Quantum Pseudorandomness and Classical Complexity

William Kretschmer
University of Texas at Austin, TX, USA

Abstract

We construct a quantum oracle relative to which $\mathsf{BQP} = \mathsf{QMA}$ but cryptographic pseudorandom quantum states and pseudorandom unitary transformations exist, a counterintuitive result in light of the fact that pseudorandom states can be "broken" by quantum Merlin-Arthur adversaries. We explain how this nuance arises as the result of a distinction between algorithms that operate on quantum and classical inputs. On the other hand, we show that *some* computational complexity assumption is needed to construct pseudorandom states, by proving that pseudorandom states do not exist if $\mathsf{BQP} = \mathsf{PP}$. We discuss implications of these results for cryptography, complexity theory, and quantum tomography.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases pseudorandom quantum states, quantum Merlin-Arthur

Digital Object Identifier 10.4230/LIPIcs.TQC.2021.2

Related Version *Previous Version*: https://arxiv.org/abs/2103.09320

Funding *William Kretschmer*: Supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship from the US Department of Defense.

Acknowledgements Thanks to Scott Aaronson for suggestions on the writing, Adam Bouland for insightful discussions, and Qipeng Liu for clarifying some questions about [16].

1 Introduction

Pseudorandomness is a key concept in complexity theory and cryptography, capturing the notion of objects that appear random to computationally-bounded adversaries. Recent works have extended the theory of computational pseudorandomness to quantum objects, with a particular focus on quantum states and unitary transformations that resemble the Haar measure [19, 13, 12].

Ji, Liu, and Song [19] define a *pseudorandom state* (PRS) ensemble as a keyed family of quantum states $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$ such that states from the ensemble can be generated in polynomial time, and such that no polynomial-time quantum adversary can distinguish polynomially many copies of a random $|\varphi_k\rangle$ from polynomially many copies of a Haar-random state. They also define an ensemble of *pseudorandom unitary transformations* (PRUs) analogously as a set of efficiently implementable unitary transformations that are computationally indistinguishable from the Haar measure. These definitions can be viewed as quantum analogues of pseudorandom generators (PRGs) and pseudorandom functions (PRFs), respectively. The authors then present a construction of PRSs assuming the existence of quantum-secure one-way functions, and also give a candidate construction of PRUs that they conjecture is secure.

Several applications of PRSs and PRUs are known. PRSs and PRUs are potentially useful in quantum algorithms: in computational applications that require approximations to the Haar measure, PRSs and PRUs can be much more efficient than t-designs, which are information-theoretic approximations to the Haar measure that are analogous to t-

wise independent functions.[1] Cryptographic applications are possible, with [19] giving a construction of a private-key quantum money scheme based on PRSs. Recent work by Bouland, Fefferman, and Vazirani [12] has also established a fundamental connection between PRSs and any possible resolution to the so-called "wormhole growth paradox" in the AdS/CFT correspondence.

1.1 Main Results

Given the importance of PRSs and PRUs across quantum complexity theory, in this work we seek to better understand the theoretical basis for the existence of these primitives. We start with a very basic question: what hardness assumptions are necessary for the existence of PRSs,[2] and which unlikely complexity collapses (such as $\mathsf{P} = \mathsf{PSPACE}$ or $\mathsf{BQP} = \mathsf{QMA}$) would invalidate the security of PRSs? Viewed another way, we ask: what computational power suffices to distinguish PRSs from Haar-random states?

At first glance, it appears that an "obvious" upper bound on the power needed to break PRSs is QMA, the quantum analogue of NP consisting of problems decidable by a polynomial-time quantum Merlin-Arthur protocol (or even QCMA, where the witness is restricted to be classical). If Arthur holds many copies of a pure quantum state $|\psi\rangle$ that can be prepared by some polynomial-size quantum circuit C, then Merlin can send Arthur a classical description of C, and Arthur can verify via the swap test that the output of C approximates $|\psi\rangle$. By contrast, most Haar-random states cannot even be approximated by small quantum circuits. So, in some sense, PRSs can be "distinguished" from Haar-random by quantum Merlin-Arthur adversaries.

There is a subtle problem here, though: QMA is defined as a set of decision problems where the inputs are *classical* bit strings, whereas an adversary against a PRS ensemble inherently operates on a *quantum* input. As a result, it is unclear whether the hardness of breaking PRSs can be related to the hardness of QMA, or any other standard complexity class. Even if we had a proof that $\mathsf{BQP} = \mathsf{QMA}$, this might not give rise to an efficient algorithm for breaking the security of PRSs.

One way to tackle this is to consider quantum adversaries that can query a classical oracle. If we can show that PRSs can be broken by a polynomial-time quantum algorithm with oracle access to some language $\mathcal{L} \subseteq \{0,1\}^*$, we conclude that if PRSs exist, then $\mathcal{L} \notin \mathsf{BQP}$. A priori, it is not immediately obvious whether oracle access to *any* language $\mathcal{L}$ suffices for a polynomial-time quantum adversary to break PRSs. For our first result, we show that a PP-complete language works. Hence, if $\mathsf{BQP} = \mathsf{PP}$, then PRSs do not exist.

▶ **Theorem 1** (Informal version of Theorem 15)**.** *There exists a polynomial-time quantum algorithm augmented with a* PP *oracle that can distinguish PRSs from Haar-random states.*

This raises the natural question of whether the PP oracle in the above theorem can be made weaker. For instance, can we break PRSs with a QCMA or QMA oracle, coinciding with our intuition that the task is solvable by a quantum Merlin-Arthur protocol? In our second result, we show that this intuition is perhaps misguided, as we construct a quantum oracle relative to which such a QMA reduction is impossible.

---

[1] t-designs are also sometimes called "pseudorandom" in the literature, e.g. [27, 14]. We emphasize that t-designs and PRSs/PRUs are fundamentally different notions and that they are generally incomparable: a t-design need not be a PRS/PRU ensemble, or vice-versa.

[2] Note that PRUs imply PRSs, so we focus only on PRSs for this part.

▶ **Theorem 2** (Informal version of Theorem 18 and Theorem 21). *There exists a quantum oracle $\mathcal{O}$ such that:*

(1) $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$, *and*

(2) *PRUs (and hence PRSs) exist relative to $\mathcal{O}$.*

Let us remark how bizarre this theorem appears from a cryptographer's point of view. If $\mathsf{BQP} = \mathsf{QMA}$, then *no* quantum-secure classical cryptographic primitives exist, because such primitives can be broken in NP. So, our construction is a black-box separation between PRUs and *all* quantum-secure classical cryptography – a relativized world in which any computationally-secure cryptography must use quantum communication. Theorem 2 thus provides a negative answer (in the quantum black box setting) to a question of Ji, Liu, and Song [19] that asks if quantum-secure one-way functions are necessary for PRSs. One could even view our result as evidence that it might be possible to base the existence of PRSs and PRUs on weaker assumptions than those usually used for classical cryptography.

1.2 Application: Hyperefficient Shadow Tomography

An immediate corollary of our results is a new impossibility result for shadow tomography. Aaronson [2] defined the shadow tomography problem as the following estimation task: given copies of an n-qubit mixed state ρ and a list of two-outcome measurements $O_1, \ldots, O_M$, estimate $\mathrm{Tr}(O_i\rho)$ for each i up to additive error ε. Aaronson showed that, remarkably, this is possible using very few copies of ρ: just $\mathrm{poly}(n, \log M, \frac{1}{\varepsilon})$ copies suffice, which is polylogarithmic in both the dimension of ρ and the number of quantities to be estimated.

Aaronson then asked in what cases shadow tomography can be made *computationally* efficient with respect to n and $\log M$. Of course, just writing down the input to the problem would take $\Omega(4^n M)$ time if the measurements are given explicitly as Hermitian matrices, and listing the outputs would also take $\Omega(M)$ time. But perhaps one could hope for an algorithm that only operates *implicitly* on both the inputs and outputs. For example, suppose we stipulate the existence of a quantum algorithm that performs the measurement O_i given input $i \in [M]$, and that this algorithm runs in time $\mathrm{poly}(n, \log M)$. Consider a shadow tomography procedure that takes a description of such an algorithm as input, and that outputs a quantum circuit C such that $|C(i) - \mathrm{Tr}(O_i\rho)| \leq \varepsilon$ for each $i \in [M]$.[3] Aaronson calls this a "hyperefficient" shadow tomography protocol if it additionally runs in time $\mathrm{poly}(n, \log M, \frac{1}{\varepsilon})$.

Aaronson gave some evidence that hyperefficient shadow tomography is unlikely to exist, by observing that if hyperefficient shadow tomography is possible, then quantum advice can always be efficiently replaced by classical advice – in other words, $\mathsf{BQP/qpoly} = \mathsf{BQP/poly}$. However, Aaronson and Kuperberg [4] showed a quantum oracle $\mathcal{U}$ relative to which $\mathsf{BQP}^{\mathcal{U}}/\mathsf{qpoly} \neq \mathsf{BQP}^{\mathcal{U}}/\mathsf{poly}$, which implies that hyperefficient shadow tomography is impossible if the observables are merely given as a black box that implements the measurement. The proof of this oracle separation amounts to showing that if the oracle $\mathcal{U}$ either (1) implements a reflection about a Haar-random n-qubit state, or (2) acts as the identity, then no $\mathrm{poly}(n)$-query algorithm can distinguish these two cases, even given a classical witness of size $\mathrm{poly}(n)$.

[3] Note the slight abuse of notation here, as the shadow tomography procedure can err with some small probability, and C itself might be a probabilistic quantum circuit. For simplicity, we assume that the shadow tomography procedure always succeeds and that C is deterministic in this exposition.

One can consider stronger forms of query access to the observables. For instance, in the common scenario where each observable measures fidelity with a pure state, meaning it has the form $O_i = |\psi_i\rangle\langle\psi_i|$, then in addition to the ability to measure overlap with $|\psi_i\rangle$, one might also have the power to produce copies of $|\psi_i\rangle$. Note that the ability to prepare $|\psi_i\rangle$ is generally much more powerful than the ability to recognize $|\psi_i\rangle$, the latter of which is equivalent to oracle access to the reflection $\mathbb{1} - 2|\psi_i\rangle\langle\psi_i|$. For example, Aaronson and Kuperberg's oracle separation of QCMA and QMA [4] amounts to building an oracle relative to which certain quantum states can be recognized efficiently but cannot be approximately prepared by small quantum circuits. Other black-box separations of state preparation and state reflection are known, e.g. [9], so one might hope that this type of query access could be substantially more powerful for shadow tomography as well.

Nevertheless, our results imply that black-box hyperefficient shadow tomography is impossible even in this setting where we have state preparation access to the observables. This follows from the simple observation that hyperefficient shadow tomography of this form would suffice to break PRS ensembles with a QCMA oracle.

▶ **Theorem 3.** *If a hyperefficient shadow tomography procedure exists that works for any list of observables of the form $|\psi_1\rangle\langle\psi_1|, \ldots, |\psi_M\rangle\langle\psi_M|$ given state preparation access to $|\psi_1\rangle, \ldots, |\psi_M\rangle$, then all PRS ensembles can be broken by polynomial-time quantum adversaries with oracle access to* QCMA*.*

Proof sketch. For a given PRS ensemble $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$, we have state preparation access to the observable list $\{|\varphi_k\rangle\langle\varphi_k|\}_{k\in\mathcal{K}}$ by way of the generating algorithm of the PRS. Hence, we can run hyperefficient shadow tomography using this observable list on copies of some unknown state $|\psi\rangle$. Suppose that this produces a quantum circuit C such that $|C(k) - \mathrm{Tr}(|\varphi_k\rangle\langle\varphi_k|\psi\rangle\langle\psi|)| \leq \frac{1}{10}$ for each $k \in \mathcal{K}$. Observe that the problem of deciding whether there exists some k such that $C(k) \geq \frac{9}{10}$ is in QCMA. If $|\psi\rangle$ is pseudorandom, then such a k always exists (whichever k satisfies $|\psi\rangle = |\varphi_k\rangle$), whereas if $|\psi\rangle$ is Haar-random, such a k exists with negligible probability over the choice of $|\psi\rangle$. Hence, these two ensembles can be distinguished by feeding C into this QCMA language. ◀

The above theorem also relativizes, in the sense that if the shadow tomography procedure only accesses the state preparation algorithm via a black box $\mathcal{O}$, then hyperefficient shadow tomography lets us break PRSs in polynomial time with oracle access to $\mathcal{O}$ and $\mathsf{QCMA}^{\mathcal{O}}$. Since Theorem 2 gives an oracle relative to which $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QCMA}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$ and PRSs exist, we conclude that hyperefficient shadow tomography is impossible with only black-box state preparation access to the observables.

1.3 Our Techniques

The starting point for the proof of Theorem 1, which gives an upper bound of PP on the power needed to break pseudorandom states, is a theorem of Huang, Kueng, and Preskill [17] that gives a simple procedure for shadow tomography.

▶ **Theorem 4** ([17])**.** *Fix M different observables $O_1, O_2, \ldots, O_M$ and an unknown n-qubit mixed state ρ. Then there exists a quantum algorithm that performs $T = O(\log(M/\delta)/\varepsilon^2 \cdot \max_i \mathrm{Tr}(O_i^2))$ single-copy measurements in random Clifford bases of ρ, and uses the measurement results to estimate the quantities $\mathrm{Tr}(O_1\rho), \mathrm{Tr}(O_2\rho), \ldots, \mathrm{Tr}(O_M\rho)$, such that with probability at least $1 - \delta$, all of the M quantities are correct up to additive error ε.*

If $\{|\varphi_k\rangle_{k\in\mathcal{K}}\}$ is a family of PRSs, then by choosing $O_k = |\varphi_k\rangle\langle\varphi_k|$ for each $k \in \mathcal{K}$ to be the list of observables, we can use the above algorithm to determine whether ρ is close to one of the states in the PRS ensemble. A Haar-random state will be far from *all* of the pseudorandom states with overwhelming probability. Hence, Theorem 4 implies the existence of an algorithm that distinguishes the pseudorandom and Haar-random ensembles, by performing a polynomial number of random Clifford measurements and analyzing the results. The key observation is that the Clifford measurements can be performed efficiently, even though the resulting analysis (which operates on purely classical information) might be computationally expensive.

Next, one could try to argue that the computationally difficult steps in the above algorithm can be made efficient with a PP oracle. However, we take a different approach. We adopt a Bayesian perspective: suppose that with 50% probability we are given copies of a Haar-random state, and otherwise with 50% probability we are given copies of a randomly chosen state from the pseudorandom ensemble. We wish to distinguish these two cases using only the results of the random Clifford measurements as observed data. One way to do this is via the Bayes decision rule: we compute the posterior probability of being Haar-random or pseudorandom given the measurements, and then guess the more likely result. In fact, the Bayes decision rule is well-known to be the *optimal* decision rule in general, in the sense that any decision rule errs at least as often as the Bayes decision rule (see e.g. [11, Chapter 4.4.1]). Hence, because the algorithm of Huang, Kueng, and Preskill (Theorem 4) distinguishes the Haar-random and pseudorandom ensembles with good probability, the Bayes decision rule conditioned on the random Clifford measurements must work *at least* as well at the same distinguishing task.

Finally, we observe that using a quantum algorithm with postselection, we can approximate the relevant posterior probabilities needed for the Bayes decision rule. This allows us to appeal to the equivalence $\mathsf{PostBQP} = \mathsf{PP}$ [1] to simulate this postselection with a PP oracle.

Technically, one challenge is that the postselected quantum algorithm requires the ability to prepare copies of a Haar-random state, even though a polynomial-time quantum algorithm cannot even approximately prepare most Haar-random states. The solution is to replace the Haar ensemble by an approximate quantum design, which we argue does not substantially change the success probability of the algorithm.

For our second result (Theorem 2), the oracle construction we use is simple to describe. The oracle $\mathcal{O}$ consists of two parts: a quantum oracle $\mathcal{U} = \{\mathcal{U}_n\}_{n\in\mathbb{N}}$, where each $\mathcal{U}_n$ consists of 2^n different Haar-random n-qubit unitary matrices, and a classical oracle $\mathcal{P}$ that is an arbitrary PSPACE-complete language. We prove that Theorem 2 holds with probability 1 over the choice of $\mathcal{U}$.

Showing that PRUs exist relative to $(\mathcal{U}, \mathcal{P})$ is reasonably straightforward. The proof uses the BBBV theorem (i.e. the optimality of Grover's algorithm) [10], and is analogous to showing that one-way functions or pseudorandom generators exist relative to a random *classical* oracle, as was shown by Impagliazzo and Rudich [18]. We only rigorously prove security against adversaries with classical advice, though we believe that the recently introduced framework of Chung, Guo, Liu, and Qian [16] should yield a security proof against adversaries with quantum advice.

Slightly more technically involved is proving that $\mathsf{BQP}^{\mathcal{U},\mathcal{P}} = \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$. To do so, we argue that a QMA verifier is not substantially more powerful than a BQP machine at learning nontrivial properties of $\mathcal{U}$. More precisely, we argue that if a QMA verifier $\mathcal{V}$ makes T queries to $\mathcal{U}_n$ for some $n \in \mathbb{N}$, then either (1) $n = O(\log T)$ is sufficiently small that $\mathrm{poly}(T)$ queries to $\mathcal{U}_n$ actually suffice to learn $\mathcal{U}_n$ to inverse-polynomial precision, or else (2) $n = \omega(\log t)$ is

sufficiently large that with high probability, the maximum acceptance probability of $\mathcal{V}$ (over the choice of Merlin's witness) is close to the average maximum acceptance probability of $\mathcal{V}$ when $\mathcal{U}_n$ is replaced by a random set of matrices sampled from the Haar measure. We prove this as a consequence of the extremely strong concentration of measure properties exhibited by the Haar measure [22].

This allows a $\mathsf{BQP}^{\mathcal{U},\mathcal{P}}$ machine to approximate the maximum acceptance probability of $\mathcal{V}^{\mathcal{U},\mathcal{P}}$ as follows. In case (1), the $\mathsf{BQP}^{\mathcal{U},\mathcal{P}}$ machine first queries $\mathcal{U}_n$ enough times to learn a unitary transformation $\widetilde{\mathcal{U}}_n$ that is close to $\mathcal{U}_n$, and then hard codes $\widetilde{\mathcal{U}}_n$ into a new $\mathsf{QMA}^{\mathcal{P}}$ verifier $\widetilde{\mathcal{V}}$ that simulates $\mathcal{V}$ by replacing queries to $\mathcal{U}_n$ with calls to $\widetilde{\mathcal{U}_n}$. In case (2), the $\mathsf{BQP}^{\mathcal{U},\mathcal{P}}$ machine similarly constructs a new $\mathsf{QMA}^{\mathcal{P}}$ verifier $\widetilde{\mathcal{V}}$, instead simulating $\mathcal{V}$ by replacing queries to $\mathcal{U}_n$ with unitaries chosen from an approximate polynomial design.[4] In both cases, $\widetilde{\mathcal{V}}$ defines a $\mathsf{QMA}^{\mathcal{P}}$ problem. Because $\mathcal{P}$ is PSPACE-complete, $\mathsf{BQP}^{\mathcal{P}} = \mathsf{QMA}^{\mathcal{P}} = \mathsf{PSPACE}$, and therefore this problem can be decided with a single query to $\mathcal{P}$.

The astute reader may notice that this proof works for more general choices of $\mathcal{P}$: it shows that for any oracle $\mathcal{P}$, if $\mathsf{BQP}^{\mathcal{P}} = \mathsf{QMA}^{\mathcal{P}}$, then $\mathsf{BQP}^{\mathcal{U},\mathcal{P}} = \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$ with probability 1 over the choice of $\mathcal{U}$. An interesting consequence is the special case when $\mathcal{P}$ is trivial.

▶ **Corollary 5.** *If* $\mathsf{BQP} = \mathsf{QMA}$*, then* $\mathsf{BQP}^{\mathcal{U}} = \mathsf{QMA}^{\mathcal{U}}$ *with probability* 1 *over the choice of* $\mathcal{U}$*.*

In words, if $\mathsf{BQP} = \mathsf{QMA}$ in the unrelativized world, then the complexity classes also coincide relative to a collection $\mathcal{U}$ of Haar-random oracles. Or, viewed another way, separating BQP from QMA relative to $\mathcal{U}$ requires separating them in the unrelativized world. This is in stark contrast to the case of random *classical* oracles, where we can prove unconditionally that for a uniformly random oracle $\mathcal{O}$, $\mathsf{BQP}^{\mathcal{O}} \neq \mathsf{QMA}^{\mathcal{O}}$ (and indeed, $\mathsf{NP}^{\mathcal{O}} \not\subset \mathsf{BQP}^{\mathcal{O}}$) with probability 1 over $\mathcal{O}$ [10].

1.4 Open Problems

Can we prove a similar result to Theorem 2 using a *classical* oracle, for either PRUs or PRSs? Attempting to resolve this question seems to run into many of the same difficulties that arise in constructing a classical oracle separation between QCMA and QMA, which also remains an open problem [4]. For one, as pointed out in [4], we do not even know whether every n-qubit unitary transformation can be approximately implemented in $\mathrm{poly}(n)$ time relative to some classical oracle. Even if one could resolve this, it is not clear whether the resulting PRUs or PRSs would be secure against adversaries with the power of QMA. For instance, we show in Appendix C that an existing construction of PRSs, whose security is provable in the random oracle model [13], can be broken with an NP oracle.

What else can be said about the hardness of learning quantum states and unitary transformations, either in the worst case or on average? A related question is to explore the hardness of problems involving quantum *meta-complexity*: that is, problems that themselves encode computational complexity or difficulty. Consider, for example, a version of the minimum circuit size problem (MCSP) for quantum states: given copies of a pure quantum state $|\psi\rangle$, determine the size of the smallest quantum circuit that approximately outputs $|\psi\rangle$. If PRSs exist, then this task should be hard, but placing an upper bound on the complexity of this task might be difficult in light of our results. We view this problem as particularly intriguing because it does not appear to have an obvious classical analogue, and also because of

[4] Technically, this requires choosing a random element of the polynomial design for each $x \in \{0,1\}^n$ by means of a random oracle, so we use Zhandry's strategy [28] to simulate T quantum queries to a random oracle using a $2T$-wise independent function.

its relevance to the wormhole growth paradox and Susskind's Complexity=Volume conjecture in AdS/CFT [12, 25, 24]. A number of recent breakthroughs in complexity theory have involved ideas from meta-complexity (see surveys by Allender [6, 7]), and it would be interesting to see which of these techniques could be ported to the quantum setting.

What other complexity-theoretic evidence can be given for the existence of PRSs and PRUs? Can we give candidate constructions of PRSs or PRUs that do not rely on the assumption $\mathsf{BQP} \neq \mathsf{QMA}$? To give a specific example, an interesting question is whether polynomial-size quantum circuits with random local gates form PRUs. Random circuits are known to information-theoretically approximate the Haar measure in the sense that they form approximate unitary designs [15], and it seems conceivable that they could also be computationally pseudorandom.

2 Preliminaries

2.1 Notation

Throughout, $[n]$ denotes the set of integers $\{1, 2, \ldots, n\}$, and $[n, m]$ denotes the set of integers $\{n, n+1, n+2, \ldots, m\}$. If $x \in \{0,1\}^n$ is a binary string, then $|x|$ denotes the length of x. For X a finite set, we let $|X|$ denote the size of X. If X is a probability distribution, then we use $x \leftarrow X$ to denote a random variable x sampled according to X. When X is a finite set, we also use $x \leftarrow X$ to indicate a random variable x drawn uniformly from X. A function $f(n)$ is *negligible* if for every constant $c > 0$, $f(n) \leq \frac{1}{n^c}$ for all sufficiently large n. We use $\mathrm{negl}(n)$ to denote an arbitrary negligible function, and $\mathrm{poly}(n)$ to denote an arbitrary polynomially-bounded function.

We use $||M||_F = \sqrt{\mathrm{Tr}\,(M^\dagger M)}$ to denote the Frobenius norm of a matrix M. We denote by $||A||_\diamond$ the diamond norm of a superoperator A acting on density matrices (see [5] for a definition). For a unitary matrix U, we use $U \cdot U^\dagger$ to denote the superoperator that maps a density matrix ρ to $U\rho U^\dagger$.

We use $\mathbb{S}(N)$ to denote the set of N-dimensional pure quantum states, and $\mathbb{U}(N)$ to denote the group of $N \times N$ unitary matrices. When $N = 2^n$, we identify these with n-qubit states and unitary transformations, respectively. We use σ_N to denote the Haar measure on $\mathbb{S}(N)$, and we let μ_N denote the Haar measure over $\mathbb{U}(N)$. We write $\mathbb{U}(N)^M$ for the space of $MN \times MN$ block-diagonal unitary matrices, where each block has size $N \times N$, and we also identify $\mathbb{U}(N)^M$ with M-tuples of $N \times N$ unitary matrices. We use μ_N^M to denote the product measure $\mu_N^M(U_1, U_2, \ldots, U_M) = \mu_N(U_1) \cdot \mu_N(U_2) \cdots \mu_N(U_M)$ on $\mathbb{U}(N)^M$.

We assume familiarity with standard complexity classes such as BQP and PP, including relativized versions of these classes that can query a quantum or classical oracle. For completeness, we define some of the relevant complexity classes and related notions in Appendix B.

We use superscript notation for algorithms that query oracles. For instance, $\mathcal{A}^{\mathcal{U}}(x, |\psi\rangle)$ denotes a quantum algorithm $\mathcal{A}$ that queries an oracle $\mathcal{U}$ and receives a classical input x and a quantum input $|\psi\rangle$.

2.2 Quantum Information

We require the following well-known fact, which bounds the distance in the diamond norm between two unitary superoperators in terms of the Frobenius norm of the difference of the two matrices. We provide a proof in Appendix A.

▶ **Lemma 6.** *Let $U, V \in \mathbb{U}(N)$. Then $||U \cdot U^\dagger - V \cdot V^\dagger||_\diamond \leq 2||U - V||_F$.*

We use the notion of an ε-approximate quantum (state) t-design, which is a distribution over quantum states that information-theoretically approximates the Haar measure over states.

▶ **Definition 7** (Approximate quantum design [8]). *A probability distribution S over $\mathbb{S}(N)$ is an* ε-approximate quantum t-design *if:*

$$(1-\varepsilon)\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma_N}|\psi\rangle\langle\psi|^{\otimes t}\leq\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow S}|\psi\rangle\langle\psi|^{\otimes t}\leq(1+\varepsilon)\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma_N}|\psi\rangle\langle\psi|^{\otimes t}$$

and:

$$\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow S}|\psi\rangle\langle\psi|=\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma_N}|\psi\rangle\langle\psi|.$$

Similarly, we require ε-approximate *unitary* t-designs, which are approximations to the Haar measure over unitary matrices. The definition of ε-approximate unitary t-designs is more technical, so we point to [15, Definition 2] for a formal definition. While there are several definitions of approximate t-designs used in the literature, for this work it is crucial that we use *multiplicative* approximate designs for both states and unitaries, meaning that the designs approximate the first t moments of the Haar measure to within a multiplicative $1\pm\varepsilon$ error (as opposed to additive error).

Efficient constructions of approximate unitary t-designs over qubits are known, as below.

▶ **Lemma 8.** *Fix $\varepsilon>0$. For each $n,t\in\mathbb{N}$, there exists $m(n)\leq\mathrm{poly}(n)$ and a $\mathrm{poly}(n,t)$-time classical algorithm $\mathcal{S}$ that takes as input a random string $x\leftarrow\{0,1\}^m$ and outputs a description of a quantum circuit on n qubits such that the circuits sampled from $\mathcal{S}$ form an ε-approximate unitary t-design over $\mathbb{U}(2^n)$.*

Proof sketch. Fix an arbitrary universal quantum gate set G with algebraic entries that is closed under taking inverses (e.g. $G=\{\mathrm{CNOT},H,T,T^\dagger\}$). Brandão, Harrow, and Horodecki [15, Corollary 7] show that n-qubit quantum circuits consisting of $\mathrm{poly}(n,t)$ random gates sampled from G, applied to random pairs of qubits, form ε-approximate unitary t-designs. So, $\mathcal{S}$ just has to sample from this distribution, which can be done with $\mathrm{poly}(n,t)$ bits of randomness. ◀

Note that this also implies an efficient construction of ε-approximate quantum (state) t-designs, as if S is an ε-approximate unitary t-design over $\mathbb{U}(N)$ then $S|\psi\rangle$ is an ε-approximate quantum t-design for any fixed $|\psi\rangle$ (e.g. $|0^n\rangle$).

Essentially the only property we need of approximate t-designs is that they can be used in place of the Haar measure in any quantum algorithm that uses t copies of a Haar-random state (or t queries to a Haar-random unitary), and the measurement probabilities of the algorithm will change by only a small multiplicative factor.

▶ **Fact 9.** *Let S be an ε-approximate quantum t-design over $\mathbb{S}(N)$, and let $\mathcal{A}$ be an arbitrary quantum measurement. Then:*

$$(1-\varepsilon)\Pr_{|\psi\rangle\leftarrow\sigma_N}\left[\mathcal{A}\left(|\psi\rangle^{\otimes t}\right)=1\right]\leq\Pr_{|\psi\rangle\leftarrow S}\left[\mathcal{A}\left(|\psi\rangle^{\otimes t}\right)=1\right]\leq(1+\varepsilon)\Pr_{|\psi\rangle\leftarrow\sigma_N}\left[\mathcal{A}\left(|\psi\rangle^{\otimes t}\right)=1\right].$$

▶ **Fact 10** ([15]). *Let S be an ε-approximate unitary t-design over $\mathbb{U}(N)$, and let $\mathcal{A}^U$ be an arbitrary quantum algorithm that makes t queries to some $U\in\mathbb{U}(N)$. Then:*

$$(1-\varepsilon)\Pr_{U\leftarrow\mu_N}\left[\mathcal{A}^U=1\right]\leq\Pr_{U\leftarrow S}\left[\mathcal{A}^U=1\right]\leq(1+\varepsilon)\Pr_{U\leftarrow\mu_N}\left[\mathcal{A}=1\right].$$

We require the following concentration inequality on the Haar measure, which is stated in terms of Lipschitz continuous functions. For a metric space M with metric d, a function $f : M \to \mathbb{R}$ is *L-Lipschitz* if for all $x, y \in M$, $|f(x) - f(y)| \leq L \cdot d(x, y)$.

▶ **Theorem 11** ([22, Theorem 5.17]). *Given $N_1, \ldots, N_k \in \mathbb{N}$, let $X = \mathbb{U}(N_1) \oplus \cdots \oplus \mathbb{U}(N_k)$ be the space of block-diagonal unitary matrices with blocks of size $N_1, \ldots, N_k$. Let $\mu = \mu_{N_1} \times \cdots \times \mu_{N_k}$ be the product of Haar measures on X. Suppose that $f : X \to \mathbb{R}$ is L-Lipschitz in the Frobenius norm. Then for every $t > 0$:*

$$\Pr_{U \leftarrow \mu}\left[f(U) \geq \mathop{\mathbb{E}}_{V \leftarrow \mu}[f(V)] + t\right] \leq \exp\left(-\frac{(N-2)t^2}{24L^2}\right),$$

where $N = \min\{N_1, \ldots, N_k\}$.

2.3 Cryptography

A family of functions $\{f_k\}_{k \in \mathcal{K}}$ where $f_k : \{0,1\}^n \to \{0,1\}^m$ is called *t-wise independent* if for every distinct $x_1, x_2, \ldots, x_t \in \{0,1\}^n$ and every (not necessarily distinct) $y_1, y_2, \ldots, y_t \in \{0,1\}^m$:

$$\Pr_{k \leftarrow \mathcal{K}}[f_k(x_i) = y_i \; \forall i \in [t]] = 2^{-mt}.$$

Efficient constructions of t-wise independent functions are known, in the sense that one can sample a random f_k from a t-wise independent function family and make queries to f_k in $\mathrm{poly}(t, n, m)$ time [28]. Our primary use of t-wise independent functions is in simulating random oracles: $2t$-wise independent functions can be used in place of a uniformly random function $\{0,1\}^n \to \{0,1\}^m$ in any quantum algorithm that makes at most t queries to the random function; see Zhandry [28, Theorem 6.1] for more details.

We use the following definitions of pseudorandom quantum states (PRSs) and pseudorandom unitaries (PRUs), which were introduced by Ji, Liu, and Song [19].

▶ **Definition 12** (Pseudorandom quantum states [19]). *Let $\kappa \in \mathbb{N}$ be the security parameter. Let D be the dimension of a quantum system and let $\mathcal{K}$ be the key set, both parameterized by κ. A keyed family of quantum states $\{|\varphi_k\rangle\}_{k \in \mathcal{K}} \subset \mathbb{S}(D)$ is* pseudorandom *if the following two conditions hold:*

(1) *(Efficient generation) There is a polynomial-time quantum algorithm G that generates $|\varphi_k\rangle$ on input k, meaning $G(k) = |\varphi_k\rangle$.*

(2) *(Computationally indistinguishable) For any polynomial-time quantum algorithm $\mathcal{A}$ and $T = \mathrm{poly}(\kappa)$:*

$$\left|\Pr_{k \leftarrow \mathcal{K}}\left[\mathcal{A}\left(|\varphi\rangle^{\otimes T}\right) = 1\right] - \Pr_{|\psi\rangle \leftarrow \sigma_D}\left[\mathcal{A}\left(|\psi\rangle^{\otimes T}\right) = 1\right]\right| = \mathrm{negl}(\kappa).$$

▶ **Definition 13** (Pseudorandom unitary transformations [19]). *Let $\kappa \in \mathbb{N}$ be the security parameter. Let D be the dimension of a quantum system and let $\mathcal{K}$ be the key set, both parameterized by κ. A keyed family of unitary transformations $\{U_k\}_{k \in \mathcal{K}} \subset \mathbb{U}(D)$ is* pseudorandom *if the following two conditions hold:*

(1) *(Efficient computation) There is a polynomial-time quantum algorithm G that implements U_k on input k, meaning that for any $|\psi\rangle \in \mathbb{S}(D)$, $G(k, |\psi\rangle) = U_k |\psi\rangle$.*

(2) *(Computationally indistinguishable) For any polynomial-time quantum algorithm $\mathcal{A}^U$ that queries $U \in \mathbb{U}(D)$:*

$$\left|\Pr_{k \leftarrow \mathcal{K}}\left[\mathcal{A}^{U_k}(1^\kappa) = 1\right] - \Pr_{U \leftarrow \mu_D}\left[\mathcal{A}^U(1^\kappa) = 1\right]\right| = \mathrm{negl}(\kappa).$$

We generally take the key set $\mathcal{K} = \{0,1\}^\kappa$ and choose $D = 2^n$ for some $n = \text{poly}(\kappa)$ in the above definitions. We sometimes call the negligible quantities in the above definitions the *advantage* of the quantum adversary $\mathcal{A}$.

In this work, we consider security against non-uniform quantum algorithms with classical advice, which means that the adversary is allowed to be a different polynomial-time quantum algorithm for each setting of the security parameter $\kappa \in \mathbb{N}$. Without loss of generality, such an adversary can always be assumed to take the form of a *uniform* $\text{poly}(\kappa)$-time quantum algorithm $\mathcal{A}(1^\kappa, x)$, where $x \in \{0,1\}^{\text{poly}(\kappa)}$ is an advice string that depends only on κ.

3 Breaking Pseudorandomness with a Classical Oracle

In this section, we prove that a polynomial-time quantum algorithm with a PP oracle can distinguish a PRS from a Haar-random state. First, we need a lemma about the overlap between a fixed state $|\varphi\rangle$ and a Haar-random state $|\psi\rangle$.

▶ **Lemma 14.** *Let $|\varphi\rangle \in \mathbb{S}(N)$, and let $\varepsilon > 0$. Then:*

$$\Pr_{|\psi\rangle \leftarrow \sigma_N}\left[|\langle\psi|\varphi\rangle|^2 \geq \varepsilon\right] \leq e^{-\varepsilon N}.$$

Proof. This follows from standard concentration inequalities, or even an explicit computation, using the fact that $|\langle\psi|\varphi\rangle|^2$ is roughly exponentially distributed for a random state $|\psi\rangle$. See e.g. [15, Equation (14)] ◀

The formal statement of our result is below.

▶ **Theorem 15.** *For any PRS ensemble $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$ of n-qubit states with security parameter κ, there exists a* PP *language $\mathcal{L}$, a $\text{poly}(\kappa)$-time quantum algorithm $\mathcal{A}^{\mathcal{L}}$, and $T = \text{poly}(\kappa)$ such that the following holds. Let $X \leftarrow \{0,1\}$ be a uniform random bit. Let $|\psi\rangle$ be sampled uniformly from the PRS ensemble if $X = 0$, and otherwise let $|\psi\rangle$ be sampled from the Haar measure σ_{2^n} if $X = 1$. Then we have:*

$$\Pr_{X,|\psi\rangle}\left[\mathcal{A}^{\mathcal{L}}\left(|\psi\rangle^{\otimes T}\right) = X\right] \geq 0.995.$$

Proof. We first describe $\mathcal{A}$. For some T to be chosen later, on input $|\psi\rangle^{\otimes T}$, $\mathcal{A}$ measures each copy of $|\psi\rangle$ in a different randomly chosen Clifford basis. Call the list of measurement bases $b = (b_1, b_2, \ldots, b_T)$ and the measurement results $c = (c_1, c_2, \ldots, c_T)$. $\mathcal{A}$ then feeds (b, c) into a single query to $\mathcal{L}$, and outputs the result of the query. This takes polynomial time because there exists an $O(n^3)$-time algorithm to sample a random n-qubit Clifford unitary, and this algorithm also produces an implementation of the unitary with $O(n^2/\log n)$ gates [20, 3].

The PP language $\mathcal{L}$ we choose for the oracle is most easily described in terms of a PostBQP algorithm $\mathcal{B}(b, c)$ (i.e. a postselected polynomial-time quantum algorithm, as in Definition 23), by the equivalence PostBQP = PP [1].[5] Let S be a $\frac{1}{17}$-approximate n-qubit quantum T-design (Definition 7) such that a state can be drawn from S in $\text{poly}(\kappa)$ time (because $n, T \leq \text{poly}(\kappa)$, the existence of such a design follows from Lemma 8). $\mathcal{B}$ begins by initializing the state:

$$\hat{\rho} = \frac{1}{2}|0\rangle\langle 0| \otimes \mathop{\mathbb{E}}_{k\leftarrow\mathcal{K}}\left[|\varphi_k\rangle\langle\varphi_k|^{\otimes T}\right] + \frac{1}{2}|1\rangle\langle 1| \otimes \mathop{\mathbb{E}}_{|\psi\rangle\leftarrow S}\left[|\psi\rangle\langle\psi|^{\otimes T}\right].$$

[5] Note that any promise problem in PostBQP is also in PP [1], and any promise problem in PP can be extended to a language in PP because PP is a syntactic class. Hence, we might as well take a language in PP instead of a promise problem.

$\mathcal{B}$ measures all but the leftmost qubit of $\hat{\rho}$ in the basis given by b, and postselects on observing c (i.e. $\mathcal{B}$ outputs $*$ if the measurements are not equal to c). Finally, conditioned on postselection succeeding, $\mathcal{B}$ measures and outputs the result of the leftmost qubit that was not measured.

It remains to show that $\mathcal{A}$ distinguishes the pseudorandom and Haar-random state ensembles. For the purpose of this analysis, it will be convenient to view $\hat{\rho}$ as an approximation to the state:

$$\rho = \frac{1}{2}|0\rangle\langle 0| \otimes \mathop{\mathbb{E}}_{k\leftarrow\mathcal{K}}\left[|\varphi_k\rangle\langle\varphi_k|^{\otimes T}\right] + \frac{1}{2}|1\rangle\langle 1| \otimes \mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma_{2^n}}\left[|\psi\rangle\langle\psi|^{\otimes T}\right],$$

where the ε-approximate T-design S is replaced by the Haar measure σ_{2^n}. Indeed, we will essentially argue the algorithm's correctness if the state $\hat{\rho}$ is replaced by ρ, and then argue that this implies the correctness of the actual algorithm.

For each $k \in \mathcal{K}$, define $O_k = |\varphi_k\rangle\langle\varphi_k|$. Note that if $X = 0$ (i.e. $|\psi\rangle$ is pseudorandom), there always exists a k such that $\mathrm{Tr}(O_k|\psi\rangle\langle\psi|) = 1$, namely whichever k satisfies $|\psi\rangle = |\varphi_k\rangle$. On the other hand, by Lemma 14 and a union bound, if $X = 1$ (i.e. $|\psi\rangle$ is Haar-random), $\mathrm{Tr}(O_k|\psi\rangle\langle\psi|) < \frac{1}{3}$ for every $k \in \mathcal{K}$, except with probability at most $|\mathcal{K}| \cdot e^{-2^n/3} \leq \mathrm{negl}(\kappa)$ over $|\psi\rangle$.

If we choose $M = |\mathcal{K}|$, $\varepsilon = \frac{1}{3}$, and $\delta = 0.001 - |\mathcal{K}| \cdot e^{-2^n/3}$, then by Theorem 4 there exists a quantum algorithm that takes as input the results (b, c) of $T = O(\log|\mathcal{K}|) = O(\kappa)$ single-copy random Clifford measurements, uses the measurement results to estimate $\mathrm{Tr}(O_k|\psi\rangle\langle\psi|)$ for each k up to additive error $\frac{1}{3}$, and is correct with probability at least $0.999 + |\mathcal{K}| \cdot e^{-2^n/3}$. In particular, this algorithm can distinguish the pseudorandom ensemble from the Haar-random ensemble, by checking if there exists a k such that the estimate for $\mathrm{Tr}(O_k|\psi\rangle\langle\psi|)$ is at least $\frac{2}{3}$. Call this algorithm $\mathcal{C}$, so that $\Pr[\mathcal{C}(b, c) = X] \geq 0.999$.

We will not actually use $\mathcal{C}$, but only its existence. By the optimality of the Bayes decision rule [11, Chapter 4.4.1], because $\mathcal{C}$ uses (b, c) to identify a state $|\psi\rangle$ as either Haar-random or pseudorandom with probability 0.999, an algorithm that computes the maximum a posteriori estimate of X also succeeds with probability at least 0.999. In symbols, let $p_i = \Pr[X = i \mid b, c]$, which we view as a random variable (depending on b and c) for each $i \in \{0, 1\}$. Then $\Pr\left[\arg\max_i p_i = X\right] \geq 0.999$.

Next, observe that $\Pr\left[\arg\max_i p_i = X\right] = \mathbb{E}\left[\Pr\left[\arg\max_i p_i = X | b, c\right]\right] = \mathbb{E}\left[\max_i p_i\right]$, by the law of total expectation. Hence, by Markov's inequality (and the fact that $\frac{1}{2} \leq \max_i p_i \leq 1$), we know that $\Pr\left[\max_i p_i \geq \frac{3}{4}\right] \geq 0.996$. In other words, the Bayes decision rule is usually confident in its predictions, so to speak.

Notice that p_i equals the probability (conditioned on postselection succeeding) that $\mathcal{B}$ outputs i if it starts with ρ in place of $\hat{\rho}$. For $i \in \{0, 1\}$, define $\hat{p}_i$ analogously as the postselcted output probabilities of $\mathcal{B}$ itself: $\hat{p}_i = \Pr\left[\mathcal{B}(b, c) = i \mid \mathcal{B}(b, c) \in \{0, 1\}\right]$. To argue that $\mathcal{A}$ is correct with 0.995 probability, it suffices to show that $\Pr\left[\max_i \hat{p}_i \geq \frac{2}{3} \wedge \arg\max_i \hat{p}_i = X\right] \geq 0.995$, as in this case the PostBQP promise is satisfied and the output of $\mathcal{L}$ agrees with X. We have that:

$$\begin{aligned}
\Pr\left[\max_i \hat{p}_i \geq \frac{2}{3} \wedge \arg\max_i \hat{p}_i = X\right] &\geq \Pr\left[\max_i p_i \geq \frac{3}{4} \wedge \arg\max_i p_i = X\right] \\
&\geq 1 - \Pr\left[\max_i p_i < \frac{3}{4}\right] - \Pr\left[\arg\max_i p_i \neq X\right] \\
&\geq 0.996 - \Pr\left[\arg\max_i p_i \neq X\right] \\
&\geq 0.995
\end{aligned}$$

Above, the first inequality follows from the assumption that S is a $\frac{1}{17}$-approximate T-design, because the acceptance probability of a postselected quantum algorithm can be viewed as the ratio of two probabilities:

$$\hat{p_i} = \frac{\Pr[\mathcal{B}(b,c) = i]}{\Pr[\mathcal{B}(b,c) \in \{0,1\}]}.$$

Fact 9 implies that both the numerator and denominator change by at most a multiplicative factor of $1 \pm \frac{1}{17}$ when switching between ρ and $\hat{\rho}$. So, if $p_i \geq \frac{3}{4}$, then $\hat{p_i} \geq \frac{3}{4} \cdot \frac{1-\frac{1}{17}}{1+\frac{1}{17}} = \frac{2}{3}$. The second inequality follows by a union bound, and the remaining inequalities were established above. ◀

We remark that the above theorem also holds relative to all oracles, in the sense that if the state generation algorithm G in the definition of the PRS (Definition 12) queries a classical or quantum oracle $\mathcal{U}$, then the corresponding ensemble of states can be distinguished from Haar-random by a polynomial-time quantum algorithm with a $\mathsf{PostBQP}^{\mathcal{U}}$ oracle.

4 Pseudorandomness from a Quantum Oracle

In this section, we construct a quantum oracle $(\mathcal{U}, \mathcal{P})$ relative to which $\mathsf{BQP} = \mathsf{QMA}$ and PRUs exist.

4.1 BQP = QMA Relative to $(\mathcal{U}, \mathcal{P})$

We start with a lemma showing that the acceptance probability of a quantum query algorithm, viewed as a function of the unitary transformation used in the query, is Lipschitz.

▶ **Lemma 16.** *Let $\mathcal{A}^U$ be quantum algorithm that makes T queries to $U \in \mathbb{U}(D)$, and define $f(U) = \Pr\left[\mathcal{A}^U = 1\right]$. Then f is $2T$-Lipschitz in the Frobenius norm.*

Proof. Suppose that $||U - V||_F \leq d$. By Lemma 6, this implies that the distance between U and V in the diamond norm is at most $2d$. The sub-additivity of the diamond norm under composition implies that as superoperators, $||\mathcal{A}^U - \mathcal{A}^V||_\diamond \leq 2Td$. By the definition of the diamond norm, it must be the case that $|f(U) - f(V)| \leq 2Td$. ◀

The next lemma extends Lemma 16 to QMA verifiers: we should think of $\mathcal{V}$ as a QMA verifier that receives a witness $|\psi\rangle$, in which case this lemma states that the maximum acceptance probability of $\mathcal{V}$ is Lipschitz with respect to the queried unitary.

▶ **Lemma 17.** *Let $\mathcal{V}^U(|\psi\rangle)$ be quantum algorithm that makes T queries to $U \in \mathbb{U}(D)$ and takes as input a quantum state $|\psi\rangle$ on some fixed (but arbitrary) number of qubits. Define $f(U) = \max_{|\psi\rangle} \Pr\left[\mathcal{V}^U(|\psi\rangle) = 1\right]$. Then f is $2T$-Lipschitz in the Frobenius norm.*

Proof. Note that f is well-defined because of the extreme value theorem. Define $f_\psi : \mathbb{U}(D) \to \mathbb{R}$ by:

$$f_\psi(U) = \Pr\left[\mathcal{V}^U(|\psi\rangle) = 1\right],$$

so that $f(U) = \max_{|\psi\rangle} f_\psi(U)$. Lemma 16 implies that f_ψ is $2T$-Lipschitz for every $|\psi\rangle$. Let $U, V \in \mathbb{U}(D)$, and suppose that $|\psi\rangle$ and $|\varphi\rangle$ are such that $f(U) = f_\psi(U)$ and $f(V) = f_\varphi(V)$.

Then:

$$\begin{aligned}|f(U)-f(V)| &= |f_\psi(U)-f_\varphi(V)| \\ &= \max\{f_\psi(U)-f_\varphi(V), f_\varphi(V)-f_\psi(U)\} \\ &\leq \max\{f_\psi(U)-f_\psi(V), f_\varphi(V)-f_\varphi(U)\} \\ &\leq 2T||U-V||_F,\end{aligned}$$

where the third line uses the fact that $f_\psi(V) \leq f_\varphi(V)$ and $f_\varphi(U) \leq f_\psi(U)$, and the last line uses the fact that f_ψ and f_φ are $2T$-Lipschitz. ◀

We are now ready to prove the first main result of this section, that $\mathsf{BQP}^{\mathcal{U},\mathcal{P}} = \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$.

▶ **Theorem 18.** *Let $\mathcal{U} = \{\mathcal{U}_n\}_{n\in\mathbb{N}}$ be a quantum oracle where each $\mathcal{U}_n$ is chosen randomly from $\mu_{2^n}^{2^n}$. Let $\mathcal{P}$ be an arbitrary* PSPACE*-complete language. Then with probability* 1 *over $\mathcal{U}$,* $\mathsf{BQP}^{\mathcal{U},\mathcal{P}} = \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$.

Proof. First, some notation. We view each $\mathcal{U}_n$ alternatively as either a unitary transformation on $2n$ qubits, or as a list of 2^n different n-qubit unitary transformations $\mathcal{U}_n = \{\mathcal{U}_{nm}\}_{m\in\{0,1\}^n}$ indexed by n-bit strings.

Let $\mathcal{L} \in \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$, which means that there exists a polynomial-time $\mathsf{QMA}^{\mathcal{U},\mathcal{P}}$ verifier $\mathcal{V}^{\mathcal{U},\mathcal{P}}(x,|\psi\rangle)$ with completeness $\frac{2}{3}$ and soundness $\frac{1}{3}$. Without loss of generality, we can amplify the completeness and soundness probabilities of $\mathcal{V}$ to $\frac{11}{12}$ and $\frac{1}{12}$, respectively. Let $p(n)$ be a polynomial upper bound on the running time of $\mathcal{V}$ on inputs of length n.

We now describe a $\mathsf{BQP}^{\mathcal{U},\mathcal{P}}$ machine $\mathcal{A}^{\mathcal{U},\mathcal{P}}(x)$ such that, with probability 1 over $\mathcal{U}$, $\mathcal{A}$ computes $\mathcal{L}$ on all but finitely many inputs $x \in \{0,1\}^*$. Let $d = \lfloor \log_2 \left(13824|x|p(|x|)^2+2\right) \rfloor$. For each $n \in [d]$, $\mathcal{A}$ performs process tomography on each $\mathcal{U}_n$, producing estimates $\widetilde{\mathcal{U}}_n$ such that $||\widetilde{\mathcal{U}}_n \cdot \widetilde{\mathcal{U}}_n^\dagger - \mathcal{U}_n \cdot \mathcal{U}_n^\dagger||_\diamond \leq \frac{1}{12p(|x|)}$ for every n, with probability at least $\frac{2}{3}$ over the randomness of $\mathcal{A}$.[6] Let $\mathcal{S}$ be the algorithm from Lemma 8 that samples from a $\frac{1}{12}$-approximate unitary $p(|x|)$-design on n qubits, given as input a random seed $r \leftarrow \{0,1\}^{k_n}$ where $k_n = \mathrm{poly}(n,|x|)$.

Consider a $\mathsf{QMA}^{\mathcal{P}}$ verifier $\widetilde{\mathcal{V}}^{\mathcal{P}}(x,|\psi\rangle)$ that simulates $\mathcal{V}^{\mathcal{U},\mathcal{P}}(x,|\psi\rangle)$ by replacing queries to $\mathcal{U}$ as follows. For each $n \in [d+1, p(|x|)]$, $\widetilde{\mathcal{V}}$ samples a function $f_n : \{0,1\}^n \to \{0,1\}^{k_n}$ from a $2p(|x|)$-wise independent function family. Then, for $n \in [d]$, queries to $\mathcal{U}_n$ are replaced by queries to $\widetilde{\mathcal{U}}_n$. For $n \in [d+1, p(|x|)]$ and $m \in \{0,1\}^n$, queries to $\mathcal{U}_{nm}$ are replaced by queries to $\mathcal{S}(f_n(m))$ (i.e. the mth unitary in $\mathcal{U}_n$ is replaced by an element of the $p(|x|)$-design, selected by $f_n(m)$). Consider the $\mathsf{QMA}^{\mathcal{P}}$ promise problem $\widetilde{\mathcal{L}}$ corresponding to $\widetilde{V}$ with completeness $\frac{2}{3}$ and soundness $\frac{1}{3}$. Since $\mathsf{QMA}^A \subseteq \mathsf{PSPACE}^A$ for all classical oracles A, $\widetilde{\mathcal{L}} \in \mathsf{PSPACE}^{\mathcal{P}} = \mathsf{PSPACE}$, so $\mathcal{A}$ can decide $\widetilde{\mathcal{L}}(x)$ with a single query to $\mathcal{P}$. $\mathcal{A}$ does this, and outputs $\widetilde{\mathcal{L}}(x)$.

We now argue that for any x, with high probability over $\mathcal{U}$, $\Pr\left[\mathcal{A}^{\mathcal{U},\mathcal{P}}(x) = \mathcal{L}(x)\right] \geq \frac{2}{3}$. It will be convenient to define several hybrid verifiers:

- $V_1 = \mathcal{V}$.
- V_2: For each $n \in [d+1, p(|x|)]$, chooses a matrix $U_n \leftarrow \mu_{2^n}^{2^n}$. Simulates V_1, replacing queries to $\mathcal{U}_n$ by U_n for $n \in [d+1, p(|x|)]$.

[6] There are many ways to accomplish this. For instance, one can use the Choi-Jamiołkowski isomorphism and quantum state tomography [23] to estimate the Choi state of $\mathcal{U}_n$ to inverse polynomial (in 2^n) error in trace distance. The estimated unitary transformation $\widetilde{\mathcal{U}}_n$ can then be compiled to a circuit using $2^{O(n)}$ 1- and 2-qubit gates [26]. Since $n \leq d = O(\log|x|)$, this can be done in polynomial time.

- V_3: For each $n \in [d+1, p(|x|)]$, samples a function $g_n : \{0,1\}^n \to \{0,1\}^{k_n}$ uniformly at random. Simulates V_2, replacing queries to U_{nm} by by queries to $\mathcal{S}(g_n(m))$ for $n \in [d+1, p(|x|)]$ and $m \in \{0,1\}^n$.
- V_4: For each $n \in [d+1, p(|x|)]$, samples a function $f_n : \{0,1\}^n \to \{0,1\}^{k_n}$ from a $2p(|x|)$-wise independent function family. Simulates V_3, replacing queries to g_n by f_n.
- V_5: Simulates V_4, replacing queries to $\mathcal{U}_n$ by queries to $\widetilde{\mathcal{U}}_n$ for $n \in [d]$. Note that V_5 and $\widetilde{\mathcal{V}}$ are equivalent, and that of these hybrids, V_5 is the only one whose output depends on the randomness of $\mathcal{A}$ (by way of $\widetilde{\mathcal{U}_n}$).

For $i \in [5]$ and a fixed choice of $\mathcal{U}$, define:

$$\mathrm{acc}(V_i) = \max_{|\psi\rangle} \Pr\left[V_i(x, |\psi\rangle) = 1\right],$$

which is well defined by the extreme value theorem. We now bound $|\mathrm{acc}(V_i) - \mathrm{acc}(V_{i-1})|$ for various i:

- By Lemma 17, because $\mathcal{V}$ makes at most $p(|x|)$ queries, we know that $\mathrm{acc}(V_1)$, viewed as a function of $(\mathcal{U}_{d+1}, \mathcal{U}_{d+2}, \ldots, \mathcal{U}_{p(|x|)})$, is $2p(|x|)$-Lipschitz in the Frobenius norm. Hence, by Theorem 11 with $N = 13824|x|p(|x|)^2 + 2$, $L = 2p(|x|)$, and $t = \frac{1}{12}$, we have that:

$$\begin{aligned}\Pr_{\mathcal{U}}\left[|\mathrm{acc}(V_1) - \mathrm{acc}(V_2)| \geq \frac{1}{12}\right] &\leq 2\exp\left(-\frac{(N-2)t^2}{24L^2}\right) \\ &= 2\exp\left(-\frac{13824|x|p(|x|)^2 \cdot \frac{1}{144}}{96p(|x|)^2}\right) \\ &= 2e^{-|x|}.\end{aligned}$$

The factor of 2 appears because Theorem 11 applies to one-sided error, but the absolute value forces us to consider two-sided error.

- Fact 10 and the assumption that $\mathcal{S}$ samples from a $\frac{1}{12}$-approximate unitary $p(|x|)$-design implies that for any fixed $|\psi\rangle$, $|\Pr[V_2(x, |\psi\rangle) = 1] - \Pr[V_3(x, |\psi\rangle) = 1]| \leq \frac{1}{12}$. This in turn implies that $|\mathrm{acc}(V_2) - \mathrm{acc}(V_3)| \leq \frac{1}{12}$.
- Zhandry [28, Theorem 6.1] shows that a quantum algorithm that makes T queries to a random function can be exactly simulated by the same algorithm with T queries to a $2T$-wise independent function, so $\mathrm{acc}(V_3) = \mathrm{acc}(V_4)$.
- Because $||\widetilde{\mathcal{U}}_n \cdot \widetilde{\mathcal{U}}_n^\dagger - \mathcal{U}_n \cdot \mathcal{U}_n^\dagger||_\diamond \leq \frac{1}{12p(|x|)}$ for each $n \in [d]$ with probability at least $\frac{2}{3}$ over $\mathcal{A}$, from the definition of the diamond norm [5] and because $\mathcal{V}$ makes at most $p(|x|)$ queries, it holds that $\Pr_{\mathcal{A}}\left[|\mathrm{acc}(V_4) - \mathrm{acc}(V_5)| \geq \frac{1}{12}\right] \leq \frac{1}{3}$.

Putting these bounds together, we have that, except with probability $2e^{-|x|}$ over $\mathcal{U}$:

$$\begin{aligned}\mathcal{L}(x) = 1 \quad &\implies \quad \Pr_{\mathcal{A}}\left[\max_{|\psi\rangle} \Pr\left[\widetilde{\mathcal{V}}(x, |\psi\rangle) = 1\right] \geq \frac{2}{3}\right] \geq \frac{2}{3} \\ \mathcal{L}(x) = 0 \quad &\implies \quad \Pr_{\mathcal{A}}\left[\max_{|\psi\rangle} \Pr\left[\widetilde{\mathcal{V}}(x, |\psi\rangle) = 1\right] \leq \frac{1}{3}\right] \geq \frac{2}{3}.\end{aligned}$$

This is to say that $\mathcal{A}$ correctly decides $\mathcal{L}(x)$, expect with probability at most $2e^{-|x|}$ over $\mathcal{U}$. By the Borel-Cantelli Lemma, because $\sum_{i=1}^{\infty} 2^i \cdot 2e^{-i} = \frac{4}{e-2} < \infty$, $\mathcal{A}$ correctly decides $\mathcal{L}(x)$ for all but finitely many $x \in \{0,1\}^*$, with probability 1 over $\mathcal{U}$. Hence, with probability 1 over $\mathcal{U}$, $\mathcal{A}$ can be modified into an algorithm $\mathcal{A}'$ that agrees with $\mathcal{L}$ on every $x \in \{0,1\}^*$, by simply hard-coding those x on which $\mathcal{A}$ and $\mathcal{L}$ disagree.

Because there are only countably many $\mathsf{QMA}^{\mathcal{U},\mathcal{P}}$ machines, we can union bound over all $\mathcal{L} \in \mathsf{QMA}^{\mathcal{U},\mathcal{P}}$ to conclude that $\mathsf{QMA}^{\mathcal{U},\mathcal{P}} \subseteq \mathsf{BQP}^{\mathcal{U},\mathcal{P}}$ with probability 1 over $\mathcal{U}$. ◀

4.2 PRUs Relative to $(\mathcal{U}, \mathcal{P})$

We proceed to the second part of the oracle construction, showing that PRUs exist relative to $(\mathcal{U}, \mathcal{P})$. We begin with a lemma establishing that the average advantage of a polynomial-time adversary is small against our PRU construction. Here, we should think of $\{U_k\}_{k\in[N]}$ as the PRU ensemble.

▶ **Lemma 19.** *Consider a quantum algorithm $\mathcal{A}^{O,U}$ that makes T queries to $O \in \mathbb{U}(D)$ and $U = (U_1, \ldots, U_N) \in \mathbb{U}(D)^N$. For fixed U, define:*

$$\mathrm{adv}(\mathcal{A}^U) := \Pr_{k\leftarrow[N]}\left[\mathcal{A}^{U_k,U} = 1\right] - \Pr_{O\leftarrow\mu_D}\left[\mathcal{A}^{O,U} = 1\right].$$

Then there exists a universal constant $c > 0$ such that:

$$\mathop{\mathbb{E}}_{U\leftarrow\mu_D^N}\left[\mathrm{adv}(\mathcal{A}^U)\right] \leq \frac{cT^2}{N}.$$

Proof. Our strategy is to reduce to the quantum query lower bound for unstructured search. Intuitively, if $\mathcal{A}$ could identify whether $O \in \{U_1, \ldots, U_N\}$ or not, then $\mathcal{A}$ could be modified into a quantum algorithm $\mathcal{B}$ that finds a single marked item from a list of size N. Then the BBBV theorem [10] forces T to be $\Omega\left(\sqrt{N}\right)$.

More formally, we construct an algorithm $\mathcal{B}(x)$ that queries a string $x \in \{0,1\}^N$ as follows. $\mathcal{B}$ draws a unitary $V = (V_0, V_1, \ldots, V_N) \in \mathbb{U}(D)^{N+1}$ from μ_D^{N+1}. Then, $\mathcal{B}$ runs $\mathcal{A}$, replacing queries to O by queries to V_0, and replacing queries to $U_k \in U$ by V_0 if $x_k = 1$ and by V_k if $x_k = 0$.

Let $e_k \in \{0,1\}^N$ be the string with 1 in the kth position and 0s everywhere else. We have that:

$$\begin{aligned}
\mathop{\mathbb{E}}_{U\leftarrow\mu_D^N}\left[\mathrm{adv}(\mathcal{A}^U)\right] &= \mathop{\mathbb{E}}_{U\leftarrow\mu_D^N}\left[\Pr_{k\leftarrow[N]}\left[\mathcal{A}^{U_k,U} = 1\right]\right] - \mathop{\mathbb{E}}_{U\leftarrow\mu_D^N}\left[\Pr_{O\leftarrow\mu_D}\left[\mathcal{A}^{O,U} = 1\right]\right] \\
&= \Pr_{k\leftarrow[N]}\left[\mathcal{B}\left(e_k\right) = 1\right] - \Pr\left[\mathcal{B}\left(0^N\right) = 1\right] \\
&\leq \frac{cT^2}{N}.
\end{aligned}$$

Above, the first line applies linearity of expectation, the second line holds by definition of $\mathcal{B}$, and the third line holds for some universal c by the BBBV theorem [10]. ◀

The next lemma uses Lemma 19 to show that the advantage of $\mathcal{A}$ is small with extremely high probability, which follows from the strong concentration properties of the Haar measure (Theorem 11).

▶ **Lemma 20.** *Consider a quantum algorithm $\mathcal{A}^{O,U}$ that makes T queries to $O \in \mathbb{U}(D)$ and $U = (U_1, \ldots, U_N) \in \mathbb{U}(D)^N$. Let $\mathrm{adv}(\mathcal{A}^U)$ be defined as in Lemma 20. Then there exists a universal constant $c > 0$ such that for any p,*

$$\Pr_{U\leftarrow\mu_D^N}\left[\left|\mathrm{adv}(\mathcal{A}^U)\right| \geq p\right] \leq 2\exp\left(-\frac{(D-2)\left(p - cT^2/N\right)^2}{384T^2}\right).$$

Proof. By Lemma 16, $\mathrm{adv}(\mathcal{A}^U)$ is $4T$-Lipschitz as a function of U, because $\mathrm{adv}(\mathcal{A}^U)$ can be expressed as the the difference between the acceptance probabilities of two algorithms that each make T queries to U. Combining Lemma 19 and Theorem 11, we obtain:

$$\Pr_{U\leftarrow\mu_D^N}\left[\mathrm{adv}(\mathcal{A}^U) \geq p\right] \leq \exp\left(-\frac{(D-2)\left(p - cT^2/N\right)^2}{384T^2}\right).$$

Similar reasoning yields the same upper bound on $\Pr_{U\leftarrow\mu_D^N}\left[\mathrm{adv}(\mathcal{A}^U)\leq -p\right]$, so we get the final bound (with an additional factor of 2) by a union bound. ◀

Completing the security proof of the PRU construction amounts to combining Lemma 20 with a union bound over all possible polynomial-time adversaries.

▶ **Theorem 21.** *Let $\mathcal{U}=\{\mathcal{U}_n\}_{n\in\mathbb{N}}$ be a quantum oracle where each $\mathcal{U}_n$ is chosen randomly from $\mu_{2^n}^{2^n}$. Let $\mathcal{P}$ be an arbitrary* PSPACE*-complete language. Then with probability 1 over $\mathcal{U}$, there exists a family of PRUs relative to $(\mathcal{U},\mathcal{P})$.*

Proof. Fix an input length $n\in\mathbb{N}$. We take the key set $\mathcal{K}=[2^n]$ and take the PRU family to be $\{U_k\}_{k\in\mathcal{K}}$, where $\mathcal{U}_n=(U_1,U_2,\ldots,U_{2^n})\in\mathbb{U}(2^n)^{2^n}$. In words, the family consists of the 2^n different Haar-random n-qubit unitaries specified by $\mathcal{U}_n$.

Without loss of generality, assume the adversary is a uniform polynomial-time quantum algorithm $\mathcal{A}^{O,\mathcal{U},\mathcal{P}}(1^n,x)$, where $x\in\{0,1\}^{\mathrm{poly}(n)}$ is the advice and $O\in\mathbb{U}(2^n)$ is the oracle that the adversary seeks to distinguish as pseudorandom or Haar-random.

By Lemma 20 with $N=D=2^n$ and $T=\mathrm{poly}(n)$, for any fixed $x\in\{0,1\}^{\mathrm{poly}(n)}$, $\mathcal{A}^{O,\mathcal{U},\mathcal{P}}(1^n,x)$ achieves non-negligible advantage with extremely low probability over $\mathcal{U}$. This is to say that for any $p=\frac{1}{\mathrm{poly}(n)}$:

$$\Pr_{\mathcal{U}_n\leftarrow\mu_{2^n}^{2^n}}\left[\left|\Pr_{k\leftarrow[2^n]}\left[\mathcal{A}^{U_k,\mathcal{U},\mathcal{P}}(1^n,x)=1\right]-\Pr_{O\leftarrow\mu_{2^n}}\left[\mathcal{A}^{O,\mathcal{U},\mathcal{P}}(1^n,x)=1\right]\right|\geq p\right]\leq\exp\left(-\frac{2^n}{\mathrm{poly}(n)}\right).$$

By a union bound over all $x\in\{0,1\}^{\mathrm{poly}(n)}$, $\mathcal{A}^{O,\mathcal{U},\mathcal{P}}(1^n,x)$ achieves advantage larger than p for *any* $x\in\{0,1\}^{\mathrm{poly}(n)}$ with probability at most $2^{\mathrm{poly}(n)}\cdot\exp\left(-\frac{2^n}{\mathrm{poly}(n)}\right)\leq\mathrm{negl}(n)$. Hence, by the Borel-Cantelli lemma, $\mathcal{A}$ achieves negligible advantage for all but finitely many input lengths $n\in\mathbb{N}$ with probability 1 over $\mathcal{U}$, as $\sum_{n=1}^{\infty}\mathrm{negl}(n)<\infty$. This is to say that $\{U_k\}_{k\in\mathcal{K}}$ defines a PRU ensemble. ◀

We expect that using the techniques of Chung, Guo, Liu, and Qian [16], one can extend Theorem 21 to a security proof against adversaries with quantum advice. Some version of [16, Theorem 5.14] likely suffices. The idea is that breaking the PRU should remain hard even if $\mathcal{A}$ could query an explicit description of O and explicit descriptions of U_k for $k\in[2^n]$, which is a strictly more powerful model. But then this corresponds to the security game defined in [16, Definition 5.12], except that the range of the random oracle is $\mathbb{U}(D)$ rather than the finite set $[M]$.

References

1 Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461:3473–3482, 2005.

2 Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 325–338, New York, NY, USA, 2018. Association for Computing Machinery.

3 Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, November 2004.

4 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007.

5 Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 20–30, New York, NY, USA, 1998. Association for Computing Machinery.

6 Eric Allender. *The Complexity of Complexity*, pages 79–94. Springer International Publishing, Cham, 2017.

7 Eric Allender. The new complexity landscape around circuit minimization. In Alberto Leporati, Carlos Martín-Vide, Dana Shapira, and Claudio Zandron, editors, *Language and Automata Theory and Applications*, pages 3–16, Cham, 2020. Springer International Publishing.

8 Andris Ambainis and Joseph Emerson. Quantum t-designs: T-wise independence in the quantum world. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity*, CCC '07, pages 129–140, USA, 2007. IEEE Computer Society.

9 Aleksandrs Belovs and Ansis Rosmanis. Tight quantum lower bound for approximate counting with quantum states, 2020.

10 Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997.

11 James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. Springer New York, 2013.

12 Adam Bouland, Bill Fefferman, and Umesh Vazirani. Computational Pseudorandomness, the Wormhole Growth Paradox, and Constraints on the AdS/CFT Duality (Abstract). In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 63:1–63:2, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

13 Zvika Brakerski and Omri Shmueli. (Pseudo) random quantum states with binary phase. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography*, pages 229–250, Cham, 2019. Springer International Publishing.

14 Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Efficient quantum pseudorandomness. *Phys. Rev. Lett.*, 116:170502, April 2016.

15 Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, 2016.

16 Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 673–684, 2020.

17 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 2020.

18 Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, pages 44–61, New York, NY, USA, 1989. Association for Computing Machinery.

19 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 126–152, Cham, 2018. Springer International Publishing.

20 Robert Koenig and John A. Smolin. How to efficiently select an arbitrary Clifford group element. *Journal of Mathematical Physics*, 55(12):122202, 2014.

21 Greg Kuperberg. How hard is it to approximate the Jones polynomial? *Theory of Computing*, 11(6):183–219, 2015.

22 Elizabeth S. Meckes. *The Random Matrix Theory of the Classical Compact Groups*. Cambridge Tracts in Mathematics. Cambridge University Press, 2019.

23 Ryan O'Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 899–912, New York, NY, USA, 2016. Association for Computing Machinery.

24 Leonard Susskind. Addendum to computational complexity and black hole horizons. *Fortschritte der Physik*, 64(1):44–48, 2016.

25 Leonard Susskind. Computational complexity and black hole horizons. *Fortschritte der Physik*, 64(1):24–43, 2016.

26 Juha J. Vartiainen, Mikko Möttönen, and Martti M. Salomaa. Efficient decomposition of quantum gates. *Phys. Rev. Lett.*, 92:177902, April 2004.

27 Yaakov S. Weinstein, Winton G. Brown, and Lorenza Viola. Parameters of pseudorandom quantum circuits. *Phys. Rev. A*, 78:052332, November 2008.

28 Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, pages 758–775, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

A Proof of Lemma 6

Proof. Let $\{\lambda_i : i \in [N]\}$ denote the eigenvalues of $UV^\dagger$. Then we have:

$$\begin{aligned}
||U-V||_F^2 &= \mathrm{Tr}\left((U-V)(U-V)^\dagger\right) \\
&= \mathrm{Tr}(2I - UV^\dagger - VU^\dagger) \\
&= 2N - \sum_{i=1}^{N}(\lambda_i + \lambda_i^*) \\
&= \sum_{i=1}^{N}(2 - 2\mathrm{Re}(\lambda_i)) \\
&\geq \max_i(2 - 2\mathrm{Re}(\lambda_i)), \qquad (1)
\end{aligned}$$

where $\mathrm{Re}(\lambda_i)$ denotes the real part of λ_i. The last line holds because the eigenvalues of a unitary matrix have absolute value 1. Aharonov, Kitaev, and Nisan [5] show that $||U \cdot U^\dagger - V \cdot V^\dagger||_\diamond = 2\sqrt{1-d^2}$, where d is the distance in the complex plane between 0 and the polygon whose vertices are $\lambda_1, \ldots, \lambda_N$. From this we may conclude:

$$\begin{aligned}
||U \cdot U^\dagger - V \cdot V^\dagger||_\diamond &\leq \max_i 2\sqrt{1 - \max\{\mathrm{Re}(\lambda_i), 0\}^2} \\
&\leq \max_i 2\sqrt{2 - 2\mathrm{Re}(\lambda_i)} \\
&\leq 2||U-V||_F,
\end{aligned}$$

where the first inequality uses the fact that either all of the eigenvalues have positive real components and therefore $d \geq \min_i \mathrm{Re}(\lambda_i)$, or else $d \geq 0$; the second inequality substitutes $1 - \max\{x, 0\}^2 \leq 2 - 2x$ which holds for all $x \in \mathbb{R}$; and the third inequality substitutes (1). ◀

B Complexity Classes

▶ **Definition 22.** *A promise problem $\mathcal{L} = (\mathcal{L}_{\mathrm{yes}}, \mathcal{L}_{\mathrm{no}})$ is in* QMA *if there exists a polynomial-time quantum algorithm $\mathcal{V}(x, |\psi\rangle)$ called a* QMA verifier *and a polynomial p such that:*

1. *(Completeness) If $x \in \mathcal{L}_{\mathrm{yes}}$, then there exists a state $|\psi\rangle$ (called a* witness *or* proof*) on $p(|x|)$ qubits such that $\Pr[\mathcal{V}(x, |\psi\rangle) = 1] \geq \frac{2}{3}$.*
2. *(Soundness) If $x \in \mathcal{L}_{\mathrm{no}}$, then for every state $|\psi\rangle$ on $p(|x|)$ qubits, $\Pr[\mathcal{V}(x, |\psi\rangle) = 1] \leq \frac{1}{3}$.*

Aaronson [1] defined PostBQP as follows, and showed that PostBQP = PP.

▶ **Definition 23.** *A promise problem* $\mathcal{L} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$ *is in* PostBQP *if there exists a polynomial-time quantum algorithm* $\mathcal{A}(x)$ *that outputs a trit* $\{0, 1, *\}$ *such that:*

1. *If* $x \in \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}}$*, then* $\Pr[\mathcal{A}(x) \in \{0,1\}] > 0$*. When* $\mathcal{A}(x) \in \{0,1\}$*, we say that* postselection succeeds*.*
2. *If* $x \in \mathcal{L}_{\text{yes}}$*, then* $\Pr[\mathcal{A}(x) = 1 \mid \mathcal{A}(x) \in \{0,1\}] \geq \frac{2}{3}$*. In other words, conditioned on postselection succeeding,* $\mathcal{A}$ *outputs* 1 *with at least* $\frac{2}{3}$ *probability.*
3. *If* $x \in \mathcal{L}_{\text{no}}$*, then* $\Pr[\mathcal{A}(x) = 1 \mid \mathcal{A}(x) \in \{0,1\}] \leq \frac{1}{3}$*. In other words, conditioned on postselection succeeding,* $\mathcal{A}$ *outputs* 1 *with at most* $\frac{1}{3}$ *probability.*

Technically, the definition of PostBQP is sensitive to the choice of universal gate set used to specify quantum algorithms, as was observed by Kuperberg [21]. However, for most "reasonable" gate sets, such as unitary gates with algebraic entries [21], the choice of gate set is irrelevant. We assume such a gate set, e.g. $\{\text{CNOT}, H, T\}$.

We consider versions of BQP, QMA, and PostBQP augmented with quantum oracles, where the algorithm (or in the case of QMA, the verifier) can apply unitary transformations from an infinite sequence $\mathcal{U} = \{\mathcal{U}_n\}_{n\in\mathbb{N}}$. We denote the respective complexity classes by $\text{BQP}^{\mathcal{U}}$, $\text{QMA}^{\mathcal{U}}$, and $\text{PostBQP}^{\mathcal{U}}$. We assume the algorithm incurs a cost of n to query $\mathcal{U}_n$ so that a polynomial-time algorithm on input x can query $\mathcal{U}_n$ for any $n \leq \text{poly}(|x|)$. In this model, a query to $\mathcal{U}_n$ consists of a single application of either $\mathcal{U}_n$, $\mathcal{U}_n^\dagger$, or controlled versions of $\mathcal{U}_n$ or $\mathcal{U}_n^\dagger$.

The quantum oracle model includes classical oracles as a special case. For a language $\mathcal{L}$, a query to $\mathcal{L}$ is implemented via the unitary transformation $\mathcal{U}$ that acts as $\mathcal{U}|x\rangle|b\rangle = |x\rangle|b \oplus \mathcal{L}(x)\rangle$.

C PRSs with Binary Phases

In this section, we sketch a proof that a PRS construction proposed by Ji, Liu, and Song [19] and shown secure by Brakerski and Shmueli [13] can be broken efficiently with an NP oracle. The PRS family is based on pseudorandom functions (PRFs). Let $\{f_k\}_{k\in\mathcal{K}}$ be a PRF family of functions $f_k : \{0,1\}^n \to \{0,1\}$ keyed by $\mathcal{K}$. The corresponding PRS family is the set of states $\{|\varphi_k\rangle\}_{k\in\mathcal{K}}$ given by:

$$|\varphi_k\rangle = \frac{1}{2^{n/2}} \sum_{x\in\{0,1\}^n} (-1)^{f_k(x)} |x\rangle .$$

For simplicity, suppose that each f_k is *balanced*, meaning that $|f_k^{-1}(0)| = |f_k^{-1}(1)| = 2^{n-1}$. Consider the quantum circuit below:

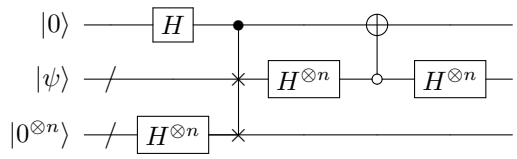

Observe that if $|\psi\rangle = |\varphi_k\rangle$, then this circuit produces the state $|0\rangle \frac{|\varphi_k\rangle|+\rangle^{\otimes n} + |+\rangle^{\otimes n}|\varphi_k\rangle}{\sqrt{2}}$ from a single copy of $|\varphi_k\rangle$. Notice that if we measure the resulting state in the computational basis, then we observe $|0\rangle|x\rangle|y\rangle$ with nonzero probability for $x, y \in \{0,1\}^n$ if and only if $f_k(x) = f_k(y)$. This is because the amplitude on this basis state is given by:

$$\langle x|\langle y| \frac{|\varphi_k\rangle|+\rangle^{\otimes n} + |+\rangle^{\otimes n}|\varphi_k\rangle}{\sqrt{2}} = \frac{(-1)^{f_k(x)} + (-1)^{f_k(y)}}{2^n\sqrt{2}}.$$

Furthermore, this shows that we in fact sample a uniformly random pair (x, y) such that $f_k(x) = f_k(y)$.

Suppose that given a state $|\psi\rangle$ which is either pseudorandom or Haar-random, we repeat this procedure $\mathrm{poly}(n)$ times to obtain a list of pairs $\{(x_i, y_i)\}$. It is an NP problem to decide whether there exists a k such that $f_k(x_i) = f_k(y_i)$ for all i. If $|\psi\rangle = |\varphi_k\rangle$ for some k then this NP language always returns true, while if $|\psi\rangle$ is Haar-random, this NP language returns true with negligible probability, so long as we take sufficiently many samples (x_i, y_i).

In the case where f_k is not perfectly balanced, we simply observe that the above procedure still works with good probability so long as f_k is *close* to a balanced function. But PRFs must be close to balanced functions, in the sense that for most $k \in \mathcal{K}$, it must be possible to change a $\mathrm{negl}(n)$ fraction of the outputs of f_k to turn it into a balanced function. Otherwise, the PRF family could be distinguished efficiently from random functions, which are $\mathrm{negl}(n)$-close to balanced with high probability.