarchy to approximate this value can be viewed as a quantum analogue of the (classical) Lasserre hierarchy used in the famous Goemans-Williamson approximation algorithm.

In [31] coauthors and I gave an alternate semidefinite programming hierarchy for approximating the minimum eigenvalue of Quantum Max Cut Hamiltonians based on two qubit swap operators instead of Pauli matrices. This hierarchy was built on the ncSoS framework, but working out the details of its implementation required understanding the representation theory of the symmetric group, in particular Schur-Weyl duality. This new hierarchy showed advantages over the "standard" Quantum Lasserre Hierarchy in numerical experiments and also allowed for new theoretical possibilities – most notably the idea of solving or approximating the Quantum Max Cut problem inside certain irreps of the symmetric group. In the same work, we also

[4] Here "approximately" means up to some constant error in total variational distance.

[5] Here by positive I mean that the map maps positive elements (i.e. sums of squares) to positive numbers.

built on the representation theory of the symmetric group to give a new proof and generalization of the Lieb-Mattis technique [20] for exactly solving certain Quantum Max Cut Hamiltonians.

### 2.4 Learning about Quantum States

One way to understand quantum mechanics is as a generalization of classical probability theory. From this viewpoint, quantum states take on the role of probability distributions and quantum systems in a given state take the role of samples drawn from that distribution. As in the classical theory, these quantum systems can be measured to learn information about the underlying state. However, unlike in the classical case, these measurements may disturb the quantum system on which they act.

How much can these effects matter? A key result in quantum mechanics known as the gentle measurement lemma gives a partial answer: stating that the disturbance any quantum measurement causes to a system is upper bounded in terms of the amount of information that measurement could provide.[6] Thus, measurements can only cause significant damage to a quantum system when their outcomes tell you something novel about the underlying state, a principle which has been referred to as the *information-disturbance tradeoff*. But this tradeoff only applies to a single measurement. When multiple measurements are made in sequence on a state no such bound applies and some sequences of measurements are known to exhibit an "anti-Xeno effect" in which they can cause large damage to a quantum system while revealing no information about the underlying state.

In [30] I, along with then an early graduate student at the University of Waterloo (and now a graduate student at Columbia) named John Bostanci, showed a version of the gentle measurement lemma could be recovered in the multi-measurement setting provided the order of measurements was randomized. We then introduced and analyzed a variety of new procedures based on this "random gentle measurement lemma" for learning partial information about a quantum state given access to a few quantum systems in that state. These results proved the correctness of a "quantum OR" procedure originally proposed by Scott Aaronson [1], answering an open question from the paper that found an error in Aaronson's original analysis [13]. Additionally, they resulted in a new procedure for the problem of shadow tomography, in which the goal is to learn the average accepting probability of a list of two outcome measurements given sample access to some unknown state [2, 3]. Our new procedure required (asymptotically) the same number of samples as the previous best known protocol [5], but required less technical machinery. Notably, the techniques used in all these proofs were elementary with the key arguments following from a careful combination of the Cauchy-Schwarz Lemma and a monotonicity argument.

## 3 Goals for Future Research

I will end this statement by discussing a few questions which I find fascinating. These aren't all questions I expect to be able to completely answer in a single paper – or even ever – but they are questions that I expect will at least partially guide my research going forwards.

- **When is there a super-polynomial difference between the sample complexities required for classical and quantum learning?** It is possible to extend the parallel between quantum states and classical probability distributions first discussed in Section 2.4. In this extended analogy, as before, quantum states take the role of probability distributions, and systems take the role of samples. Additionally, two outcome measurements take the role of classical true/false events that can be evaluated on samples from the distribution. We can then compare the number of samples of the classical probability distribution required to learn some feature of the events to the number of samples of the quantum state required to learn some feature of the measurements.

  Often, these two sample complexities are polynomially related. This is true when the goal is just to determine if there is a single likely event (Quantum OR) or when the goal is to fully characterize the quantum state (Quantum State Tomography). But for the well-studied shadow tomography problem (discussed in Section 2.4) the best known quantum algorithm requires a number of samples that

[6] Formally, the trace distance between the state of the initial and post measurement system is bounded by the square root of the probability of seeing an unlikely outcome.

scales with the dimension of the Hilbert space the quantum state lives in, while the classical learning algorithms do not. Why this discrepancy? Is this just a failure of our ability to come up with good quantum algorithms? If so, what are we missing? If not, what is the difference between shadow tomography and the previously discussed problems? I find these questions conceptually fascinating.

- **How (else) can we round from low-degree SoS approximations to quantum states?** A key step in the celebrated Goemans-Williamson algorithm is a "rounding" step, which takes a low degree sum of squares representation of a probability distribution to an actual probability distribution on boolean variables. Analogues of this procedure have been studied in the case of the quantum Lasserre hierarchy and Quantum Max Cut. When the goal is to round to a product (i.e. unentangled) state an optimal rounding procedure is known [24] and it bears some resemblance to the original Goemans-Williamson procedure. However, when the goal is to round to an entangled state the situation is much more opaque. Current state of the art techniques are based on tuning parameters in short depth circuits [19, 18]. While the analysis of these algorithms is impressive, rounding algorithms of this form involve some seemingly "ad hoc" choice of parameterized circuit to optimize. Can we improve on these? One fascinating challenge to doing so comes from the fact that general quantum states are not efficiently describable. Thus, optimal rounding algorithms may be necessarily non-constructive (instead just guaranteeing the existence of some hard-to-describe quantum state) in the quantum setting.

- **How can measurement effects in shallow quantum circuits be used for practical benefits?** A key ingredient in the shallow circuit separations discussed in Section 2.2 is a phenomenon where measurements on spatially separated qubits, initially independent, become correlated after conditioning on the outcome of measurements on qubits connecting them. This can also be understood as a special case of quantum teleportation which, in some cases, allows shallow quantum circuits to approximate the behavior of much deeper circuits. We have a few examples of how this result can be used to prove complexity-theoretic separations. But what about more practical ones? To give one particular example, I'm very interested in how intermediate measurements and feed-forward circuits may be used to reduce gate count in quantum compiling algorithms.

- **What lower bounds can we prove on the depth required to compile unitaries?** In the study of classical circuits, no explicit lower bounds are known beyond log-depth. Put differently, it remains possible that log depth classical circuits may be able to perform all known polynomial time algorithms, although this is commonly believed to be very unlikely. While a straightforward counting argument shows many functions cannot be implemented on log depth circuits, no explicit example is known. Can we hope to prove anything stronger in the quantum case? At first the situation may seem even more hopeless, but there are some unique quantum phenomena, such as teleportation and no-cloning, which may make proving lower bounds easier. Can we exploit these? While this is a highly speculative research direction, it is one I hope to look into in the future.

# References


[1] S. Aaronson. Qma/qpoly/spl sube/pspace/poly: de-merlinizing quantum protocols. In *21st Annual IEEE Conference on Computational Complexity (CCC'06)*, pages 13–pp. IEEE, 2006.

[2] S. Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pages 325–338, 2018.

[3] S. Aaronson and G. N. Rothblum. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 322–333, 2019.

[4] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[5] C. Bădescu and R. O'Donnell. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1398–1411, 2021.

[6] A. Bene Watts and J. W. Helton. 3xor games with perfect commuting operator strategies have perfect tensor product strategies and are decidable in polynomial time. *Communications in Mathematical Physics*, pages 1–61, 2023.

[7] S. Bravyi, D. Gosset, and R. König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018.

[8] J. Briët and T. Vidick. Explicit lower and upper bounds on the entangled value of multiplayer xor games. *Communications in Mathematical Physics*, 321(1):181–207, 2013.

[9] X. S. Cai, G. Perarnau, B. Reed, and A. B. Watts. Acyclic edge colourings of graphs with large girth. *Random Structures & Algorithms*, 50(4):511–533, 2017.

[10] R. Cleve, L. Liu, and W. Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1), 2017.

[11] A. Coladangelo and J. Stark. Unconditional separation of finite and infinite-dimensional quantum correlations. *arXiv preprint arXiv:1804.05116*, 2018.

[12] S. Gharibian and O. Parekh. Almost optimal classical approximation algorithms for a quantum generalization of max-cut. *arXiv preprint arXiv:1909.08846*, 2019.

[13] A. W. Harrow, C. Y.-Y. Lin, and A. Montanaro. Sequential measurements, disturbance and property testing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1598–1611. SIAM, 2017.

[14] J. Helton and S. McCullough. A positivstellensatz for non-commutative polynomials. *Transactions of the American Mathematical Society*, 356(9):3721–3737, 2004.

[15] W. Helton, K. P. Meyer, V. I. Paulsen, and M. Satriano. Algebras, synchronous games and chromatic numbers of graphs. *arXiv preprint arXiv:1703.00960*, 2017.

[16] P. Høyer and R. Špalek. Quantum fan-out is powerful. *Theory of computing*, 1(1):81–103, 2005.

[17] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen. Mip*= re. *arXiv preprint arXiv:2001.04383*, 2020.

[18] R. King. An improved approximation algorithm for quantum max-cut on triangle-free graphs. *Quantum*, 7:1180, 2023.

[19] E. Lee. Optimizing quantum circuit parameters via sdp. *arXiv preprint arXiv:2209.00789*, 2022.

[20] E. Lieb and D. Mattis. Ordering energy levels of interacting spin systems. *Journal of Mathematical Physics*, 3(4):749–751, 1962.

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

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

[23] O. Parekh and K. Thompson. Application of the level-2 quantum lasserre hierarchy in quantum approximation algorithms. *arXiv preprint arXiv:2105.05698*, 2021.

[24] O. Parekh and K. Thompson. An optimal product-state approximation for 2-local quantum hamiltonians with positive terms. *arXiv preprint arXiv:2206.08342*, 2022.

[25] D. Pérez-García, M. M. Wolf, C. Palazuelos, I. Villanueva, and M. Junge. Unbounded violation of tripartite bell inequalities. *Communications in Mathematical Physics*, 279(2):455–486, 2008.

[26] R. Raussendorf. Measurement-based quantum computation with cluster states. *International Journal of Quantum Information*, 7(06):1053–1203, 2009.

[27] W. Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 33(1):1–56, 2020.

[28] B. S. Tsirel'son. Quantum analogues of the Bell inequalities. The case of two spatially separated domains. *Journal of Mathematical Sciences*, 36(4):557–570, 1987.

[29] E. Viola. The complexity of distributions. *SIAM Journal on Computing*, 41(1):191–218, 2012.

[30] A. B. Watts and J. Bostanci. Quantum event learning and gentle random measurements. *arXiv preprint arXiv:2210.09155*, 2022.

[31] A. B. Watts, A. Chowdhury, A. Epperly, J. W. Helton, and I. Klep. Relaxations and exact solutions to quantum max cut via the algebraic structure of swap operators. *arXiv preprint arXiv:2307.15661*, 2023.

[32] A. B. Watts, A. W. Harrow, G. Kanwar, and A. Natarajan. Algorithms, bounds, and strategies for entangled xor games. *arXiv preprint arXiv:1801.00821*, 2018.

[33] A. B. Watts, J. W. Helton, and I. Klep. Noncommutative nullstellens\" atze and perfect games. *arXiv preprint arXiv:2111.14928*, 2021.

[34] A. B. Watts, J. W. Helton, and Z. Zhao. Satisfiability phase transtion for random quantum 3xor games. *arXiv preprint arXiv:2209.04655*, 2022.

[35] A. B. Watts, R. Kothari, L. Schaeffer, and A. Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 515–526, 2019.

[36] A. B. Watts, S. Norin, and L. Yepremyan. A turán theorem for extensions via an erdős-ko-rado theorem for lagrangians. *Combinatorica*, 39(5):1149–1171, 2019.

[37] A. B. Watts and N. Parham. Unconditional quantum advantage for sampling with shallow circuits. *arXiv preprint arXiv:2301.00995*, 2023.