

The Hardness of Learning Quantum Circuits and Its Cryptographic Applications

Bill Fefferman  

University of Chicago, IL, USA

Soumik Ghosh 

University of Chicago, IL, USA

Makrand Sinha 

University of Illinois Urbana-Champaign, IL, USA

Henry Yuen  

Columbia University, New York, NY, USA

Abstract

We show that concrete hardness assumptions about learning or cloning the output state of a random quantum circuit can be used as the foundation for secure quantum cryptography. In particular, under these assumptions we construct secure one-way state generators (OWSGs), digital signature schemes, quantum bit commitments, and private key encryption schemes. We also discuss evidence for these hardness assumptions by analyzing the best-known quantum learning algorithms, as well as proving black-box lower bounds for cloning and learning given state preparation oracles.

Our random circuit-based constructions provide concrete instantiations of quantum cryptographic primitives whose security do not depend on the existence of one-way functions. The use of random circuits in our constructions also opens the door to NISQ-friendly quantum cryptography. We discuss noise tolerant versions of our OWSG and digital signature constructions which can potentially be implementable on noisy quantum computers connected by a quantum network. On the other hand, they are still secure against noiseless quantum adversaries, raising the intriguing possibility of a useful implementation of an end-to-end cryptographic protocol on near-term quantum computers. Finally, our explorations suggest that the rich interconnections between learning theory and cryptography in classical theoretical computer science also extend to the quantum setting.

2012 ACM Subject Classification Security and privacy → Cryptography; Theory of computation → Quantum complexity theory

Keywords and phrases quantum learning, quantum circuits, cryptographic hardness, one-way state generators

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.56

Related Version *Full Version:* <https://arxiv.org/abs/2504.15343> [28]

Funding B.F. and S.G. acknowledge support from AFOSR (FA9550-21-1-0008). This material is based upon work partially supported by the National Science Foundation under Grant CCF-2044923 (CAREER), by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers (Q-NEXT) and by the DOE QuantISED grant DE-SC0020360. H.Y. acknowledges support from AFOSR award FA9550-23-1-0363, NSF CAREER award CCF-2144219, NSF award CCF-2329939, and the Sloan Foundation. M.S. acknowledges support from the NSF award QCIS-FF: Quantum Computing & Information Science Faculty Fellow at the University of Illinois Urbana-Champaign (NSF 1955032). This work was done in part while some of the authors were visiting the Simons Institute for the Theory of Computing, supported by NSF QLCI Grant No. 2016245.

Acknowledgements We thank John Bostanci, Jonas Helsen, Yihui Quek, and Abhinav Deshpande for helpful discussions.



© Bill Fefferman, Soumik Ghosh, Makrand Sinha, and Henry Yuen; licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 56; pp. 56:1–56:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The hardness of learning and modern cryptography are inextricably linked in the world of classical computing. On one hand, hard cryptographic problems have served as the basis of showing that various learning tasks are intractable (see e.g., [40, 44, 25]). Conversely, the hardness of various learning tasks have been used to construct fundamental cryptographic primitives, like pseudorandom functions [37, 13, 62], and practical cryptographic protocols [65, 52]. Furthermore, this connection between learning and cryptography extends further and has shed light on fundamental questions in the related areas of pseudorandomness [60] and circuit lower bounds [50, 20].

In the quantum world, our understanding of the analogous interconnections is quite lacking. In one direction, some prior works have used classical cryptographic assumptions, like quantum-secure one-way functions or LWE, to argue about the hardness of learning quantum states and circuits in different contexts (see [9, 70]), but the complexity of many fundamental quantum learning tasks remain open. However, the *converse* direction – using hardness of quantum learning as a foundation for cryptography – has not received much attention, unlike the classical case¹. Arguably, one difficulty in establishing such a connection is that classical cryptographic primitives appear insufficient to fully capture the inherently quantum nature of tasks involving learning quantum states or unitaries [23]. This raises an interesting question that forms the basis of this paper:

Can we base quantum cryptography on the hardness of quantum learning?

Our results

We propose several fine-grained assumptions about quantum learning tasks related to random quantum circuits, give evidence for it in the black-box model, and use these assumptions to construct several quantum cryptographic primitives. Incidentally, this exploration also leads to several interesting results such as “NISQ-friendly” quantum cryptography and a deeper understanding of various quantum learning tasks and cryptographic primitives. In more detail:

- **Quantum hardness of learning assumptions.** We posit that it is computationally intractable to learn the output states of *unknown* random quantum circuits (Computational No-Learning Assumption) or to clone them (Computational No-Cloning Assumption). We formulate these conjectures and discuss the differences between them in Section 1.1. We also give evidence for these conjectures by proving lower bounds in the black-box model.
- **Cryptography from hardness of quantum learning.** We show that these hardness of learning assumptions can be used to construct quantum cryptographic primitives. In particular, we construct one way state generators (OWSGs) and quantum-secure digital signatures based on the Computational No-Learning Assumption, and we construct quantum bit commitments based on the Computational No-Cloning Assumption. Recent results have uncovered the tantalizing possibility of reducing the hardness assumptions behind such primitives beyond what is possible in classical cryptography [8, 56, 19, 46, 47]. Our hardness assumptions about random circuits do not appear to imply the existence of classical one-way functions, and thus our constructions concretely instantiate this possibility.

¹ As of the last several months, this is starting to change; see Section 1.5 for discussion of recent work on this topic.

- **Quantum cryptography for the NISQ era.** The use of random circuits in our constructions also opens the door to *NISQ-friendly quantum cryptography*. We demonstrate that our OWSG and digital signature constructions can be made noise tolerant and thus implementable on a noisy quantum computer, provided we have a quantum channel to pass quantum states as cryptographic keys. On the other hand, they are still secure against *noiseless* quantum adversaries. This raises the intriguing possibility of a near-term demonstration of *quantum cryptographic advantage*: using a NISQ device to run a cryptographic protocol whose security relies on fewer assumptions than what is possible in classical cryptography.

1.1 Our hardness of learning assumptions

We propose two main assumptions regarding the hardness of quantum learning: the Computational No-Learning Assumption and the Computational No-Cloning Assumption. We precisely define these conjectures and discuss evidence for them.

1.1.1 The Computational No-Learning Assumption

In what follows, we let \mathcal{C} denote a class of quantum circuits; for example, a concrete choice of \mathcal{C} is the ensemble of 1D brickwork circuits with depth $d = \log^2 n$.²

► **Conjecture 1** (Computational No-Learning). *Let \mathcal{C} be a class of circuits. The $(\mathcal{C}, \varepsilon, \delta)$ -Computational No-Learning Assumption stipulates the following. For all polynomial-time quantum algorithms A , for all sufficiently large n , we have that*

$$\mathbb{P} \left[\begin{array}{c} D \in \mathcal{C}_n \\ |\langle C|D \rangle|^2 \geq \varepsilon \end{array} : \begin{array}{c} C \leftarrow \mathcal{C}_n \\ D \leftarrow A(|C\rangle^{\otimes \text{poly}(n)}) \end{array} \right] \leq \delta$$

Here, $C \leftarrow \mathcal{C}_n$ denotes that C is sampled uniformly from the subset of circuits in \mathcal{C} that act on n qubits, and D is sampled from the output of the algorithm A given polynomially-many copies of $|C\rangle = C|0^n\rangle$. The output D is interpreted as a description of an n -qubit quantum circuit, and $|D\rangle = D|0^n\rangle$.

In words, the Computational No-Learning Assumption says that it a polynomial-time quantum algorithm A , given only polynomially-many copies of the output $|C\rangle$ of a randomly-chosen circuit C from the circuit ensemble \mathcal{C} , cannot produce a classical description of another circuit D from the same class \mathcal{C}_n and whose output state $|D\rangle$ approximates $|C\rangle$. We note that the algorithm A only has access to copies of the state³ generated by the random circuit and the circuit description itself is *not* known to the algorithm. Furthermore, Conjecture 1 actually refers to an entire *family* of assumptions, one for each choice of circuit class \mathcal{C} and parameters (ε, δ) . Throughout this paper we consider a fixed circuit family for convenience (e.g., 1D brickwork circuits with depth $\log^2 n$ unless otherwise specified). We elaborate on the different parameter settings for (ε, δ) below. For clarity we abbreviate “Computational No-Learning” as “No-Learning”. Oftentimes we will abbreviate “ $(\varepsilon, \varepsilon)$ -No-Learning” as “ ε -No-Learning”.

² We pick out $d = \log^2 n$ here because it happens to be a depth at which the existing state-of-the-art learning algorithms, like those in [35, 49, 43], fail to be efficient. Even if the learning algorithms were to be improved, any depth beyond which the efficiency of those learning algorithms fail is a good design choice for our assumptions and cryptographic constructions.

³ We stress that $|C\rangle$ does not mean the classical description of C , but rather the state generated by the circuit C on the all zero input.

1.1.1.1 Relationship between parameter values

(ε, δ) -No-Learning implies (ε', δ') -No-Learning for all $\varepsilon' \geq \varepsilon$ and $\delta' \geq \delta$. This is argued by taking the contrapositive: if there exists an efficient algorithm A that produced an ε' -fidelity approximation $|D\rangle$ of $|C\rangle$ with probability at least δ' , then A *also* produced (with the same D) a ε -fidelity approximation with probability at least δ .

A conservative assumption would be to conjecture $(1 - 1/\text{poly}(n))$ -No-Learning (i.e., it is hard for a polynomial-time algorithm to produce a high-fidelity approximation of a state with high probability). On the other hand, given our current understanding of quantum state learning techniques, it also seems plausible to conjecture $2^{-\Omega(n)}$ -No-Learning (i.e., it is hard for a polynomial-time algorithm to produce even an *exponentially bad* approximation with *exponential small* probability). That said, we remark that for the purposes of implementing the cryptographic constructions in this paper, the parameters we require are fairly mild. For instance, we merely need the $(1 - 1/\text{poly}(n))$ -No-Learning assumption for the cryptographic constructions that assume a noise-free circuit, and even for our NISQ-friendly constructions, the ε -No-Learning assumption for any negligible⁴ function $\varepsilon(n)$ is sufficient.

1.1.1.2 Edge cases

The hardness assumption cannot hold for $\varepsilon \ll 2^{-n}$; this is because a quantum algorithm can output a uniformly random circuit D , and $|D\rangle$ will have fidelity at least 2^{-n} with $|C\rangle$ in expectation. Similarly, the hardness assumption cannot hold for $\delta \leq |\mathcal{C}_n|^{-1} = 2^{-\Omega(nd)}$, because the algorithm can always simply guess the description C of the state $|C\rangle$.

1.1.1.3 Proper vs. improper learning

We point out that the assumption is about the hardness of producing a circuit description D that *comes from the same class* \mathcal{C} as C . In the learning theory literature this is known as the setting of *proper* learning, where the learner has to output a hypothesis from the same concept class as the true underlying concept. One can also consider the setting of *improper* learning, where the output description D can come from a more general class of circuits. For example, the circuit learning algorithms of [36, 49] are improper learners: given copies of output states of depth d circuits, they produce descriptions of circuits with depth greater than d . One can consider hardness of learning conjectures in the improper setting as well; ours is the more *conservative* assumption (because the learner has the additional constraint of outputting a circuit from \mathcal{C}). As will be clear later, even if an improper learner exists for our class of circuits, our main cryptographic primitive designed on the hardness of learning – a random circuit based one-way state generator – will remain secure. Qualitatively, this is because the honest parties know the depth of the circuit. So, even if an adversary tries to cheat by using an improper learner, the honest parties can easily detect the mismatch in depth from the circuit description sent by the adversary.

1.1.2 The Computational No-Cloning assumption

In this section, we postulate a computational version of the famous No-Cloning principle of quantum mechanics. At a high level, it stipulates that given polynomially copies of the output state of a random quantum circuit, it is intractable to produce an additional copy of the state, or even an approximation of it.

⁴ A function is negligible if it goes to zero faster than any inverse polynomial.

► **Conjecture 2** (Computational No-Cloning). *Let \mathcal{C} denote a class of circuits. The $(\mathcal{C}, \varepsilon, \delta)$ -Computational No-Cloning Assumption stipulates the following. For all polynomial-time quantum algorithms A , for all polynomials $k(n)$, for all sufficiently large n , we have that*

$$\mathbb{P} \left[\left| \langle C |^{\otimes(k+1)} | \phi \rangle \right|^2 \geq \varepsilon : \begin{array}{l} C \leftarrow \mathcal{C}_n \\ |\phi\rangle \leftarrow A(|C\rangle^{\otimes k}) \end{array} \right] \leq \delta$$

Here, $C \leftarrow \mathcal{C}_n$ denotes that C is sampled uniformly from the subset of circuits in \mathcal{C} that act on n qubits, and a $(k+1)n$ qubit state $|\phi\rangle$ is sampled from the output of the algorithm A , given polynomially-many copies of $|C\rangle = C|0^n\rangle$.

As before, the algorithm A only has access to copies of the state generated by the random circuit and not the circuit description itself. We also abbreviate “Computational No-Cloning” as “No-Cloning” and also abbreviate “ $(\varepsilon, \varepsilon)$ -No-Cloning” as “ ε -No-Cloning”.

► **Corollary 3.** *The ε -No-Cloning Assumption (Conjecture 2) implies the ε^c -No-Learning Assumption (Conjecture 1) for some $0 < c < 1$.*

Proof Sketch. This follows from taking the contrapositive. If the description of the circuit C could be approximately learned with squared fidelity ε^c by some quantum polynomial-time algorithm A (i.e., quantum state learning is easy), we can devise an efficient cloning algorithm B that first runs A in a coherent fashion on the k copies to learn an approximation D , synthesizes an extra copy of $|D\rangle$ (which is supposed to represent the $(k+1)$ st copy of $|C\rangle$), and then uncomputes A to recover the original k copies of $|C\rangle$. The fact that this strategy obtains ε fidelity with $|C\rangle^{\otimes(k+1)}$ follows from a calculation virtually identical to that of Claim 4.10 of [28], which is part of the proof that the No-Cloning Assumption implies the existence of quantum commitments. This violates the ε -No-Cloning Assumption. ◀

Cloning is potentially easier than learning

Corollary 3 means that learning is potentially an *easier* task than cloning. It could be that we can clone just one more copy of the state but we still cannot learn the state.

There is some suggestive evidence of this fact from the work of Nehoran and Zhandry [59], who construct a quantum oracle with respect to which a collection of states is efficiently clonable, but it is not efficiently “telegraphable,” given only one quantum sample. Informally, telegraphing means getting a classical string out of one copy of a quantum state by “deconstructing” it, from which one copy of the state can again be “reconstructed.” While the notion is qualitatively reminiscent of a learning task, this is formally different from the notion of learning in our paper.

1.2 Barriers and evidence for the hardness assumptions

1.2.1 The barriers

What are the prospects of proving Conjecture 1 or Conjecture 2 outright? First, we note that the conjectures are necessarily *computational*, meaning that they are intrinsically about the limits of efficient quantum computation. This is because it *is* possible for an exponential-time quantum algorithm to learn a circuit description with high probability, given only polynomially-many quantum samples (and thus also solve the cloning task with similar sample complexity). One method is to use the classical shadows protocol of [34]; we describe this in more detail in [28]. We note that Morimae and Yamakawa already observed that

there is always an exponential-time attack on one-way state generators [55] via shadow tomography; this is essentially the same observation.

Furthermore, we note that proving our hardness conjectures would have dramatic implications for *classical* complexity theory. The classical shadows-based algorithm described in [28] can be implemented in polynomial time assuming the complexity inclusion $\text{NP}^{\#P} \subseteq \text{BQP}$.⁵ Therefore Conjecture 1 implies $P \neq \text{PSPACE}$. While this doesn't imply (say) $P \neq \text{NP}$, for all intents and purposes this would represent a similar breakthrough in mathematics and complexity theory. The longstanding difficulty of proving such complexity separations pose a barrier to proving our hardness assumptions.

1.2.2 The evidence

We now discuss the evidence for our hardness assumptions. First we discuss the No-Learning assumption.

1.2.2.1 Examining best known learning algorithms

Recently there has been progress on obtaining efficient algorithms for learning states of bounded circuit complexity [35, 29, 49]. The algorithm of Fefferman, Ghosh, and Zhan [29] requires the use of oracle access to the circuit itself, rather than only having copies of the output state $C|0^n\rangle$, which we do not consider in this paper. For algorithms that learn using copies of the output state alone, the current state-of-the-art is due to Landau and Liu [49], who show the following:

► **Theorem 4** ([49]). *Fix an integer $k > 0$. There is an algorithm that, for all $\varepsilon > 0, \delta > 0$, given copies of an unknown quantum state $|C\rangle = C|0^n\rangle$ for some depth- d circuit acting on a k -dimensional lattice with two-qubit gates, outputs the description of a depth $(2k + 1)d$ circuit D such that $\| |C\rangle - |D\rangle \| \leq \varepsilon$ with probability $1 - \delta$. Furthermore, the algorithm uses*

$$\frac{n^4 \cdot 2^{O(d^k)}}{\varepsilon^4} \log \frac{n}{\delta}$$

copies of $|C\rangle$, and runs in time

$$\frac{n^4 \cdot 2^{O(d^k)}}{\varepsilon^4} \log \frac{n}{\delta} + \left(\frac{nd}{\varepsilon}\right)^{O(d^{k+1})}.$$

Here, the $O(\cdot)$ suppresses dependence on k , which we treat as a constant.

For intuition, consider some parameter settings. Suppose ε, δ are some fixed constants (like 0.001).

1. When $d = O(1)$, then the sample complexity is $O(n^4 \log n)$, and the time complexity is $n^{O(d^k)}$.
2. When $d = O(\log^{1/k} n)$, then the sample complexity is still $\text{poly}(n)$, but the time complexity becomes quasipolynomial $2^{\text{poly} \log n}$.
3. When $d \gg \log n$, then both the sample and time complexity become superpolynomial.

⁵ It was recently shown by Hiroka and Hsieh [32] that efficient state learning is possible if $\text{PP} \subseteq \text{BQP}$, which represents an improved complexity upper bound.

In more detail, the learning algorithms of both Huang, et al. [35] and Landau and Liu [49] revolve around the idea of learning so-called “local inversions” of the circuit C , which are small subcircuits V that “undo” part of the overall circuit to be learned: $(V^\dagger \otimes \mathbb{I})|C\rangle \approx |0^k\rangle \otimes |\theta\rangle$ for some $n - k$ qubit state $|\theta\rangle$. In other words, some small number k of qubits have been disentangled from the state $|\psi\rangle$.

If the circuit depth d is small, then the local inversions have size $O(d^k)$ (assuming a k -dimensional circuit architecture), and the learning algorithm can brute force over all possible subcircuits of size $O(d^k)$. Once the local inversions have been learned, the challenge is to “stitch” all of the local inversions (which overlap with each other) in a consistent way.

The complexity of learning a single local inversion takes time at least $2^{O(d^k)}$. Thus with a randomly chosen circuit of depth that asymptotically grows faster than $O(\log n)$ (e.g., $d = \log^2 n$), there are many more possibilities than is possible to consider in polynomial time. Furthermore, searching through candidate local inversions of a random circuit appears to be an “all or nothing” task: either a candidate successfully inverts a local patch, or it will likely scramble the state even further. Thus it does not seem possible to make the learning algorithms of [35, 49] “gracefully fail” by dialing back their runtime to being polynomial in n . Looking at the prototypical algorithmic ideas in these papers suggests that if one is in fact limited to polynomial time algorithms, then the typical fidelity will likely be exponentially small.

There are other papers, like [43], that take as input copies of the local density matrix on $\mathcal{O}(1)$ -sized patches as input and output the description of the global unitary. [43] uses an information-theoretic criteria and learns local quantum channels that can extend the state. However, this approach is also similar in spirit to the local inversion technique and the time complexity has a factor that scales as $2^{O(d^k)}$, which means that, similar to [35], these algorithms are inefficient for any depth that grows asymptotically larger than $\mathcal{O}(\log n)$.

1.2.2.2 Worst-case hardness from post-quantum assumptions

The No-Learning Assumption is an *average-case* hardness assumption, because it is about learning over the uniform distribution over circuits. One can also wonder about the hardness of *worst-case* learning. Zhao, et al. [70] showed that, assuming that subexponential-time quantum computers cannot solve the RingLWE problem, any quantum algorithm that learns the circuit description given copies of the n -qubit state requires at least $\exp(\Omega(\min\{G, n\}))$ time, where G is size of the quantum circuit to be learned. RingLWE is a version of the Learning With Errors (LWE) problem, whose assumed hardness underlies most proposals for post-quantum cryptography (i.e., cryptography that can be run on classical computers, but are secure against quantum computers). When the number of gates G is superlogarithmic, then the time complexity lower bound is superpolynomial.

One may wonder why, given the results of Zhao, et al. [70], we need to make a separate assumption about the hardness of learning quantum circuits. Given the widespread belief in the security of various versions of LWE (for example this underlies much of NIST’s recent standardization of recommended post-quantum cryptosystems [30, 58]), it would seem that hardness of (Ring)LWE automatically implies the hardness of learning circuits from [70] and therefore all our applications of Conjecture 1 follow.

There are two issues with this reasoning. The first is that the result of Zhao, et al. [70] implies the hardness of learning under a very particular distribution of quantum circuits: namely, ones that encode the RingLWE problem, which are quite structured⁶ and are statistically far from truly random quantum circuits. The second and most important issue

⁶ In more detail, these are based on constructions of *pseudorandom states* from one-way functions [39].

is that, while the RingLWE hardness can be viewed as evidence for our hardness conjectures, it appears to be a significantly harder statement to prove. For one, the hardness of RingLWE trivially implies $P \neq NP$, one of the major open questions in mathematics. On the other hand, $P \neq NP$ is *not* known to be implied by Conjecture 1. Indeed, there is emerging evidence that Conjecture 1 is a reduced assumption compared to $P \neq NP$ (i.e., there are mathematical worlds in which $P = NP$ but tasks related to learning quantum circuits is still hard [46]). Since one of our primary motivations is to base cryptography on the fewest mathematical assumptions possible, we treat our hardness assumptions as being more basic and plausible than any post-quantum hardness assumption.

1.2.2.3 Evidence for No-Cloning assumption

We can similarly examine the evidence for the No-Cloning assumption. As mentioned above, since the cloning is potentially an easier task, the No-Cloning conjecture is a *stronger* assumption than No-Learning. However, as far as we are aware, there are no known algorithms for cloning that perform significantly better than ones for learning. One could consider, for example, the optimal pure-state cloning map that was analyzed by Werner [69]. This map, which is essentially to project the input state onto the $(k + 1)$ -fold symmetric subspace, is provably the most sample-efficient procedure for taking copies $|\psi\rangle^{\otimes k}$ of an *arbitrary* input state $|\psi\rangle$ (not necessarily one generated by a polynomial-size circuit) and producing an approximation of $|\psi\rangle^{\otimes k+1}$. Werner [69] showed the best achievable fidelity in general is

$$\frac{\binom{2^n+k-1}{k}}{\binom{2^n+k}{k+1}} = \frac{k+1}{2^n}$$

which is exponentially small for $k = \text{poly}(n)$.

Aside from Werner’s optimal cloning map, as far as we are aware the best algorithms for the cloning task are based on first solving the learning task. As we discussed before, such algorithmic ideas, if we restrict them to run in polynomial time, seem unable to achieve anything better than exponentially small fidelity. We leave it is an interesting open problem to come up with an algorithms – even ones that run in subexponential time – for cloning states of random circuits that do not learn the circuit description first.

1.2.2.4 Black-box lower bounds

We give more evidence to support the No-Learning and No-Cloning assumptions by proving lower bounds in the black-box model. We desire a black-box model which captures the analogous properties in the white-box setting. For instance, one property we would like to capture is the typical distribution of the output state of a random quantum circuit, which mimics the distribution of a Haar-random state. Another property we would like to capture in the black-box model is the existence of a learning algorithm that uses only polynomially-many copies of the state but is allowed to make exponentially many black-box queries – this would be analogous to the existence of the exponential-time, polynomial-sample complexity learning algorithm.

With such considerations in mind, we formalize a *state preparation oracle* model where the oracle \mathcal{O} takes as input a state $|i\rangle|0^n\rangle$ and outputs $|i\rangle|\psi_i\rangle$ for some Haar-random state $|\psi_i\rangle$. The oracle can be accessed in superposition, and there is no guarantee about what the oracle does when the second register is initialized to something other than all zeroes. Intuitively, each index i corresponds to a different polynomial-size circuit, but the algorithm is not allowed to exploit the structure of the circuit except to prepare the resulting output state. We prove the following:

► **Protocol 6. Random Circuit OWSG**

Generation algorithm: Given input key $C \in \{0, 1\}^{r(n)}$, interpret it as a description of an n -qubit circuit C from the ensemble \mathcal{C}_n . Output $|C\rangle = C|0^n\rangle$.

Verification procedure: Given input $C \in \{0, 1\}^{r(n)}$ and a state $|D\rangle$ on n qubits, apply C^\dagger to the state, and measure. Accept if the result is all zeroes, and reject otherwise.

■ **Figure 1** Construction of one-way state generator from random circuits.

► **Theorem 5** (Black-box lower bounds for cloning). *There exists a state preparation oracle \mathcal{O} such that all T -query quantum query algorithms getting k copies of $|\psi_J\rangle$ for a uniformly random index $J \in [2^n]$ satisfy*

$$F(\rho, |\psi_J\rangle\langle\psi_J|^{\otimes k+1}) \leq 2^{-n/4}(2T + k + 1)$$

where ρ is the output of the query algorithm and $F(\cdot)$ denotes the fidelity function.

In other words, unless either the number of queries T or the number of copies k are $2^{\Omega(n)}$, the expected fidelity of cloning is exponentially small. Since the ability to learn implies the ability to clone, this also shows a similar lower bound for the learning task. We elaborate on the model and prove Theorem 5 in [28].

1.3 Cryptography from our hardness assumptions

We now summarize our main cryptographic applications using the hardness of quantum learning assumptions from Section 1.1.

1.3.1 One-way state generators from hardness of learning

A *one way state generator* (OWSG) is an efficient algorithm that takes as input a classical key k , and outputs a quantum state $|\psi_k\rangle$ from that key. Anyone with the key k can efficiently verify that $|\psi_k\rangle$ is the correct output. Furthermore, given polynomially many copies of the $|\psi_k\rangle$ for a randomly chosen k , it should be computationally hard for an adversary to produce another key k' such that the corresponding state $|\psi_{k'}\rangle$ is close to the original output $|\psi_k\rangle$. (For a formal definition, see [28]). Introduced by Morimae and Yamakawa [56], OWSGs are a quantum analogue of one-way functions, which are functions efficiently computable in the forwards direction but computationally difficult to invert.

The No-Learning Assumption (Conjecture 1) is essentially *equivalent* to the existence of a OWSG, namely the Random Circuit OWSG described below in Figure 1. For a range of parameters, this is immediate: for negligible ε (i.e., ε goes to 0 faster than any inverse polynomial), the ε -No-Learning Assumption is easily seen to be equivalent to the security of the Random Circuit OWSG. When ε is larger, say even up to $1 - 1/\text{poly}(n)$, the equivalence still holds; this relies on hardness amplification techniques for OWSGs [55, 17]. We note that this equivalence was also independently observed by Hiroka and Hsieh in a recent preprint [32].

Although a OWSG is not immediately cryptographically useful by itself, it is now known that OWSGs can be used as a primitive building block for a variety of quantum cryptosystems. In their papers defining OWSGs [56, 55], Morimae and Yamakawa showed

that OWSGs can be used to build bounded-time-secure digital signature schemes. In a breakthrough work, Khurana and Tomer showed that OWSGs imply the existence of quantum bit commitments [41], which in turn imply other functionalities such as quantum zero knowledge proofs for NP and secure multiparty computation [19]. Therefore, using the Random Circuit OWSG, we obtain concrete implementations of these cryptosystems on quantum computers, without relying on the use of one-way functions.

An attractive feature of our Random Circuit OWSG is that it seems potentially amenable to implementation on noisy quantum computers, and thus the corresponding cryptosystems may be realizable in the near- and medium-term. We discuss noise tolerant versions of our random circuit-based cryptographic protocols in Section 1.4.

1.3.2 Simple quantum commitments from the hardness of cloning

A commitment scheme enables two parties (known as a “committer” and a “receiver”) to perform the cryptographic equivalent of putting a message in a sealed envelope that is opened later. In a quantum commitment scheme, a committer upon getting a bit b generates a bipartite pure state $|\psi_b\rangle_{AB}$, and sends the register B to the receiver; this is the “commitment phase” of the protocol and is the analogue of sending the sealed envelope. At this point the receiver should not be able to tell what the bit b is.

Later, in the “reveal phase” of the protocol, the committer announces the bit b and sends the remaining register A of $|\psi_b\rangle$ to the receiver, who can uncompute the state to check its validity. The security of the commitment scheme ensures that the committer cannot “change his mind” in between the commit and reveal phases to convince the receiver he had committed to the opposite bit $1 - b$.

Recently, quantum commitments have become a centerpiece of the zoo of quantum cryptographic primitives [8, 56, 19, 41]. As mentioned, Khurana and Tomer showed that OWSGs can be used to construct quantum bit commitments in a generic way [41]. Therefore our No-Learning Assumption implies the existence of secure quantum commitments. However the construction of commitments in [41] is quite involved, requiring an intricate sequence of transformations mimicking the classical transformation from one-way functions to pseudorandom generators [31].

We construct a simple quantum bit commitments based on the Computational No-Cloning Assumption (Conjecture 2), and furthermore the analysis is fairly direct and straightforward. We describe the construction below. As discussed in Section 1.1, the No-Cloning Assumption implies the No-Learning Assumption. The ease of obtaining a commitment scheme from the No-Cloning Assumption suggests that the No-Cloning Assumption could be strictly stronger than the No-Learning Assumption (because it appears that obtaining commitments from OWSGs requires an intricate analysis [41]).

We present and analyze the bit commitment scheme in detail in [28].

1.4 NISQ-friendly quantum cryptography

An important challenge in the field of quantum computing is to find a practical use case for noisy, near-term quantum (i.e., NISQ) computers. Although great strides have been made recently in demonstrating principles of error-correction and fault-tolerance on quantum devices [14, 4], large-scale implementations of many quantum algorithms of interest (e.g., Shor’s, Grover’s, Hamiltonian simulation, etc) still seem quite a ways off. Thus there is significant interest in finding an application that (a) can be implemented on an a NISQ

► **Protocol 7. Commitment scheme based on random circuits.**

Commitment phase. To commit to bit $b = 0$, the committer prepares the state

$$|\psi_0\rangle_{AB} := \frac{1}{\sqrt{|\mathcal{C}_n|}} \sum_{C \in \mathcal{C}_n} (|C\rangle_A^{\otimes k} \otimes |0^n\rangle_A) \otimes |\hat{C}\rangle_B$$

where $|\hat{C}\rangle$ represents the *classical description* of the circuit C .

To commit to bit $b = 1$, then prepare the state

$$|\psi_1\rangle_{AB} := \frac{1}{\sqrt{|\mathcal{C}_n|}} \sum_{C \in \mathcal{C}_n} |C\rangle_A^{\otimes(k+1)} \otimes |\hat{C}\rangle_B .$$

The committer sends register B of the state $|\psi_b\rangle$ to the receiver.

Reveal phase. The committer reveals the bit b to the receiver, and also sends the remaining register A of $|\psi_b\rangle$. To verify, the receiver will uncompute the unitary that synthesizes $|\psi_b\rangle$ and check that the all zeroes state is obtained.

■ **Figure 2** Commitment scheme based on random circuits.

device, (b) has some advantage over classical computers/protocols, and (c) is practically useful. We show that our hardness assumptions yield quantum cryptosystems that satisfy these three criteria.

Our random circuits-based cryptosystems are arguably NISQ-friendly. Random quantum circuits have been investigated extensively on real hardware platforms ever since Google’s original quantum supremacy announcement in 2019 [10]. The lack of structure in random quantum circuits is advantageous for maximizing quantum advantage while minimizing the burden on the NISQ device. Furthermore, the effective noise model when executing random quantum circuits often becomes quite simple [10, 24].

Furthermore, implementing cryptosystems based on our hardness assumptions would concretely realize the possibility of having secure quantum cryptography using reduced assumptions as compared to classical cryptography (for example, we do not need to assume $P \neq NP$ or that one-way functions exist) [45, 8, 56, 19, 42]. This would represent what we call “quantum cryptographic advantage.”

We presented protocols that solve useful cryptographic tasks: digital signatures, bit commitments, encryption, and more. Although these protocols are not immediately NISQ-friendly out of the box, we show how to “NISQ-ify” some of them so that they are.

1.4.1 NISQ-friendly one-way state generators

While implementing the Random Circuit OWSG on a realistic quantum computer, noise can degrade the fidelity of the output state. For certain choices of noise parameters and depth regimes, the fidelity can be inverse polynomial in the number of qubits. If that happens, the success probability of any verification procedure would degrade similarly. Although the OWSG may be secure against any polynomial-time adversary, it may not be very useful if the “honest user” (e.g., someone with the key) tries to run it on a noisy quantum computer. On the other hand, making a OWSG tolerant to noise may give greater leeway to break its security.

However, if we make a sufficiently strong assumption (e.g., the ε -No Learning assumption for negligible ε), there is a noticeable gap between the success probability of a *noisy* verifier (who has the circuit description) and any *noiseless* polynomial-time adversary (who doesn't have the circuit description): the adversary cannot produce any approximation to the state except with negligible fidelity. We can amplify this gap to obtain a NISQ-friendly OWSG, where both the generation and verification algorithms can be run on a noisy quantum computer, but the OWSG also retains its security against noiseless adversaries. We achieve this amplification via a *computational Chernoff bound*.

Computational Chernoff bounds

A standard way to amplify the security of a OWSG is via *parallel repetition* [55, 17]: the key to the amplified OWSG corresponds to a t -tuple of independently chosen random circuits (C_1, \dots, C_t) , and the output is the tensor product of the corresponding states $|C_1\rangle \otimes \dots \otimes |C_t\rangle$. Verification proceeds by checking that the i 'th block of n qubits is in the state $|C_i\rangle$ for each $i = 1, \dots, t$. Intuitively, if it is somewhat hard for the adversary to learn the output of one random quantum circuit, then it should be *very* hard for the adversary to simultaneously learn the output of *many* random quantum circuits. However, because the noisy fidelity is small, standard parallel repetition theorems do not directly work because the fidelity of the overall state on nt qubits would degrade exponentially with t , making it $\text{negl}(n)$, which means even honest verifiers would fail to see a non-trivial signal.

Our technical innovation is to use *threshold parallel repetition* for amplification – this is a OWSG $G^{t,k}$ consisting of t independent copies of G , but instead of verifying that all t copies have been inverted, the verification algorithm checks that at least k out of the t have been inverted. To bound its security we prove new *computational Chernoff bounds* for OWSGs. A detailed discussion, with the theorem statements, can be found in [28].

1.4.2 NISQ-friendly digital signatures

As a direct application of our NISQ-friendly OWSG we obtain a NISQ-friendly quantum digital signature scheme. At a high level, a digital signature scheme (with quantum public keys) is a method for a *signer* to generate a *signature* for a message in a way that a third party *verifier* (using a quantum public key posted by the user beforehand) can verify that the signature belongs to the message (and in particular, the message or the signature have not been changed).

Morimae and Yamakawa [56] showed that such quantum digital signature schemes can be directly constructed from OWSGs by adapting the famous Lamport construction of digital signatures [48]. We show that by plugging in the NISQ-friendly random circuit OWSG, their digital signature scheme becomes NISQ-friendly as well. This gives example of an end-to-end cryptographic task for NISQ-devices whose hardness relies on an innately quantum conjecture.

We present the scheme and the analysis in detail in [28].

1.4.3 On noise assumptions and asymptotics

Note that the only noise assumption we need is that the fidelity of the signal should at most be inverse polynomially large. However, an observant reader would notice that \mathcal{C}_n is an ensemble of sufficiently deep circuits. For certain noise models, for e.g. constant rate of depolarizing noise per gate, or more generally, constant rate of unital noise per gate, the output converges to the maximally mixed state [6, 26] exponentially fast in the depth of

the circuit, which would cause inverse superpolynomially large signal decay at large depths. However, there are three perspectives on why our proposal is still relevant to near-term experiments.

Firstly, the results proving convergence to the maximally mixed state [26, 7] are asymptotic statements, whereas real experiments have finite system sizes. Thus, there is a discrepancy between what theoretically happens when we scale up the system size and what is experimentally observed for a fixed system size. For example, in quantum advantage demonstrations with random circuits, the experimentalists have observed signatures of long range entanglement and evidence of convergence to the Porter-Thomas distribution when the output state is measured in the standard basis [10, 57]. Note that Porter-Thomas distribution is far in total variation distance from the uniform distribution, where the latter is what we get when we measure a maximally mixed state in the standard basis. If the system were indeed close to being maximally mixed, we would neither see long-range entanglement nor Porter-Thomas type behavior. Thus, sampling from the uniform distribution is not a good approximation to the realistic output distribution, even though the distributions are close asymptotically [6, 26].

There are other classical samplers, such as the recent one proposed in [7], which differ from simply sampling the uniform distribution. However, the sampler in [7] achieves a smaller total variation distance than the uniform distribution only at a depth of approximately $\sim \log n$. Properties of real experiments – such as the presence of long-range entanglement – indicate that they do not operate at logarithmic depth but rather in a much deeper regime. Thus, in the same way as the uniform distribution, it is unclear if the sampler of [7] is directly applicable to real experiments, even though its classical spoofing distribution is again asymptotically close to the noisy target distribution.

Secondly, note that the fidelity of the output state of a noisy circuit, comprised of two qubit gates and a single-qubit, uncorrelated noise channel acting upon each qubit after the application of each gate, is

$$F = (1 - \epsilon)^{2s},$$

or the probability that no errors occurred anywhere in the system. Here, where s is the circuit size and ϵ is the noise rate per qubit. If ϵ is at most $\sim \frac{1}{n}$, $F \approx e^{-\Theta(s \cdot \epsilon)}$. Hence, one valid regime for inverse polynomial decay in fidelity is when ϵ is $\sim \frac{1}{n \log n}$ and the system size s is $\sim n \log^2 n$. In certain models, the structure of the output state becomes even simpler. One of the a noise models for which this is manifestly true is the white noise model [10, 14, 24]. According to this model, if the noise per gate is unital, and if the noise rate ϵ is at most $\sim \frac{1}{n}$, then the output state of the circuit can be written as

$$\rho_{\text{out}} = F \rho_{\text{noiseless}} + (1 - F) \frac{\mathbb{I}}{2^n},$$

that is, as a linear combination of $\rho_{\text{noiseless}}$, which is what the output state would have been if there were no noise, and the maximally mixed state. While the noise rate per gate going down with n is unrealistic for extremely large system sizes, it is nonetheless a reasonable model of real experiments as structural properties of experimental output states match the white noise output state. In fact, judging by the recent progress in random quantum circuit experiments, e.g. [10, 57, 66], it seems quite reasonable to model realistic noise as going down with system size.

Thirdly, note that in all of the previous discussions, we have assumed noise to be unital. But real noise is complicated and it is unclear if any of the above results (e.g., convergence to the maximally mixed state or the white noise model) are realistic for near-term experiments. As an example of surprising behavior with more general noise models, researchers have recently

studied the effects of non-unital noise channels in random quantum circuits [27, 53, 61]. Non-unital noise is ubiquitous in real world experiments, as witnessed by, e.g., readout errors, T_1 decay for superconducting systems, and photon loss in bosonic systems [67, 68, 72]. In particular, certain structural properties that are true for unital noise at some regimes, like anti-concentration or convergence to the maximally mixed state, fail in the presence of any constant rate non-unital noise channel [27], and hardness or easiness results that assume these properties, like [1, 5, 18, 7], do not seem to work.

1.5 Related work

We discuss the relationship between our work and some concurrent, independent works that recently appeared.

Cryptography from assumptions about random circuits

Khurana and Tomer [42] also study the quantum cryptographic implications of hardness assumptions about random circuits. Their hardness assumption posits that it is $\#P$ -hard to estimate the output probabilities of a random quantum circuit, given a classical description of the circuit. This is a well-studied hardness assumption and is the theoretical basis for many quantum supremacy proposals (e.g., [1, 18, 10]). Combined with the complexity-theoretic assumption that $P^{\#P} \not\subseteq \text{ioBQP}/\text{qpoly}$, [42] show how to construct *quantum one-way puzzles*, which in turn implies the existence of quantum bit commitments through their previous work [41].

While random circuits are at the core of the hardness assumptions of our paper as well as [42], the similarities quickly end. Khurana and Tomer are positing the hardness of a *classical-input, classical-output* task: given the description of a quantum circuit and a string, estimate the probability of outputting the string. Our hardness assumptions, on the other hand, are about *quantum-input* tasks: given *quantum* copies of the output state of a random circuit, either learn or clone it. Importantly, the circuit description is *not* known to the adversary.

Furthermore, our motivations have some differences: they are motivated by basing quantum cryptography on separations between *decision* complexity classes, whereas we are primarily motivated by the connection between quantum learning problems (which involve quantum inputs) and quantum cryptography.

On a different note, in Bostanci, Haferkamp, Hangleiter, and Poremba [16], the authors construct quantum cryptography from a suite of assumptions about random IQP circuits. Depending on the type of assumption, the authors get quantum trapdoor functions, quantum pseudoentanglement, and candidate constructions of efficient pseudorandom unitaries. There are large differences between this work and our work, in terms of the nature of assumptions, the justifications for hardness, the flavours of cryptography that one gets, and the query lower bounds. These two works represent complementary explorations into cryptography from two different random ensembles, ours involving brickwork random circuits, and theirs involving IQP circuits.

Cryptography from assumptions about quantum states, protocols, and noise

Qian, Raizes, and Zhandry [64] study the quantum cryptographic implications of a new “search-type” assumption they call *classical \rightarrow quantum extrapolation*, where the goal is to extrapolate the rest of a bipartite pure state given the first register measured in the computational basis. They show that the hardness of this extrapolation task implies the existence of quantum bit commitments and is implied by the existence of various quantum

public-key primitives. We view their work as studying the cryptographic implications of a (conceptually new) “generic” assumption, where they do not specify how exactly to generate the hard, inextrapolable bipartite states. On the other hand, we are focused on the cryptographic implications of a “concrete” assumption, where we instantiate the underlying primitive (OWSG, quantum commitment) with a concrete algorithmic implementation. This is similar to the difference between assuming that *some* one-way function exists, versus assuming that a *specific* one-way function exists (e.g., the RSA function or the LWE function).

Morimae, Shirakawa, and Yamakawa [54] give a characterization of the complexity assumptions needed for a class of protocols for proofs of quantumness; in particular they show that one-way puzzles (the same primitive constructed by [42]) are necessary and sufficient. Their goal is squarely aimed at understanding the complexity of proofs of quantumness in the abstract, and less by having concrete instantiations of quantum cryptographic primitives. For related papers on the interplay between one-way puzzles, proofs of quantumness, and quantum cryptography, also see [22, 33, 21].

Hiroka and Hsieh [32] studies the hardness of learning efficiently generatable pure states. The main focus of this paper is a PP upper bound on this task. They also base some cryptography on their hardness assumptions, like the existence of one-way state generators. The crucial difference between this work and ours is that they consider one particular learning assumption, whereas we consider a suite of learning and distinguishing assumptions, basing different flavours of cryptography on each. We also discuss in detail how to make our protocols NISQ-implementable.

Poremba, Quek, and Shor [63] put forward a new quantum-inspired primitive called Learning Stabilizers with Noise (LSN), which deals with decoding a random stabilizer code in the presence of local depolarizing noise. Their primitive implies (statistically hiding and computationally binding) bit commitments. Their goal is to construct a new natively quantum assumption for quantum cryptography, as opposed to NISQ-friendliness of their protocols.

NISQ-friendly cryptography

Finally, we comment on the relationship between the proposals by [3, 11] to use NISQ devices to perform certifiable randomness generation. Similarly to our work, they propose a cryptographic task that can be performed on a NISQ device, and whose security can be based on hardness assumptions about random circuits.

However a significant difference is that the verification procedure in their protocols are inherently inefficient; it requires exponential time even using a quantum computer as it requires approximating output probabilities of a random quantum circuit. In contrast, our NISQ-friendly digital signature scheme is efficiently implementable.

1.6 Summary

Our exploration also uncovers a deeper understanding of various quantum learning tasks and cryptographic primitives.

Fine-grained distinctions in learning and cryptography

Our work connects fine-grained learning tasks to understanding the relative power of different cryptographic primitives. This opens up a potential new approach to understand both. For instance, cloning is potentially an easier task than learning⁷. Our work shows that the

⁷ There is some suggestive evidence here in the form of a black-box separation for a related task [59].

hardness of learning assumptions is essentially equivalent to the existence of OWSGs from random circuits, while the hardness of cloning can be used to construct a fairly simple bit commitment scheme. This indicates that commitments are likely to be a more minimalistic cryptographic primitive than OWSGs. Our conclusion is consistent with oracular evidence in [12, 15].

Practical applications of random quantum circuits

Random circuits have been extensively studied in the context of quantum advantage in the NISQ-era. Several experimental groups around the world (for e.g., [10, 57, 73, 71, 72]) have claimed practical demonstrations of quantum advantage with sampling tasks involving random circuits. Furthermore, there also has been extensive effort to build theoretical foundations of quantum advantage based on such sampling tasks (see e.g., [2, 18, 7, 27]), but even if we are able to demonstrate practical advantage with such circuits, one major challenge that remains is to use NISQ devices or random quantum circuits to solve practically useful problems.

By proposing useful cryptographic applications of random circuits that are “NISQ-friendly”, our work takes one further step in addressing this challenge and complements the recent work on certified random number generation with random circuits ([3, 11]).

Minimal assumptions for cryptography

Our work contributes to understanding the minimal theoretical assumptions needed for quantum cryptography. While in the classical world the existence of one-way functions is widely believed to be necessary for cryptography [38], in the quantum context this may not be the case. In particular, recent work has given black-box evidence in which $P = NP$ and yet single-copy secure pseudorandom quantum states still exist [46, 45]. This suggests that certain quantum cryptographic primitives are possible even without the existence of one-way functions, for e.g., see [8, 19, 51, 47]. On the other hand, in the white-box setting, all currently known constructions of such quantum pseudorandom states require the existence of quantum-secure one-way functions [39]. Consequently, a major challenge that remains is to construct quantum cryptographic primitives which do not rely on the existence of one-way functions and are based on concrete hardness assumptions. Our work addresses this by constructing cryptography based on the hardness of quantum learning.

1.7 Future directions

The connections between the hardness of quantum learning and cryptography leads to several interesting directions that require more exploration:

- **Relations between learning and cloning.** We posed two concrete assumptions about the hardness of quantum learning for random circuits. The natural open question that remains is to understand the differences between these different learning tasks (No-Learning versus No-cloning), their relative hardness, and to understand the security parameters needed for the hardness assumptions.
- **Quantum cryptography from concrete hardness assumptions.** Innately quantum hardness of learning assumptions, like the No-Learning and No-Cloning assumptions, give a natural direction to give concrete instantiations of other cryptographic primitives on assumptions that might be weaker than one-way functions.

- **NISQ-friendly cryptography and practical applications.** While several cryptographic constructions presented in this paper are NISQ-friendly, others, such as for quantum bit commitments, involve operations that are not realistic for near-term devices, for instance, coherently implementing a superposition over all quantum circuits. This leads to the tantalizing possibility of finding other NISQ-friendly cryptographic primitives.
- **Quantum pseudorandomness and connections to complexity.** There are fundamental open questions regarding pseudorandomness properties of states produced by random quantum circuits. Using hardness of learning to probe such questions is an interesting open direction. Along this line of inquiry, one might further expect to find deeper connections between learning, cryptography, pseudorandomness, state complexity classes and circuit lower bounds.
- **Realization using classical communication and quantum nodes.** There are open questions regarding whether our cryptographic protocols, like the digital signature scheme and the bit-commitment scheme, are realizable by quantum end-nodes and classical communication, obviating the need for a quantum network to transport the states. This will make the protocols even more NISQ-friendly, as noise-robust quantum networks are hard to build in practice.

References

- 1 Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011. doi:10.1145/1993636.1993682.
- 2 Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference (CCC 2017)*, pages 22:1–22:67. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.CCC.2017.22.
- 3 Scott Aaronson and Shih-Han Hung. Certified randomness from quantum supremacy. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 933–944. ACM, 2023. doi:10.1145/3564246.3585145.
- 4 Rajeev Acharya, Laleh Aghababaie-Beni, Igor Aleiner, Trond I Andersen, Markus Ansmann, Frank Arute, Kunal Arya, Abraham Asfaw, Nikita Astrakhantsev, Juan Atalaya, et al. Quantum error correction below the surface code threshold. *arXiv preprint*, 2024. arXiv:2408.13687.
- 5 Dorit Aharonov. A simple proof that Toffoli and Hadamard are quantum universal. *arXiv preprint*, 2003. arXiv:quant-ph/0301040.
- 6 Dorit Aharonov and Michael Ben-Or. Polynomial simulations of decohered quantum computers. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 46–55. IEEE Comput. Soc. Press, 1996. doi:10.1109/sfcs.1996.548463.
- 7 Dorit Aharonov, Xun Gao, Zeph Landau, Yunchao Liu, and Umesh Vazirani. A polynomial-time classical algorithm for noisy random circuit sampling. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1307–1318. ACM, 2023. doi:10.1145/3564246.3585161.
- 8 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Advances in Cryptology – CRYPTO 2022*, volume 13508 of *Lecture Notes in Computer Science*, pages 208–236. Springer, 2022. doi:10.1007/978-3-031-15979-4_8.
- 9 Srinivasan Arunachalam, Alex B. Grilo, and Aarthi Sundaram. Quantum hardness of learning shallow classical circuits. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020)*, pages 291–304, New York, NY, USA, 2020. ACM. doi:10.1145/3357713.3384300.

- 10 Frank Arute, Kunal Arya, Ryan Babbush, and et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510, 2019. doi:10.1038/s41586-019-1666-5.
- 11 Roozbeh Bassirian, Adam Bouland, Bill Fefferman, Sam Gunn, and Avishay Tal. On certified randomness from fourier sampling or random circuit sampling, 2024. arXiv:2111.14846.
- 12 Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. A new world in the depths of microcrypt: Separating OWSGs and quantum money from QEFID. Cryptology ePrint Archive, Paper 2024/1567, 2024. URL: <https://eprint.iacr.org/2024/1567>.
- 13 Avrim Blum, Merrick Furst, Michael Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO ’93: 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22–26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 278–291. Springer Berlin Heidelberg, 1993. doi:10.1007/3-540-48329-2_24.
- 14 Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, et al. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626(7997):58–65, 2024.
- 15 John Bostanci, Boyang Chen, and Barak Nehoran. Oracle separation between quantum commitments and quantum one-wayness. *arXiv preprint*, 2024. arXiv:2410.03358.
- 16 John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. Efficient quantum pseudorandomness from hamiltonian phase states, 2024. doi:10.48550/arXiv.2410.08073.
- 17 John Bostanci, Luowen Qian, Nicholas Spooner, and Henry Yuen. An efficient quantum parallel repetition theorem and applications. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*. ACM, 2024. doi:10.1145/TBD.
- 18 Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159–163, 2019. doi:10.1038/s41567-018-0318-2.
- 19 Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, pages 24:1–24:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ITCS.2023.24.
- 20 Marco Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Learning algorithms from natural proofs. In *Proceedings of the 31st Computational Complexity Conference (CCC 2016)*, volume 50 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:24, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2016.10.
- 21 Bruno P. Cavalari, Eli Goldin, Matthew Gray, and Peter Hall. A meta-complexity characterization of quantum cryptography, 2024. doi:10.48550/arXiv.2410.04984.
- 22 Nai-Hui Chia, Honghao Fu, Fang Song, and Penghui Yao. A Cryptographic Perspective on the Verifiability of Quantum Advantage. *arXiv*, October 2023. doi:10.48550/arXiv.2310.14464.
- 23 Nai-Hui Chia, Daniel Liang, and Fang Song. Quantum State Learning Implies Circuit Lower Bounds. *arXiv*, May 2024. arXiv:2405.10242.
- 24 Alexander M Dalzell, Nicholas Hunter-Jones, and Fernando GSL Brandão. Random quantum circuits transform local noise into global white noise. *Communications in Mathematical Physics*, 405(3):78, 2024.
- 25 Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity, 2014. arXiv:1311.2272.
- 26 Abhinav Deshpande, Pradeep Niroula, Oles Shtanko, Alexey V. Gorshkov, Bill Fefferman, and Michael J. Gullans. Tight bounds on the convergence of noisy random circuits to the uniform distribution. *PRX Quantum*, 3(4), December 2022. doi:10.1103/prxquantum.3.040329.

- 27 Bill Fefferman, Soumik Ghosh, Michael Gullans, Kohdai Kuroiwa, and Kunal Sharma. Effect of nonunitary noise on random-circuit sampling. *PRX Quantum*, 5(3), July 2024. doi:10.1103/prxquantum.5.030317.
- 28 Bill Fefferman, Soumik Ghosh, Makrand Sinha, and Henry Yuen. The hardness of learning quantum circuits and its cryptographic applications, 2025. doi:10.48550/arXiv.2504.15343.
- 29 Bill Fefferman, Soumik Ghosh, and Wei Zhan. Anti-concentration for the unitary Haar measure and applications to random quantum circuits, 2024. doi:10.48550/arXiv.2407.19561.
- 30 NIST FIPS203. Module-Lattice-based Key-Encapsulation Mechanism Standard. *Federal Information Processing Standards Publication*, 2023.
- 31 Johan Håstad, Russell Impagliazzo, Leonid A Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.
- 32 Taiga Hiroka and Min-Hsiu Hsieh. Computational complexity of learning efficiently generatable pure states. *arXiv preprint*, 2024. arXiv:2410.04373.
- 33 Taiga Hiroka and Tomoyuki Morimae. Quantum cryptography and meta-complexity, 2024. arXiv:2410.01369.
- 34 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020. doi:10.1038/s41567-020-0932-7.
- 35 Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau, and Jarrod R. McClean. Learning shallow quantum circuits. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1343–1351. ACM, 2023. doi:10.1145/3564246.3585165.
- 36 Hsin-Yuan Huang, John Preskill, and Mehdi Soleimanifar. Certifying almost all quantum states with few single-qubit measurements, 2024. doi:10.48550/arXiv.2404.07281.
- 37 R. Impagliazzo and L.A. Levin. No better ways to generate hard NP instances than picking uniformly at random. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, pages 812–821 vol.2, 1990. doi:10.1109/FSCS.1990.89604.
- 38 Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th Annual Structure in Complexity Theory Conference*, pages 134–147. IEEE Computer Society, 1995. doi:10.1109/SCT.1995.514853.
- 39 Zhengfeng Ji, Yi-Kai Liu, Fang Song, John Watrous, and Henry Yuen. Pseudorandom quantum states. In *Advances in Cryptology – CRYPTO 2018*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018. doi:10.1007/978-3-319-96881-0_5.
- 40 Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994. doi:10.1145/174644.174647.
- 41 Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 968–978, 2024. doi:10.1145/3618260.3649654.
- 42 Dakshita Khurana and Kabir Tomer. Founding Quantum Cryptography on Quantum Advantage, or, Towards Cryptography from $\#P$ -Hardness. *arXiv preprint*, 2024. arXiv:2409.15248.
- 43 Hyun-Soo Kim, Isaac H. Kim, and Daniel Ranard. Learning state preparation circuits for quantum phases of matter, 2024. arXiv:2410.23544.
- 44 Adam R. Klivans and Alexander A. Sherstov. Cryptographic hardness for learning intersections of halfspaces. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 553–562. IEEE, 2006. doi:10.1109/FOCS.2006.24.
- 45 William Kretschmer. Quantum pseudorandomness and classical complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, pages 2:1–2:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.TQC.2021.2.

- 46 William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1589–1602. ACM, 2023. doi:10.1145/3564246.3585172.
- 47 William Kretschmer, Luowen Qian, and Avishay Tal. Quantum-computable one-way functions without one-way functions, 2024. doi:10.48550/arXiv.2411.02554.
- 48 Leslie Lamport. Constructing digital signatures from a one-way function. Technical Report CSL-98, SRI International Computer Science Laboratory, 1979. URL: <https://www.microsoft.com/en-us/research/publication/constructing-digital-signatures-from-a-one-way-function/>.
- 49 Zeph Landau and Yunchao Liu. Learning quantum states prepared by shallow circuits in polynomial time. *arXiv preprint*, 2024. arXiv:2410.23618, doi:10.48550/arXiv.2410.23618.
- 50 Nati Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, Fourier transform, and learnability. *Journal of the Association for Computing Machinery*, 40(3):607–620, 1993. doi:10.1145/174130.174138.
- 51 Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24–28, 2024*, pages 979–990. ACM, 2024. doi:10.1145/3618260.3649650.
- 52 Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010. doi:10.1007/978-3-642-13190-5_1.
- 53 Antonio Anna Mele, Armando Angrisani, Soumik Ghosh, Sumeet Khatri, Jens Eisert, Daniel Stilck França, and Yihui Quek. Noise-induced shallow circuits and absence of barren plateaus, 2024. arXiv:2403.13927.
- 54 Tomoyuki Morimae, Yuki Shirakawa, and Takashi Yamakawa. Cryptographic characterization of quantum advantage, 2024. arXiv:2410.00499.
- 55 Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint*, 2022. arXiv:2210.03394.
- 56 Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Annual International Cryptology Conference*, pages 269–295. Springer, 2022. doi:10.1007/978-3-031-15802-5_10.
- 57 A. Morvan, B. Villalonga, X. Mi, et al. Phase transition in random circuit sampling. *Nature*, 616:70–76, 2023. doi:10.1038/s41586-023-05781-7.
- 58 National Institute of Standards and Technology. FIPS 204: Module-Lattice-Based Digital Signature Standard, 2024. URL: <https://csrc.nist.gov/pubs/fips/204/final>.
- 59 Barak Nehoran and Mark Zhandry. A computational separation between quantum no-cloning and no-telegraphing. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 82:1–82:23. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ITCS.2024.82.
- 60 Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994. doi:10.1016/S0022-0000(05)80043-1.
- 61 Changhun Oh, Liang Jiang, and Bill Fefferman. On classical simulation algorithms for noisy boson sampling, 2023. arXiv:2301.11532.
- 62 Igor C. Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds and pseudorandomness, 2016. arXiv:1611.01190.
- 63 Alexander Poremba, Yihui Quek, and Peter Shor. The learning stabilizers with noise problem, 2024. doi:10.48550/arXiv.2410.18953.
- 64 Luowen Qian, Justin Raizes, and Mark Zhandry. Hard quantum extrapolations in quantum cryptography. *arXiv preprint*, 2024. arXiv:2409.16516.

- 65 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography, 2024. doi:10.48550/arXiv.2401.03703.
- 66 C. Ryan-Anderson, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, E. Nielsen, D. Regalado, S. Ryan, J. Sedlacek, R. T. Sutherland, E. Tirrito, C. Volin, T. Walker, K. White, J. Wootton, and K. Wright. Demonstration of logical qubits and repeated error correction with better-than-physical error rates. *Nature*, 618:500–505, 2023. doi:10.1038/s41586-023-06096-x.
- 67 Mark Saffman et al. Quantum computation with programmable neutral-atom arrays. *Nature Physics*, 2023. URL: <https://quera.ai/research>.
- 68 Christian Weedbrook et al. Quantum advantage with gaussian boson sampling in photonic quantum computers. *Nature Physics*, 2022. Preprint. URL: <https://xanadu.ai/research>.
- 69 Reinhard F Werner. Optimal cloning of pure states. *Physical Review A*, 58(3):1827, 1998.
- 70 Haimeng Zhao, Laura Lewis, Ishaan Kannan, Yihui Quek, Hsin-Yuan Huang, and Matthias C Caro. Learning quantum states and unitaries of bounded gate complexity. *PRX Quantum*, 5(4):040306, 2024.
- 71 Qi Zhao, Hao Chen, Xiao Yuan, et al. Quantum computational advantage via 62-qubit superconducting processor. *Physical Review Letters*, 127(18):180501, 2021. doi:10.1103/PhysRevLett.127.180501.
- 72 Han-Sen Zhong, Yuan Li, et al. Phase-programmable gaussian boson sampling using stimulated squeezed light. *Physical Review Letters*, 127(18):180502, 2021. doi:10.1103/PhysRevLett.127.180502.
- 73 Han-Sen Zhong, Hui Wang, Yi-Han Deng, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. doi:10.1126/science.abe8770.