Pseudorandom Strings from Pseudorandom Quantum States

Prabhanjan Ananth
Department of Computer Science, University of California Santa Barbara, CA, USA

Yao-Ting Lin
Department of Computer Science, University of California Santa Barbara, CA, USA

Henry Yuen
Department of Computer Science, Columbia University, New York, NY, USA

Abstract

We study the relationship between notions of pseudorandomness in the quantum and classical worlds. Pseudorandom quantum state generator (PRSG), a pseudorandomness notion in the quantum world, is an efficient circuit that produces states that are computationally indistinguishable from Haar random states. PRSGs have found applications in quantum gravity, quantum machine learning, quantum complexity theory, and quantum cryptography. Pseudorandom generators, on the other hand, a pseudorandomness notion in the classical world, is ubiquitous to theoretical computer science. While some separation results were known between PRSGs, for some parameter regimes, and PRGs, their relationship has not been completely understood.

In this work, we show that a natural variant of pseudorandom generators called *quantum pseudorandom generators (QPRGs)* can be based on the existence of logarithmic output length PRSGs. Our result along with the previous separations gives a better picture regarding the relationship between the two notions. We also study the relationship between other notions, namely, pseudorandom function-like state generators and pseudorandom functions. We provide evidence that QPRGs can be as useful as PRGs by providing cryptographic applications of QPRGs such as commitments and encryption schemes.

Our primary technical contribution is a method for pseudodeterministically extracting uniformly random strings from Haar-random states.

2012 ACM Subject Classification Security and privacy → Mathematical foundations of cryptography; Theory of computation → Quantum computation theory

Keywords and phrases Quantum Cryptography

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.6

Related Version *Full Version*: https://arxiv.org/abs/2306.05613

1 Introduction

Deterministically generating long pseudorandom strings from a few random bits is a fundamental task in classical cryptography. Pseudorandom generators (PRGs) are a primitive that achieves this task and are ubiquitous throughout cryptography. Beyond cryptography, pseudorandom generators have found applications in complexity theory [31, 21] and derandomization [29, 18].

The concept of pseudorandomness has also been explored in other contexts. Of interest is the notion of pseudorandom quantum states, a popular pseudorandomness notion studied in the quantum setting. Pseudorandom quantum states (PRS), introduced by Ji, Liu, and Song [19], are efficiently computable states that are computationally indistinguishable from Haar-random states. PRS have found numerous applications in other areas such as physics [7, 6], quantum machine learning [17], and cryptography [2, 25].

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).
Editor: Venkatesan Guruswami; Article No. 6; pp. 6:1–6:22
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

While some recent works make progress towards understanding the feasibility of PRSGs (PRS generators), its relationship with PRGs[1] and its implications to cryptography, there are still some important gaps that are yet to be filled. While the directions listed below might come across as seemingly unrelated, we will discuss how our work addresses all these three different directions.

Direction 1. Separating PRSGs and PRGs. We summarize the implication from PRGs to PR**S**Gs[2] and back in the table below. [19] showed that any quantum-query secure PRF (implied by PRGs) implies the existence of $\omega(\lambda)$-output length PR**S**Gs, where λ is the seed length. On the other hand, Kretschmer [20] showed a separation between $\omega(\lambda)$-length PR**S**Gs and PRGs. In the $c \cdot \log(\lambda)$-regime, where $c \in \mathbb{R}$ and $c \ll 1$, [8] showed that $c \cdot \log(\lambda)$-length PR**S**Gs can be constructed unconditionally. Hence there is a trivial separation between $c \cdot \log(\lambda)$-length PR**S**Gs, with $c \ll 1$, and PRGs since the latter requires computational assumptions. In the same work, [8] showed that $c \cdot \log(\lambda)$-length PR**S**Gs, when $c \gg 1$, can be constructed from PRGs. However, whether PR**S**Gs with output length at least $\log(\lambda)$ imply PRGs or are separated from it is currently unknown.

Output Length of PRSG	**Implied by PRG?**	**Implication to PRG?**
$\omega(\log(\lambda))$	Yes [19]	Black-box separation [20]
$c \cdot \log(\lambda)$), $c \geq 1$	Yes [8]	**unknown**
$c \cdot \log(\lambda)$), $c \ll 1$	N/A (Information-theoretic [8])	Separation (trivial)

Thus, the following question has been left open.

Does $\log(\lambda)$*-length PRS imply PRGs or is it (black-box) separated from PRGs?*

Direction 2. Hybrid Cryptography. While the recent results demonstrate constructions of quantum cryptographic tasks such as commitments, zero-knowledge and secure computation from assumptions potentially weaker than one-way functions, the main drawback of these constructions is that they require the existence of quantum communication channels, an undesirable feature. Starting with the work of Gavinsky [12], there has been an effort in building quantum cryptographic primitives using classical communication channels. We can thus characterize the class of cryptographic primitives into three categories: classical cryptography (that uses only classical resources), hybrid cryptography (uses quantum computing but classical communication channels) and quantum cryptography (no restrictions). Hybrid cryptography has the advantage that the primitives in this category could be based on assumptions weaker than classical cryptography but on the other hand, has the advantage that we only need classical communication channels. Towards a deeper understanding of hybrid cryptography, the following is a pertinent question:

Identify foundational primitives in hybrid cryptography and understand their relationship with classical and quantum cryptographic primitives.

[1] We are interested in PRGs that guarantee security against efficient quantum adversaries. Typically, such PRGs are referred to as post-quantum PRGs. For brevity, henceforth, by PRGs we will mean post-quantum PRGs.

[2] The "S" is emphasized in this paragraph to highlight the difference between PR**S**Gs and PRGs.

Direction 3. Domain Extension, Generically. Given any pseudorandom generator of output length m, we know how to generically transform it into another secure pseudorandom generator of length ℓ, for any $\ell > m$. On the other hand, we have limited results for pseudorandom quantum states. Recently, [16] showed that multi-copy $\omega(\lambda)$-length PRS implies a single-copy PRS with large output length. However, we are not aware of any length extension transformation that preserves the number of copies. Investigating this question will help us understand the relationship between PRSs of different output lengths. This leads to the following question:

Can we generically transform a multi-copy n-output PRS into a multi-copy ℓ-output PRS, where $\ell \gg n$? Or is there a black-box separation?

Our Work

Towards simultaneously addressing all three directions above, we introduce the notion of *quantum pseudorandom generators (QPRGs)*: which are like classical PRGs in that the input is a short classical string and the output is a longer classical string that is computationally indistinguishable from uniform, but (a) generation algorithm is a quantum algorithm, and (b) the mapping from seed to output only has to be *pseudodeterministic* (i.e., for a fixed seed, the output is a fixed string with high probability). We first show that assumptions that are plausibly weaker than the existence of classical OWFs/PRGs can be used to build QPRGs: we show that QPRGs can be constructed from logarithmic-output PRSGs. In other words, we can generate pseudorandom strings using pseudorandom quantum states in a (pseudo-)deterministic fashion. We then present cryptographic applications of QPRGs and highlight some implications for the structure of classical versus quantum cryptography.

The reader might wonder whether the notion of quantum generation of classical pseudorandomness is trivial. After all, since quantum computation is inherently probabilistic and can generate unlimited randomness starting from a fixed input, why would one need *pseudo*randomness? However, for cryptographic applications having a source of randomness is not enough; it is important that some random-looking string can be *deterministically generated* using a secret key.

1.1 Our Results

Quantum PRGs from PRS

Informally, a $(1-\varepsilon)$-pseudodeterministic QPRG is a quantum algorithm G where

- (*Pseudodeterminism*) For $1-\varepsilon$ fraction of seeds $k \in \{0,1\}^\lambda$ outputs a fixed string $y_k \in \{0,1\}^n$ with probability at least $1-\varepsilon$, and
- (*Pseudorandomness*) For all efficient quantum distinguishers A,

  $$\left| \Pr_{k \leftarrow \{0,1\}^\lambda}[A(G(k)) = 1] - \Pr_{y \leftarrow \{0,1\}^n}[A(y) = 1] \right| \leq \mathsf{negl}(\lambda) \ .$$

  In other words, no efficient quantum adversary can distinguish between the output of the generator and a uniformly random string.

(See Section 4 for a formal definition of QPRGs). Our first result is the following:

▶ **Theorem 1** (Informal). *Assuming the existence of logarithmic PRS, there exist $\left(1 - \frac{1}{\mathsf{poly}(\lambda)}\right)$-pseudodeterministic QPRGs.*

Our result has several implications.

- **Implication to Direction 1 (Separating PRSGs and PRGs).** Perhaps surprisingly, our result suggests that in the logarithmic output regime, PRS and PRGs are not separated. This is unlike the super-logarithmic output regime and sub-logarithmic output regime both of which are separated from PRGs. In contrast, in the classical setting, achieving more pseudorandom bits is harder (in some cases strictly [30]) than achieving a few pseudorandom bits. Our result when combined with prior results [19, 8] gives a better picture on the relationship between PRS and PRGs (refer to the table in the introduction).
- **Implication to Direction 2 (Hybrid Cryptography).** Quantum pseudorandom generator is a hybrid cryptographic primitive since it has a quantum generation algorithm but classical inputs and outputs. Later, we show that QPRGs imply many other hybrid cryptographic primitives such as quantum bit commitments with classical communication, quantum encryption with classical ciphertexts, and so on. Our results suggest that QPRGs could play a similar role in hybrid cryptography as PRGs did in classical cryptography. Proving whether QPRGs is a minimal assumption in hybrid cryptography, akin to how PRGs is a minimal assumption in classical cryptography [13], is an interesting open question. Furthermore, our result above highlights connections between hybrid and quantum cryptography.
- **Implication to Direction 3 (Domain Extension, Generically).** In the above result, unfortunately, we only obtain QPRGs with inverse polynomial pseudodeterminism error. Suppose we can reduce the error to be negligible then we claim that $O(\log(\lambda))$-output PRS can be generically transformed into $\omega(\log(\lambda))$-output PRS. This can be achieved by appropriately instantiating the construction of Ji, Liu, and Song [19] using quantum pseudorandom generators, which in turn can be built from $O(\log(\lambda))$-output PRS. Thus, the question of whether we can generically increase the output length of PRS is related to the question of reducing pseudodeterminism error in the above theorem.
- **Other Implications: New Approach to Pseudorandomness.** One implication of our QPRG construction is that it demonstrates an "inherently quantum" way to generate classical pseudorandomness. There are plausible candidates for PRS (even the logarithmic-length ones) that don't seem to involve any classical OWFs in them at all; for example, it is conjectured that random polynomial-size quantum circuits generate pseudorandom states [2].

Applications of Quantum PRGs

Next we investigate the cryptographic applications of QPRGs. We demonstrate that QPRGs can effectively replace classical pseudorandom generators in some applications; although the QPRGs are not entirely deterministic, being $(1 - \frac{1}{\mathsf{poly}(\lambda)})$-pseudodeterministic is good enough.

Concretely, we explore two applications: statistically binding and computationally hiding commitments, and pseudo one-time pads. While [1] previously demonstrated that these applications can be based on logarithmic PRS, we provide alternate proofs assuming the existence of QPRGs combined with Theorem 1. Moreover, our constructions resemble the textbook constructions of classical commitments and pseudo one-time pads and thus, are conceptually simpler than the ones presented by [1].

A statistically binding commitment scheme is a fundamental cryptographic notion where a sender commits to a value such that it is infeasible, even if it is computationally unbounded, for them to change their commitment to a different value. Statistically binding quantum commitments have been a critical tool to achieve another fundamental notion in cryptography, namely secure computation [4, 15]. We demonstrate that statistically binding and computationally hiding commitments can be constructed from QPRGs.

▶ **Theorem 2** (Informal). *Assuming the existence of* $(1 - \frac{1}{\mathsf{poly}(\lambda)})$*-pseudodetermininistic QPRGs, there exist statistically binding and computationally hiding quantum commitments with classical communication.*

It is worth mentioning that there is another recent work [5] that also builds quantum commitments with classical communication albeit from incomparable assumptions[3].

Pseudo one-time pads are a variation of the one-time pad encryption scheme, where the encryption key is much smaller than the message length. As demonstrated by [2], pseudo one-time pads are useful for constructing classical garbling schemes [3] and quantum garbling schemes [9], which have numerous applications in cryptography. We demonstrate that pseudo one-time pads can be constructed from QPRGs.

▶ **Theorem 3** (Informal). *Assuming the existence of* $(1 - \frac{1}{\mathsf{poly}(\lambda)})$*-pseudodetermininistic QPRGs, there exist pseudo one-time pads.*

Quantum PRFs

In addition to the above, we also explore pseudorandom functions with a quantum generation algorithm, which we call quantum pseudorandom functions (QPRFs). We show the following theorem:

▶ **Theorem 4** (Informal). *Assuming the existence of* $(\omega(\log \lambda), O(\log \lambda))$*-PRFS, there exists a quantum pseudorandom function (QPRF) satisfying determinism with probability at least* $\left(1 - \frac{1}{\mathsf{poly}(\lambda)}\right)$.

In the above theorem, we use pseudorandom function-like states [2], a quantum analog of pseudorandom functions, to accomplish this. The notion of pseudorandom function-like states says the following: t copies of states $(|\psi_{x_1}\rangle, \ldots, |\psi_{x_q}\rangle)$ are computationally indistinguishable from t copies of q Haar states, where $|\psi_{x_i}\rangle$, for every $i \in [q]$, is produced using an efficient PRFS generator that receives as input a key $k \in \{0,1\}^\lambda$, picked uniformly at random, and an input $x_i \in \{0,1\}^\lambda$. Just like in the case of QPRGs, in the above theorem, we require PRFS with logarithmic input length.

We show how to leverage QPRFs to achieve private-key encryption with QPT algorithms and classical communication, which is the first result to achieve this notion from assumptions potentially weaker than one-way functions.

1.2 Future Directions

Our research raises several important open questions that remain to be explored. Below, we highlight two particularly interesting ones.

Separating QPRGs and QPRFs from Classical Cryptography

While QPRGs and QPRFs are similar in flavor to their classical counterparts, their ability to generate quantum states suggests that they may be based on weaker assumptions than classical pseudorandom generators and functions. A key question is whether there is a fundamental separation between QPRGs and PRGs, as well as between QPRFs and PRFs.

[3] They consider a variant of PRS referred to as PRS with proof of deletion. On one hand, they don't have restriction on the output length like we do and on the other hand, they assume that PRS satisfies the additional proof of deletion property whereas we don't.

Proving that there is no separation would require a mechanism to efficiently dequantize the generation algorithm, which is a challenging task. This is especially true if the quantum generation algorithm involves running a quantum algorithm that is believed to be difficult to efficiently dequantize; for example, Shor's algorithm.

Reducing Determinism Error

One limitation of both QPRGs and QPRFs is that they suffer from inverse polynomial determinism error. It would be interesting to explore whether this error can be reduced to negligible, or whether a negative result can be proven. Understanding the fundamental limits of determinism in quantum pseudorandomness could have important implications. For instance, due to the inverse polynomial error, it is unclear how to apply the GGM transformation [14] to go from quantum pseudorandom generators to quantum pseudorandom functions.

1.3 Technical Overview

We summarise our technical contributions below:

- We identify the definition of pseudodeterministic extractor that gives quantum pseudorandom generators. We then realize the notion of pseudodeterministic extractors; this is our core technical contribution and it involves using interesting properties about the Haar measure in the analysis.
- We define and realize quantum pseudorandom functions from logarithmic output pseudorandom function-like states. Defining quantum PRFs turns out to be subtle.
- We demonstrate applications of quantum pseudorandom generators and functions to commitments and encryption schemes. Especially, in commitment schemes, it turns out to be tricky to argue security due to the inverse polynomial determinism error.

1.3.1 Core Contribution: Pseudodeterministic Extractor

We focus on the goal of building a quantum pseudorandom generator from $O(\log(\lambda))$-qubit pseudorandom quantum states. Towards this goal, we identify the important step as follows: extracting $\mathsf{poly}(d(\lambda))$-length binary strings from $\log(d(\lambda))$-qubit Haar states in such a way, the following key properties are satisfied:

- **Pseudodeterminism:** Running the extraction process on $\mathsf{poly}(d(\lambda))$-copies of $|\psi\rangle$ should give the same string y with a very high probability. Ideally, with probability at least $1 - \frac{1}{\mathsf{poly}(d(\lambda))}$,
- **Efficiency:** The extraction process should run in time $\mathsf{poly}(d(\lambda))$,
- **Statistical Indistinguishability:** The string y is statistically close to the uniform distribution over $\{0,1\}^{\mathsf{poly}(d(\lambda))}$ as long as $|\psi\rangle$ is sampled from the Haar distribution. Here, we allow the total variation distance error to be as large as $O(\frac{1}{d})$.

It turns out most of the work goes in achieving the pseudodeterminism property.

Toy Case

Towards designing an extractor satisfying the above three properties, we first consider an alternate task. Instead of $\mathsf{poly}(d(\lambda))$-copies of the $\log(d(\lambda))$-qubit state $|\psi\rangle$, we are given all the amplitudes of $|\psi\rangle$, say $(\alpha_1, \ldots, \alpha_{d(\lambda)})$, in the clear. Can we extract true randomness from this? For instance, we could extract $b_1, \ldots, b_{d(\lambda)}$, where b_i is the first bit of the real component of α_i. Firstly, it is not even clear that b_i is distributed to according to the uniform distribution over $\{0,1\}$. Moreover, all the bits $b_1,\ldots,b_{d(\lambda)}$ are not independent and in fact, are correlated with each other due to the normalization condition $\sum_i |\alpha_i|^2 = 1$.

Fortunately, we can rely upon a result in random matrix theory [24], that states the following: suppose $(\alpha_1,\ldots,\alpha_{d(\lambda)})$ are drawn from a Haar measure in $\mathcal{S}(\mathbb{R}^d)$ then it holds that any $o(d(\lambda))$ co-ordinates of $(\alpha_1,\ldots,\alpha_{d(\lambda)})$ are $1/o(d(\lambda))$-close in total variation distance with $o(d(\lambda))$-dimensional vector where each component is drawn from i.i.d Gaussian $\mathcal{N}(0,\frac{1}{d})$.

We generalize this result to the case when $(\alpha_1,\ldots,\alpha_{d(\lambda)})$ are drawn from a Haar measure in $\mathcal{S}(\mathbb{C}^d)$, and not $\mathcal{S}(\mathbb{R}^d)$ (see Corollary 17) at the cost of reducing the standard deviation from $\frac{1}{d}$ to $\frac{1}{2d}$. We then use our observation to come up with an extractor as follows. The extractor takes as input[4] $(\alpha_1,\ldots,\alpha_{d(\lambda)})$,

- Choose the first $k = o(d(\lambda))$ entries among $(\alpha_1,\ldots,\alpha_{d(\lambda)})$.
- Rounding step: for every $i \in [k]$, if $\mathsf{Re}(\alpha_i) > 0$, then set $b_i = 0$. Otherwise, set $b_i = 1$.
- Output $b_1 \cdots b_k$.

From our observation and the symmetricity of $\mathcal{N}(0,\frac{1}{d})$, it follows that when $(\alpha_1,\ldots,\alpha_{d(\lambda)})$ is drawn from a Haar distribution on $\mathcal{S}(\mathbb{C}^d)$ then the output of the extractor is $o(\frac{1}{d(\lambda)})$-close to uniform distribution on $\{0,1\}^k$. Moreover, the above procedure is deterministic.

Challenges

Our hope is to leverage the above ideas to design an extractor that can extract given $\mathsf{poly}(d(\lambda))$-copies of a $O(\log(d(\lambda)))$-qubit Haar state $|\psi\rangle$. We encounter a couple of challenges.

1. First challenge: We have access only to the copies of $|\psi\rangle\langle\psi|$ without the amplitudes given to us in plain text, making it infeasible to implement the previously described method. However, we can still carry out tomography and retrieve an estimated version of the matrix $|\psi\rangle\langle\psi|$. If the amplitudes of $|\psi\rangle$ are $\{\alpha_x\}_{x\in[d]}$ then the $(x,y)^{th}$ entry in the density matrix $|\psi\rangle\langle\psi|$ is $\alpha_x\alpha_y^*$. We need to analyze the distribution corresponding to $\alpha_x\alpha_y^*$ and, design an approach for obtaining a uniform distribution from it.
2. Second challenge: Tomography is inherently a probabilistic technique, and hence, each time tomography is executed on multiple copies of $|\psi\rangle$, the output obtained may differ. Additionally, the trace distance between the density matrix obtained via tomography and the original density matrix is inversely proportional to the dimension, which is polynomial in this case, and this may be significant. Both of these factors collectively affect the determinism guarantees of the extractor. In general, it is not feasible to partition $\mathcal{S}(\mathbb{C}^d)$ into regions labeled by a bitstring such that given multiple copies of a state in a region, the corresponding bitstring can be deterministically recovered.

We tackle the above challenges using the following insights.

Addressing the first challenge

We first tackle the first bullet above. Notice that the diagonal entries in the density matrix $|\psi\rangle\langle\psi|$ is $\{|\alpha_i|^2\}_{i\in[d(\lambda)]}$, where $|\psi\rangle = \sum_i \alpha_i |i\rangle$. If $\alpha_j = a_j + ib_j$ then $|\alpha_j|^2 = a_j^2 + b_j^2$. Given our earlier observation about the closeness of $o(d(\lambda))$ entries in a vector drawn from $\mathcal{S}(\mathbb{C}^d)$ with iid Gaussian, we will make the following simplifying assumption. We assume that $(\alpha_1,\ldots,\alpha_k)$, where $k = o(d)$, is sampled such that for every $i \in [k]$, a_i and b_i are distributed according to i.i.d Gaussian $\mathcal{N}(0,\frac{1}{2d})$. From this, it follows that $|\alpha_i|^2$ is distributed according to

[4] For the current discussion, we assume that the extractor has an infinite input tape that allows for storing infinite bits of precision of the complex numbers.

a *chi-squared* distribution with 2 degrees of freedom. Unfortunately, chi-squared distribution does not have the same nice symmetricity property as a Gaussian distribution. So we will instead extract randomness in a different way.

We divide $(|\alpha_1|^2, \ldots, |\alpha_k|^2)$ into blocks of size r and denote ℓ to be the number of blocks, where $r, \ell = o(d)$. Then, add the elements in a block. Call the resulting elements $q_1, \ldots, q_\ell$. From central limit theorems [32], one can show that $q_1, \ldots, q_\ell$ are $O(1/\sqrt{r})$-close to ℓ samples drawn i.i.d from $\mathcal{N}(\frac{r}{d}, \frac{r}{d^2})$. Thus, using central limit theorem, we are back to the normal distribution, except that the mean is shifted to $\frac{r}{d}$ rather than 0. This gives rise to a natural rounding mechanism.

We will check if $q_i > \frac{r}{d}$ and if so, we set a bit $b_i = 0$ and if not, we set it to be 0. By carefully choosing the parameters k and ℓ and combining the above observations, we can argue that $b_1, \ldots, b_k$ is $O(d^{-1/6})$-close to the uniform distribution on $\{0,1\}^\ell$.

To summarise, the informal description of the extractor is as follows: given $\mathsf{poly}(d)$ copies of a d-dimensional state $|\psi\rangle$,

- First perform tomography to recover a matrix $M \in \mathbb{C}^d \times \mathbb{C}^d$ that is an approximation of $|\psi\rangle\langle\psi|$.
- Then, pick $o(d)$ diagonal entries in M and break this into ℓ blocks of size r.
- Sum up all the entries in each block to get ℓ values $q_1, \ldots, q_\ell$. Round every q_i to get b_i.
- Output $b_1, \ldots, b_\ell$.

Addressing the second challenge

While the above construction seems promising, we still have not addressed the second challenge pertaining to the determinism property. It could be the case that all the q_is are very close to the mean and due to the tomography error, every time we try to extract we set $b_i = 0$ sometimes and $b_i = 1$ the rest of the time. This should not be surprising as we said earlier, that it should not be possible to partition $\mathcal{S}(\mathbb{C}^d)$ such that for every $|\psi\rangle$, there is a bitstring b_ψ such that given many copies of $|\psi\rangle$, the extractor always outputs the same bitstring b_ψ.

In fact, we can identify a *forbidden region* in $\mathcal{N}(\frac{r}{d}, \frac{r}{d^2})$ (see Figure 1 below) such that if q_i falls into the forbidden region then there is a significant chance that q_i will be classified as either 0 or 1. The forbidden region has width $\frac{1}{d}$ on either side of the mean. Given this, we give up all hope of achieving perfect determinism and instead shoot for determinism with $o(1/d)$ error.

We identify a set of $d(\lambda)$-dimensional states $\mathcal{G}_\Delta$, where $\Delta = \frac{1}{d}$, such that if a state $|\psi\rangle$ is in $\mathcal{G}_\Delta$ then it holds that none of $q_1, \ldots, q_\ell$, generated from $|\psi\rangle$, lies in the forbidden region. The setting of Δ is carefully chosen to accommodate for the error in tomography.

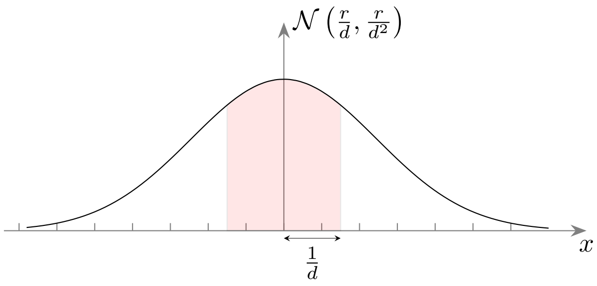

Figure 1 The red region denotes the *forbidden region*.

Once $\mathcal{G}_\Delta$ is identified, we prove two things:

- Firstly, if a state is sampled from the Haar distribution on $\mathcal{S}(\mathbb{C}^{d(\lambda)})$ then with at least $1 - o\left(\frac{1}{d}\right)$ probability, $|\psi\rangle \in \mathcal{G}_\Delta$.
- Secondly, for every $|\psi\rangle \in \mathcal{G}_\Delta$, the probability that the extractor, given $\mathsf{poly}(d(\lambda))$-copies of $|\psi\rangle$, outputs the same string twice is at least $1 - o\left(\frac{1}{d}\right)$. Roughly, this follows from the fact that $(q_1, \ldots, q_\ell)$, generated from $|\psi\rangle$, gets misclassified with very small probability.

We can leverage the above two observations to show that our extractor satisfies determinism with probability at least $1 - o\left(\frac{1}{d}\right)$.

1.3.2 From Pseudodeterminism Extractor to Quantum PRGs

With the pseudodeterministic extractor in hand, we propose the following construction of quantum pseudorandom generators: on input a seed $k \in \{0,1\}^\lambda$,

- Perform this procedure polynomially many times: run the PRS generator on k to produce the PRS state $|\psi\rangle$,
- Run the pseudodeterministic extractor (discussed in Section 1.3.1) on polynomially many copies of $|\psi\rangle$ to obtain a binary string y,
- Output y.

Recall that the guarantees of the pseudodeterministic extractor only hold for Haar states. In the above construction, we are invoking the extractor on PRS states. However, we can invoke the security of PRSGs to replace PRS states with Haar states and then invoke the guarantees of the extractor. Just like the extractor, the above quantum PRG would also suffer from an inverse polynomial determinism error. Similarly, the above construction would suffer from inverse polynomial error in security; that is, the output of the above QPRG (on a random seed) cannot be distinguished from a uniformly random output with at most inverse polynomial error. We call a QPRG that satisfies inverse polynomial determinism error and inverse polynomial security error to be a *weak* QPRG and a QPRG that satisfies negligible security error to be a *strong* QPRG.

While we currently do not know how to reduce the pseudodeterminism error, there is still hope to reduce the security error. Indeed, there are security amplification techniques that are well studied in the classical cryptography literature.

From Weak QPRG to Strong QPRG

The naive approach of going from a weak QPRG to a strong QPRG is to use parallel repetition: on input a seed of length $s \cdot \lambda$, break the seed into s parts, apply QPRG on each of them and then XOR the outputs. While this should help security (as we see below), it hurts pseudodeterminism. Using union bound, we can argue that the pseudodeterminism error increases by a multiplicative factor of λ. Thus, if we start with a weak PRG with a sufficiently small pseudodeterminism error then this multiplicative factor won't hurt us.

To prove that the security error is negligible, we will use XOR-based security amplification techniques [11, 22, 23]. The analysis in the amplification theorems [11, 22, 23] was initially tailored to classical settings, using a careful analysis, we show that the analysis also extends to the quantum setting.

1.3.3 Quantum Pseudorandom Functions

We also demonstrate the connections between quantum pseudorandom functions and logarithmic-output pseudorandom function-like states [1].

Roughly speaking, a quantum pseudorandom function is a pseudorandom function except that the generation algorithm is quantum and moreover, we allow for inverse polynomial determinism error. Defining the security of this notion requires some care. If we allow the adversary to make arbitrary queries to the oracle (that is either the QPRF or the random function) then such a notion is clearly impossible to achieve. For instance, the adversary can query the same input twice; in the case when the oracle implements a random function, we always get the same output, and in the QPRF case, there is an inverse polynomial probability with which we get different outputs. Hence, we restrict our attention to the setting when the adversary only makes selective and distinct queries. We show that this definition is sufficient for applications.

The construction of QPRFs from logarithmic-output PRFS follows along similar lines as the construction of QPRGs from pseudodeterminism extractors.

1.3.4 Applications

We show that QPRGs imply both statistically binding commitments and pseudo-one-time pads. The constructions are similar to the existing constructions from (classical) pseudorandom generators except that we need to contend with the inverse polynomial determinism error. For most of the applications, naive parallel repetition and a majority argument are sufficient to circumvent the determinism issue. However, for the application of commitments, the analysis turns out to be relatively more complicated.

Commitments

The construction of statistically binding commitments from QPRGs is inspired by Naor commitments [27].

To recall Naor's construction: the receiver sends a random string r of length 3λ to the sender who applies a classical pseudorandom generator with output length 3λ on a random seed of length λ. Depending on the message bit, the sender either XORs the output with r or sends the PRG output as-is. The proof of binding relies upon the fact that the number of pairs of keys whose outputs when XORed with each other lead to r is precisely upper bounded by $2^{2\lambda}$, which is negligible in comparison with all possible values of r.

A natural modification to the above construction is to replace the PRG with quantum pseudorandom generator. The immediate issue that arises here is correctness due to the inverse polynomial determinism error. Again, using naive parallel repetition we can resolve the determinism error: where the sender computes many QPRG outputs on independent seeds and depending on the message, the outputs are either XORed with r or kept as-is. In the modified construction, arguing hiding is fairly straightfoward. However, arguing the binding property requires some care.

Naor's binding argument cannot be immediately generalized to the QPRG setting since it has inverse polynomial determinism error. However, we come up with a different argument in this setting: in two technical claims (see full version) we prove that the statistical binding property still holds. Roughly speaking, the intuition behind the argument is as follows. Suppose Bad be the set of QPRG seeds where the pseuodeterminism error is too high; larger than any inverse polynomial and Good be the set containing the rest of the QPRG seeds. If the adversarial sender chooses from Bad in the commit phase (or even in the opening phase), it could only hurt itself because it will not be able to control the output of the QPRG during the verification process executed by the receiver in the opening phase. On the other hand, if the adversarial sender commits to seed from Good in the commit phase and sends (a possibly different) seed from Good in the opening phase then using the fact that the outputs are mostly deterministic, we can argue that with overwhelming probability over r, the XOR of the two seeds does not equal r.

2 Preliminaries

We refer the reader to [28] for a comprehensive reference on the basics of quantum information and quantum computation. We use I to denote the identity operator. We use $\mathcal{S}(\mathcal{H})$ to denote the set of unit vectors in the Hilbert space $\mathcal{H}$. We use $\mathcal{D}(\mathcal{H})$ to denote the set of density matrices in the Hilbert space $\mathcal{H}$. Let P, Q be distributions. We use $d_{TV}(P, Q)$ to denote the total variation distance between them. Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be density matrices. We write $\mathrm{TD}(\rho, \sigma)$ to denote the trace distance between them, i.e.,

$$\mathrm{TD}(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$$

where $\|X\|_1 = \mathrm{Tr}(\sqrt{X^\dagger X})$ denotes the trace norm. We denote $\|X\| := \sup_{|\psi\rangle}\{\langle\psi|X|\psi\rangle\}$ to be the operator norm where the supremum is taken over all unit vectors. For a vector $|x\rangle$, we denote its Euclidean norm to be $\||x\rangle\|_2$. We use the notation $M \geq 0$ to denote the fact that M is positive semi-definite.

Haar Measure

The Haar measure over $\mathbb{C}^d$, denoted by $\mathscr{H}(\mathbb{C}^d)$ is the uniform measure over all d-dimensional unit vectors. One useful property of the Haar measure is that for all d-dimensional unitary matrices U, if a random vector $|\psi\rangle$ is distributed according to the Haar measure $\mathscr{H}(\mathbb{C}^d)$, then the state $U|\psi\rangle$ is also distributed according to the Haar measure. For notational convenience we write $\mathscr{H}_m$ to denote the Haar measure over m-qubit space, or $\mathscr{H}((\mathbb{C}^2)^{\otimes m})$.

2.1 Quantum Algorithms

A quantum algorithm A is a family of generalized quantum circuits $\{A_\lambda\}_{\lambda\in\mathbb{N}}$ over a discrete universal gate set (such as $\{CNOT, H, T\}$). By generalized, we mean that such circuits can have a subset of input qubits that are designated to be initialized in the zero state, and a subset of output qubits that are designated to be traced out at the end of the computation. Thus a generalized quantum circuit A_λ corresponds to a *quantum channel*, which is a is a completely positive trace-preserving (CPTP) map. When we write $A_\lambda(\rho)$ for some density matrix ρ, we mean the output of the generalized circuit A_λ on input ρ. If we only take the quantum gates of A_λ and ignore the subset of input/output qubits that are initialized to zeroes/traced out, then we get the *unitary part* of A_λ, which corresponds to a unitary operator which we denote by $\hat{A}_\lambda$. The *size* of a generalized quantum circuit is the number of gates in it, plus the number of input and output qubits.

We say that $A = \{A_\lambda\}_\lambda$ is a *quantum polynomial-time (QPT) algorithm* if there exists a polynomial p such that the size of each circuit A_λ is at most $p(\lambda)$. We furthermore say that A is *uniform* if there exists a deterministic polynomial-time Turing machine M that on input 1^λ outputs the description of A_λ.

We also define the notion of a *non-uniform* QPT algorithm A that consists of a family $\{(A_\lambda, \rho_\lambda)\}_\lambda$ where $\{A_\lambda\}_\lambda$ is a polynomial-size family of circuits (not necessarily uniformly generated), and for each λ there is additionally a subset of input qubits of A_λ that are designated to be initialized with the density matrix ρ_λ of polynomial length. This is intended to model nonuniform quantum adversaries who may receive quantum states as advice. Nevertheless, the reductions we show in this work are all uniform.

The notation we use to describe the inputs/outputs of quantum algorithms will largely mimick what is used in the classical cryptography literature. For example, for a state generator algorithm G, we write $G_\lambda(k)$ to denote running the generalized quantum circuit G_λ on input $|k\rangle\langle k|$, which outputs a state ρ_k.

Ultimately, all inputs to a quantum circuit are density matrices. However, we mix-and-match between classical, pure state, and density matrix notation; for example, we may write $A_\lambda(k, |\theta\rangle, \rho)$ to denote running the circuit A_λ on input $|k\rangle\langle k| \otimes |\theta\rangle\langle\theta| \otimes \rho$. In general, we will not explain all the input and output sizes of every quantum circuit in excruciating detail; we will implicitly assume that a quantum circuit in question has the appropriate number of input and output qubits as required by context.

2.2 Pseudorandomness Notions

The notion of pseudorandom quantum states was first introduced by Ji, Liu, and Song in [19]. We present the following relaxed definition of pseudorandom state (PRS) generators.[5] We note that the relaxation is due to [2].

▶ **Definition 5** (Pseudorandom State (PRS) Generator). *We say that a QPT algorithm G is a* pseudorandom state (PRS) generator *if the following holds.*

1. **State Generation.** *For all $\lambda \in \mathbb{N}$ and all $k \in \{0,1\}^\lambda$, the algorithm G behaves as $G_\lambda(k) = \rho_k$ for some $n(\lambda)$-qubit (possibly mixed) quantum state ρ_k.*
2. **Pseudorandomness.** *For all polynomials $t(\cdot)$ and any (non-uniform) QPT distinguisher A, there exists a negligible function $\varepsilon(\cdot)$ such that for all $\lambda \in \mathbb{N}$, we have*

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda}[A_\lambda(G_\lambda(k)^{\otimes t(\lambda)}) = 1] - \Pr_{|\vartheta\rangle \leftarrow \mathscr{H}_{n(\lambda)}}[A_\lambda(|\vartheta\rangle^{\otimes t(\lambda)}) = 1] \right| \leq \varepsilon(\lambda).$$

*We also say that G is an $n(\lambda)$-*PRS generator *to succinctly indicate that the output length of G is $n(\lambda)$.*

▶ **Definition 6** (Selectively Secure Pseudorandom Function-Like State (PRFS) Generators). *We say that a QPT algorithm G is a* selectively secure pseudorandom function-like state (PRFS) generater *if for all polynomials $q(\cdot), t(\cdot)$, any (non-uniform) QPT distinguisher A, and any family of pairwise distinct indices $\left(\{x_1, \ldots, x_{q(\lambda)}\} \subseteq \{0,1\}^{m(\lambda)}\}\right)_\lambda$, there exists a negligible function $\varepsilon(\cdot)$ such that for all $\lambda \in \mathbb{N}$,*

$$\Big| \Pr_{k \leftarrow \{0,1\}^\lambda}\left[A_\lambda(x_1, \ldots, x_{q(\lambda)}, G_\lambda(k, x_1)^{\otimes t(\lambda)}, \ldots, G_\lambda(k, x_{q(\lambda)})^{\otimes t(\lambda)}) = 1\right]$$
$$- \Pr_{|\vartheta_1\rangle, \ldots, |\vartheta_{q(\lambda)}\rangle \leftarrow \mathscr{H}_{n(\lambda)}}\left[A_\lambda(x_1, \ldots, x_{q(\lambda)}, |\vartheta_1\rangle^{\otimes t(\lambda)}, \ldots, |\vartheta_{q(\lambda)}\rangle^{\otimes t(\lambda)}) = 1\right] \Big| \leq \varepsilon(\lambda).$$

*We also say that G is an $(m(\lambda), n(\lambda))$-*PRFS generator *to succinctly indicate that its input length is $m(\lambda)$ and its output length is $n(\lambda)$.*

2.3 Basics of Statistics and Haar Measure

A simple yet useful observation is that for any two density matrices, the difference between any of their diagonal entries is bounded above by their trace distance.

▶ **Fact 7.** *For any density matrices $\rho, \sigma \in \mathcal{D}(\mathbb{C}^d)$, it holds that $\max_{i \in [d]} |\rho_{ii} - \sigma_{ii}| \leq \mathrm{TD}(\rho, \sigma)$, where ρ_{ii}, σ_{ii} denote the i-th diagonal entry of ρ, σ respectively, i.e., $\rho_{ii} = \langle i|\rho|i\rangle$ and $\sigma_{ii} = \langle i|\sigma|i\rangle$.*

[5] In [19], the output of the generator needs to be pure; while we allow it to be mixed.

Proof. Note that the trace distance has the following variational form:

$$\mathrm{TD}(\rho,\sigma) = \max_{0 \leq M \leq I} \mathrm{Tr}(M(\rho-\sigma)).$$

Furthermore, trace distance is symmetric. Therefore, taking $M := |i\rangle\langle i|$ for $i \in [d]$, we have $\mathrm{TD}(\rho,\sigma) \geq \max(\rho_{ii}-\sigma_{ii}, \sigma_{ii}-\rho_{ii}) = |\rho_{ii}-\sigma_{ii}|$ as desired. ◀

▶ **Fact 8.** *Let X, Y be random variables and f be a function. Then $d_{TV}(f(X), f(Y)) \leq d_{TV}(X, Y)$.*

▶ **Lemma 9** (Chernoff-Hoeffding Inequality)**.** *Let $X_1, X_2, \ldots, X_n$ be independent random variables, such that $0 \leq X_i \leq 1$ for all $i \in [n]$. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Then for any $\varepsilon > 0$,*

$$\Pr[|X-\mu| > \varepsilon)] \leq 2e^{-\frac{2\varepsilon^2}{n}}.$$

2.3.1 Chi-Squared Distributions

We present the definition and properties of the chi-squared distribution in the following.

▶ **Definition 10** (Chi-Squared Distribution)**.** *Let $Z_1, \ldots, Z_k$ be i.i.d. Gaussian random variables $\mathcal{N}(0,1)$. The random variable*

$$Q := \sum_{i \in [k]} Z_i^2.$$

is distributed according to the chi-squared distribution with k degrees of freedom*, denoted by $Q \sim \chi_k^2$.*

▶ **Fact 11.** *Let $Z \sim \mathcal{N}(0,1)$. Z^2 has a finite third moment.*

▶ **Fact 12.** *For all $k \in \mathbb{N}$, the following holds. Let $Q \sim \chi_k^2$. The mean of Q is k and the variance of Q is $2k$. Moreover, suppose $Q_1 \sim \chi_{k_1}^2$ and $Q_2 \sim \chi_{k_2}^2$, then $Q_1 + Q_2 \sim \chi_{k_1+k_2}^2$. When $k = 1$, we often omit the subscript and denote it by χ^2.*

We introduce a strong version of the central limit theorem that characterizes the *total variation distance* between the sum of i.i.d. *absolutely continuous*[6] random variables and Gaussian random variables. Note that most versions of central limit theorems state only the convergence in *cumulative density function*, which is not sufficient for our purpose.

▶ **Lemma 13** ([32, Theorem 1], restated)**.** *Let $X_1, \ldots, X_k$ be i.i.d. random variables. If X_1 is absolutely continuous and has a finite third moment, then*

$$d_{TV}\left(\frac{\sum_{i\in[k]}(X_i-\mu)}{\sqrt{k}\sigma}, Z\right) = O\left(\frac{1}{\sqrt{k}}\right),$$

where μ is the mean of X_1, σ is the standard deviation of X_1 and $Z \sim \mathcal{N}(0,1)$. Equivalently, $d_{TV}(\sum_{i\in[k]} X_i, Z') = O(1/\sqrt{k})$, where $Z' \sim \mathcal{N}(k\mu, k\sigma^2)$.

Since a random variable with a chi-squared distribution is the sum of squared i.i.d. Gaussian random variables, we have the following immediate corollary.

[6] A random variable X is absolutely continuous if there exists a (probability density) function $f : \mathbb{R} \to [0,1]$ such that $\Pr[X \leq x] = \int_{-\infty}^{x} f(t)\,dt$ for all $x \in \mathbb{R}$ and $\int_{-\infty}^{\infty} f(t)\,dt = 1$.

▶ **Corollary 14.** *Let Q be a random variable with a distribution χ^2_k. Then $d_{TV}(Q,Z) = O(1/\sqrt{k})$, where $Z \sim \mathcal{N}(k,2k)$.*

Proof. By definition, $Q = \sum_{i\in[k]} Z_i^2$ where $Z_i \sim \mathcal{N}(0,1)$. It immediately follows from the facts that Z_i^2 is absolutely continuous, $\mathbb{E}\left[Z_i^2\right] = 1$, $\mathsf{Var}\left(Z_i^2\right) = 2$, the third moment of Z_i^2 is finite and Lemma 13. ◀

2.3.2 Haar Measure

Given a d-dimensional Haar state, all coordinates of the state are correlated due to the unit-norm condition. The following theorem states that the *joint distribution* of $k = o(d)$ fraction of the coordinates in $\mathcal{S}(R^d)$ is statistically close to a random vector with i.i.d. Gaussian entries. The theorem was first proven in [10]. We will use the version stated in [24].

▶ **Theorem 15** ([24, Theorem 2.8]). *For every integer $d \geq 5$ and every $k \in \mathbb{N}$ that satisfies $1 \leq k \leq d-4$, let $X = (X_1,\ldots,X_d)$ be a uniform point on $\mathcal{S}(\mathbb{R}^d)$. Let Z be a random vector in $\mathbb{R}^k$ with i.i.d. Gaussian entries $\mathcal{N}(0,1/d)$. Then*

$$d_{TV}\left((X_1,\ldots,X_k),Z\right) \leq \frac{2(k+2)}{d-k-3}.$$

The above lemma can be extended to uniformly random vectors on $\mathcal{S}(\mathbb{C}^d)$. For a complex number α, we denote by $\mathsf{Re}(\alpha)$ and $\mathsf{Im}(\alpha)$, in order, the real part and imaginary part of α.

▶ **Lemma 16.** *Let $|\psi\rangle = \sum_{i\in[d]} \alpha_i |i\rangle$ be a uniform point on $\mathcal{S}(\mathbb{C}^d)$. Then $(\mathsf{Re}(\alpha_1),\ldots,\mathsf{Re}(\alpha_d), \mathsf{Im}(\alpha_1),\ldots,\mathsf{Im}(\alpha_d))$ is a uniform point on $\mathcal{S}(\mathbb{R}^{2d})$.*

Proof. First proposed by Muller [26], a uniform point on $\mathcal{S}(\mathbb{R}^{2d})$ can be sampled via the following procedures:

1. For $i \in [2d]$, sample $a_i \leftarrow \mathcal{N}(0,\sigma^2)$.
2. Output $\sum_{i\in[2d]} \frac{a_i}{\sqrt{\sum_{j\in[2d]} a_j^2}} |i\rangle$.

Where σ^2 in step 1 could be an arbitrary positive number.
On the other hand, a uniform point on $\mathcal{S}(\mathbb{C}^d)$ can be sampled as follows:

1. For $i \in [d]$, sample $\alpha_i \sim \mathbb{C}\mathcal{N}(0,1)$.
2. Output $\sum_{i\in[d]} \frac{\alpha_i}{\sqrt{\sum_{j\in[d]} |\alpha_j|^2}} |i\rangle$.

Particularly, in step 1, sampling $\alpha \sim \mathbb{C}\mathcal{N}(0,1)$ is equivalent to sampling $\alpha = a+ib$ according to $a \sim \mathcal{N}(0,1/2)$, $b \sim \mathcal{N}(0,1/2)$ by the definition of the complex normal distribution. Hence, picking $\sigma^2 = 1/2$ completes the proof. ◀

▶ **Corollary 17.** *For every integer $d \geq 3$ and every $k \in \mathbb{N}$ that satisfies $1 \leq 2k \leq 2d-4$, let $|\psi\rangle = \sum_{i\in[d]} \alpha_i |i\rangle$ be a random point on $\mathcal{S}(\mathbb{C}^d)$. Let Z be a random vector in $\mathbb{R}^{2k}$ with i.i.d. Gaussian entries $\mathcal{N}(0,1/(2d))$. Then*

$$d_{TV}\left((\mathsf{Re}(\alpha_1),\ldots,\mathsf{Re}(\alpha_k),\mathsf{Im}(\alpha_1),\ldots,\mathsf{Im}(\alpha_k)),Z\right) \leq \frac{2(2k+2)}{2d-2k-3}.$$

Proof. It immediately follows from Theorem 15 and Lemma 16. ◀

Next, the following simple fact gives an upper bound of the probability that a Gaussian random variable takes values near its mean.

▶ **Fact 18.** *Let $Z \sim \mathcal{N}(\mu, \sigma^2)$. For any $\Delta > 0$,*

$$\Pr[|Z - \mu| \le \Delta] \le \sqrt{\frac{2}{\pi}} \frac{\Delta}{\sigma}.$$

Proof. Let $f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$ be the probability density function of $\mathcal{N}(\mu, \sigma^2)$. The probability $\int_{\mu-\Delta}^{\mu+\Delta} f(x)\, dx$ can be upper-bounded by $f(\mu) \cdot 2\Delta = \sqrt{\frac{2}{\pi}} \frac{\Delta}{\sigma}$. ◀

2.4 Quantum State Tomography

▶ **Lemma 19** ([1, Corollary 7.6])**.** *There exists a tomography procedure* Tomography *that satisfies the following. For any error tolerance $\delta = \delta(\lambda) \in (0, 1]$ and any dimension $d = d(\lambda) \in \mathbb{N}$, given at least $t = t(\lambda) := 36\lambda d^3/\delta$ copies of a d-dimensional density matrix ρ,* $\mathsf{Tomography}(\rho^{\otimes t})$ *outputs a matrix $M \in \mathbb{C}^{d\times d}$ such that the following holds:*

$$\Pr\left[\|M - \rho\|_F^2 \le \delta : M \leftarrow \mathsf{Tomography}(\rho^{\otimes t})\right] \ge 1 - \mathsf{negl}(\lambda),$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Moreover, the running time of Tomography *is polynomial in $1/\delta, d$ and λ.*

By using the fact that $\|A\|_1 \le \sqrt{d}\|A\|_F$, we have the following immediate corollary.

▶ **Corollary 20.** *There exists a tomography procedure* Tomography *that satisfies the following. For any error tolerance $\delta = \delta(\lambda) \in (0, 1]$ and any dimension $d = d(\lambda) \in \mathbb{N}$, given at least $t = t(\lambda) := 144\lambda d^4/\delta^2$ copies of a d-dimensional density matrix ρ,* $\mathsf{Tomography}(\rho^{\otimes t})$ *outputs a matrix $M \in \mathbb{C}^{d\times d}$ such that the following holds:*

$$\Pr\left[\mathrm{TD}(M, \rho) \le \delta : M \leftarrow \mathsf{Tomography}(\rho^{\otimes t})\right] \ge 1 - \mathsf{negl}(\lambda).$$

Moreover, the running time of Tomography *is polynomial in $1/\delta, d$ and λ.*

3 Deterministically Extracting Classical Strings from Quantum States

In this section, we show how to pseudodeterministically extract classical strings from $O(\log(\lambda))$-qubit quantum states in *polynomial* time. We first present the outline of our construction.

1. Take as input $t(\lambda)$ copies of a $d(\lambda)$-dimensional (possibly mixed) quantum state ρ. Note that for our applications, we require $d(\lambda) = \mathsf{poly}(\lambda)$ and $t(\lambda) = \mathsf{poly}(\lambda)$.
2. Perform Tomography on the input $\rho^{\otimes t(\lambda)}$ to get an approximation $M \in \mathbb{C}^{d\times d}$ of its classical description.
3. Pick the first $k = o(d)$ diagonal entries of M, denoted by $p_1, \dots, p_k$. Divide them into ℓ groups where each of them is of size r (namely, $k = \ell \cdot r$).
4. In each group, consider the sum of all the elements. By q_i we denote the sum of the i-th group.
5. For each q_i, we round it to a bit, called b_i, according to which side it deviates from r/d.
6. Output the concatenation of every bit $b_1|| \dots ||b_\ell$.

In particular, we are interested in the case where the input is (polynomially many copies of) a *Haar state*. Informally, a Haar state can be thought of as a uniformly random point on a high-dimensional sphere. We can partition the sphere into many regions and assign each region a unique bitstring. Given the input quantum state, the goal of the extractor is to find the corresponding bitstring. Hence, Haar states can be viewed as a natural source of randomness. Below, we present our main theorem.

▶ **Theorem 21.** *There exists a quantum algorithm* Ext *such that for all $d(\cdot)$, there exists a (deterministic) function $f : \mathcal{D}(\mathbb{C}^{d(\lambda)}) \to \{0,1\}^{\ell(\lambda)}$ associated with* Ext*, where $\ell(\lambda) = d(\lambda)^{1/6}$. On input $t(\lambda) = \mathsf{poly}(d(\lambda), \lambda)$ copies of a $d(\lambda)$-dimensional density matrix ρ, the algorithm* Ext *outputs an $\ell(\lambda)$-bit string y and satisfies the following conditions.*

- **Efficiency:** *The running time is polynomial in d and λ.*
- **Correctness:** *For all $\lambda \in \mathbb{N}$, there exists a set $\mathcal{G}_\Delta \subseteq \mathcal{S}(\mathbb{C}^{d(\lambda)})$ such that*
  1. $\Pr\left[|\psi\rangle \in \mathcal{G}_\Delta : |\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})\right] \geq 1 - O(d(\lambda)^{-1/6})$.
  2. *For all $|\psi\rangle \in \mathcal{G}_\Delta$,*
  $$\Pr\left[y = f(|\psi\rangle\langle\psi|) : y \leftarrow \mathsf{Ext}\left(|\psi\rangle\langle\psi|^{\otimes t(\lambda)}\right)\right] \geq 1 - \mathsf{negl}(\lambda),$$
  where the probability is over the randomness of the extractor Ext.
- **Statistical Closeness to Uniformity:** *For sufficiently large $\lambda \in \mathbb{N}$,*
  $$d_{TV}(Y_\lambda, U_{\ell(\lambda)}) \leq O(d(\lambda)^{-1/6}),$$
  where $U_{\ell(\lambda)}$ is the uniform distribution over all $\ell(\lambda)$-bit strings and the random variable Y_λ is defined by the following process:
  $$|\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)}), Y_\lambda \leftarrow \mathsf{Ext}\left(|\psi\rangle\langle\psi|^{\otimes t(\lambda)}\right).$$

Proof. Here we present our construction of the extractor Ext.

▶ **Construction 22** (The Extractor Ext).

- *Input: $t(\lambda) := 144\lambda d(\lambda)^8$ copies of a $d(\lambda)$-dimensional quantum state $\rho \in \mathcal{D}(\mathbb{C}^{d(\lambda)})$.*
- *Perform* $\mathsf{Tomography}(\rho^{\otimes t(\lambda)})$ *with error tolerance $\delta(\lambda) := d(\lambda)^{-5/3}$ to get the classical description $M \in \mathbb{C}^{d(\lambda)\times d(\lambda)}$ that approximates ρ.*
- *Run* $\mathsf{Round}(M)$ *to get $y \in \{0,1\}^{\ell(\lambda)}$.*
- *Output y.*

The classical post-processing procedure $\mathsf{Round}(M)$ is defined as follows:

$\mathsf{Round}(M)$:

- Input: a matrix $M \in \mathbb{C}^{d(\lambda)\times d(\lambda)}$.
- Set parameters $k(\lambda) := d(\lambda)^{5/6}$, $r(\lambda) := d(\lambda)^{2/3}$ and $\ell(\lambda) := d(\lambda)^{1/6}$.
- Let $p_1, \ldots, p_{d(\lambda)}$ be the diagnal entries of M. For $i \in \{1, \ldots, \ell(\lambda)\}$, let
  $$q_i := \sum_{j=1}^{r} p_{(i-1)r+j}.$$
- For $i \in \{1, \ldots, \ell(\lambda)\}$, define
  $$b_i = \begin{cases} 0, & \text{if } q_i < r/d \\ 1, & \text{if } q_i > r/d. \end{cases}$$
- Output $b_1||\ldots||b_{\ell(\lambda)}$.

By Corollary 20 and the fact that $t(\lambda) = \mathsf{poly}(d, \lambda)$ and $\delta(\lambda) = 1/\mathsf{poly}(d)$, it is easy to see that the running time of the extractor Ext is polynomial in d and λ. Before proving the correctness and statistical closeness to uniformity, we present several statistical properties.

First, the distribution of the real and imaginary parts of any $k = o(d)$ coordinates of a Haar state $|\psi\rangle \sim \mathscr{H}(\mathbb{C}^d)$ is statistically close to a random vector with i.i.d. Gaussian entries.

▷ Claim 23. Let $|\psi\rangle = \sum_{i=1}^{d} \alpha_i |i\rangle$ be a uniformly random point on $\mathcal{S}(\mathbb{C}^d)$. Then

$$d_{TV}\left((\mathsf{Re}(\alpha_1), \mathsf{Im}(\alpha_1), \ldots, \mathsf{Re}(\alpha_k), \mathsf{Im}(\alpha_k)), Z\right) = O\left(k/d\right),$$

where Z is a random variable in $\mathbb{R}^{2k}$ with i.i.d. Gaussian entries $\mathcal{N}(0, 1/(2d))$.

Proof. It immediately follows from Corollary 17. ◁

Next, since the i-th diagonal entry p_i of $|\psi\rangle\langle\psi|$ is the squared absolute value of the i-th coordinate α_i of $|\psi\rangle$, the joint distribution of $(p_1, \ldots, p_k)$ is statistically close to a random vector in $\mathbb{R}^k$ with i.i.d. χ_2^2 entries.

▷ Claim 24. $d_{TV}\left((p_1, \ldots, p_k), Q/(2d)\right) = O(k/d)$ where Q is a random variable in $\mathbb{R}^k$ with i.i.d. χ_2^2 entries.

Proof. For $i \in [k]$, each diagonal entry $p_i = |\alpha_i|^2 = \mathsf{Re}(\alpha_i)^2 + \mathsf{Im}(\alpha_i)^2$. By Claim 23, the total variation distance induced by replacing the real and imaginary parts of the amplitudes with i.i.d. Gaussians $\mathcal{N}(0, 1/(2d))$ is $O(k/d)$. Then by setting $f(x_1, \ldots, x_{2k}) := \left(x_1^2 + x_2^2, \ldots, x_{2k-1}^2 + x_{2k}^2\right)$ in Fact 8 and the definition of χ_2^2, we complete the proof. ◁

Now, we consider the distribution of q_i's. Note that the sum of r i.i.d. χ_2^2 random variables is identically distributed to χ_{2r}^2 by the property of the χ^2-distribution in Fact 12. Namely, the joint distribution of $(q_1, \ldots, q_\ell)$ is statistically close to a random vector in $\mathbb{R}^\ell$ with i.i.d. χ_{2r}^2 entries.

▷ Claim 25. $d_{TV}\left((q_1, \ldots, q_\ell), R/(2d)\right) = O(k/d)$ where R is a random variable in $\mathbb{R}^\ell$ with i.i.d. χ_{2r}^2 entries.

Proof. Recall that $q_i := \sum_{j=1}^{r} p_{(i-1)r+j}$. From Claim 24, we have $d_{TV}\left((p_1, \ldots, p_k), Q/(2d)\right) = O(k/d)$, where $Q = (Q_1, \ldots, Q_k)$ is a random variable in $\mathbb{R}^k$ with i.i.d. χ_2^2 entries. Hence by the data processing inequality (Fact 8) and setting

$$f(x_1, \ldots, x_k) := \left(\sum_{j=1}^{r} x_j, \sum_{j=1}^{r} x_{r+j}, \ldots, \sum_{j=1}^{r} x_{(\ell-1)r+j}\right),$$

we have $d_{TV}\left((q_1, \ldots, q_\ell), R/(2d)\right) = O(k/d)$, where we use the fact that the sum of r i.i.d. χ_2^2 random variables is identically distributed to χ_{2r}^2. ◁

Moreover, a χ_{2r}^2 random variable is the sum of r i.i.d. absolutely continuous random variables. Hence, relying on the aforementioned central limit theorem, it is statistically close to a Gaussian distribution.

▶ **Lemma 26.** *$d_{TV}\left((q_1, \ldots, q_\ell), Z/(2d)\right) = O(k/d) + O(\ell/\sqrt{r})$ where Z is a random variable in $\mathbb{R}^\ell$ with i.i.d. $\mathcal{N}(2r, 4r)$ entries, i.e., $Z/(2d)$ has i.i.d. $\mathcal{N}(r/d, r/d^2)$ entries.*

Proof. By Corollary 14 and hybrids over every coordinate for $i \in [\ell]$, we have $d_{TV}(R/(2d), Z/(2d)) = O(\ell/\sqrt{r})$, where R is defined in Claim 25. Together with Claim 25 finishes the proof. ◀

Now, we are ready to prove the correctness and the statistical closeness to uniform properties.

Correctness. First, define the function $f : \mathcal{D}(\mathbb{C}^{d(\lambda)}) \to \{0,1\}^{\ell(\lambda)}$ associated with the extractor as

$$f(\sigma) := \mathsf{Round}(\sigma).$$

Due to the continuous nature of quantum states, it is impossible to discretize them perfectly. For any $\sigma \in \mathcal{D}(\mathbb{C}^d)$, consider the corresponding $q_1, \dots, q_\ell$ defined in Construction 22. If *all* $q_1, \dots, q_\ell$ are sufficiently away from r/d, then the extractor is able to output the correct string with high probability by the correctness of Tomography. Here, we define the set $\mathcal{G}_\Delta$ of "good states" whose $q_1, \dots, q_\ell$ are *all* Δ-away from r/d (the parameter $\Delta(\lambda)$ will be chosen later). The following claim characterizes the probability of a Haar random state being in $\mathcal{G}_\Delta$.

▷ Claim 27. Let the set $\mathcal{G}_\Delta \subseteq \mathcal{S}(\mathbb{C}^{d(\lambda)})$ be

$$\mathcal{G}_\Delta := \left\{ \ | \ \psi\rangle \in \mathcal{S}(\mathbb{C}^{d(\lambda)}) : \forall i \in [\ell], \ \left|q_i - \frac{r}{d}\right| > \Delta \right\},$$

where each q_i is defined on the matrix $|\psi\rangle\langle\psi|$. It holds that

$$\Pr\left[|\psi\rangle \in \mathcal{G}_\Delta : |\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})\right] \geq 1 - O\left(\frac{k}{d}\right) - O\left(\frac{\ell}{\sqrt{r}}\right) - O\left(\frac{\Delta \ell d}{\sqrt{r}}\right).$$

Proof. By Lemma 26, the total variation distance between $(q_1, \dots, q_\ell)$ and the random variable $Z = (Z_1, \dots, Z_\ell)$ with i.i.d. Gaussian entries $Z_i \sim \mathcal{N}(r/d, r/d^2)$ is $O(k/d) + O(\ell/\sqrt{r})$. Hence,

$$\begin{aligned}\Pr\left[|\psi\rangle \in \mathcal{G}_\Delta : |\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})\right] &= \Pr\left[\forall i \in [\ell], \ \left|q_i - \frac{r}{d}\right| > \Delta\right] \\ &\geq \Pr\left[\forall i \in [\ell], \ \left|Z_i - \frac{r}{d}\right| > \Delta\right] - O\left(\frac{k}{d}\right) - O\left(\frac{\ell}{\sqrt{r}}\right).\end{aligned}$$

Moreover, by Fact 18, for every coordinate $i \in [\ell]$, it holds that

$$\Pr\left[\left|Z_i - \frac{r}{d}\right| \leq \Delta\right] \leq O\left(\frac{\Delta}{\sqrt{r/d}}\right).$$

By a union bound over $i \in [\ell]$, with all but $O\left(\frac{\Delta \ell d}{\sqrt{r}}\right)$ probability every Z_i is Δ-away from r/d. Collecting the probabilities completes the proof of Claim 27. ◁

Hence, by setting $\Delta(\lambda) = 1/d(\lambda)$, the choice of parameters $r(\lambda) = d(\lambda)^{2/3}$, $\ell(\lambda) = d(\lambda)^{1/6}$, $k(\lambda) = d(\lambda)^{5/6}$ and Claim 27, we have

$$\Pr\left[|\psi\rangle \in \mathcal{G}_\Delta : |\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})\right] \geq 1 - O(d(\lambda)^{-1/6}).$$

Next, given a state which is in $\mathcal{G}_\Delta$, the output bitstring extracted from it will be $f(|\psi\rangle\langle\psi|)$ with overwhelming probability by the correctness of Tomography in Corollary 20.

▷ Claim 28. If $|\psi\rangle \in \mathcal{G}_\Delta$, then running Ext in Construction 22 with error tolerance $\delta(\lambda) = d^{-5/3} = \Delta(\lambda)/r(\lambda)$ for Tomography satisfies

$$\Pr\left[y = f(|\psi\rangle\langle\psi|) : y \leftarrow \mathsf{Ext}\left(|\psi\rangle\langle\psi|^{\otimes t(\lambda)}\right)\right] \geq 1 - \mathsf{negl}(\lambda),$$

where the probability is over the randomness of the extractor Ext.

Proof. Let M be the classical description obtained by running $\mathsf{Tomography}(|\psi\rangle\langle\psi|^{\otimes t(\lambda)})$ with error tolerance δ and $t \geq 144\lambda d^8 \geq 144\lambda d^4/\delta^2$. Let $\hat{p}_i$'s and $\hat{q}_j$'s be the corresponding diagonal entries and sums of M. By Corollary 20 , $\mathrm{TD}(|\psi\rangle\langle\psi|, M) \leq \delta$ holds with overwhelming probability. For the rest of the proof, we assume that this event holds. Then by Fact 7, we have $|p_i - \hat{p}_i| \leq \delta$ for every $i \in [k]$. Since $|\psi\rangle \in \mathcal{G}_\Delta$, we have $|q_i - r/d| > \Delta$ for every $i \in [\ell]$. We now show that if $q_i > r/d + \Delta$, then $\hat{q}_i > r/d$. For every $i \in [\ell]$, by the triangle inequality and the fact that $\delta = \Delta/r$, we have

$$\hat{q}_i = q_i - (q_i - \hat{q}_i) > \left(\frac{r}{d} + \Delta\right) - \sum_{j=1}^{r} \left|p_{(i-1)r+j} - \hat{p}_{(i-1)r+j}\right| \geq \frac{r}{d} + \Delta - r \cdot \delta = \frac{r}{d}.$$

Similarly, we have $q_i < r/d - \Delta$ implies that $\hat{q}_i < r/d$. Hence, this ensures the consistency between $\mathsf{Round}(M)$ and $\mathsf{Round}(|\psi\rangle\langle\psi|)$ and completes the proof of Claim 28. ◁

Statistical Closeness to Uniformity. We finish the proof with a hybrid argument:

- H_1 : In the first hybrid, the output is generated according to Construction 22.
  1. Sample $|\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})$.
  2. Perform $\mathsf{Tomography}(\rho^{\otimes t(\lambda)})$ with $t(\lambda) := 144\lambda d(\lambda)^8$ and error tolerance $\delta(\lambda) := d(\lambda)^{-5/3}$ to get the classical description $M \in \mathbb{C}^{d(\lambda)\times d(\lambda)}$ that approximates ρ.
  3. Output $y = \mathsf{Round}(M)$.
- H_2 : In the second hybrid, the input of Round is changed to the exact description of the quantum state.
  1. Sample $|\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^{d(\lambda)})$.
  2. Output $y = \mathsf{Round}(|\psi\rangle\langle\psi|)$.
- H_3 : In the third hybrid, the output is generated by rounding i.i.d. Gaussians.
  1. Sample $z_1, \ldots, z_\ell \leftarrow \mathcal{N}(r/d, r/d^2)$.
  2. For $i \in [\ell]$,

$$b_i = \begin{cases} 0, & \text{if } z_i < r/d \\ 1, & \text{if } z_i > r/d. \end{cases}$$

  3. Output $b_1||\ldots||b_{\ell(\lambda)}$.

We can bound the total variation distance between H_1 and H_2 by $O(\delta) = O(d^{-5/3})$ using Corollary 20 and our chosen error tolerance. Additionally, the total variation distance between H_2 and H_3 is at most $O(d^{-1/6})$ from Lemma 26. Finally, since Gaussians are symmetric about the mean, the output string in H_3 is uniformly and randomly distributed. This finishes the proof of Theorem 21. ◀

4 Quantum PRGs and PRFs

In this section, we present our main application of the extractor in Section 3. We introduce the notion of *pseudodeterministic quantum pseudorandom generators* (QPRGs). As the name suggests, QPRGs is a pseudorandom generator with quantum generation satisfying only *pseudodeterminism* property. To be more precise, by pseudodeterminism we mean that there exist some constant $c > 0$ and at least $1 - O(\lambda^{-c})$ fraction of "good seeds" for which the output is almost certain. That is, for each good seed, the probability (over the randomness of the QPRG) of the most likely output is at least $1 - O(\lambda^{-c})$. Due to space limitations, we postpone the details to the full version.

5 Applications

In this section, we present applications based on sQPRGs and selectively secure QPRFs introduced in Section 4. One key advantage of using sQPRGs or selectively secure QPRFs as the starting point is that we can build higher-level primitives simply by following the classical construction with a slight modification, and then security will follow from the same reasoning. However, we must address the issue of correctness since sQPRGs and selectively secure QPRFs have only $1 - O(\lambda^{-c})$ pseudodeterminism. To resolve the issue, we apply a simple parallel repetition followed by a majority vote to boost correctness at the expense of increased communication complexity and key length. Due to space limitations, we postpone the details to the full version.

References

1 Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In *Theory of Cryptography Conference*, pages 237–265. Springer, 2022.

2 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *CRYPTO*, 2022.

3 Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *computational complexity*, 15(2):115–162, 2006.

4 James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 467–496. Springer, 2021. doi:10.1007/978-3-030-84242-0_17.

5 Amit Behera, Zvika Brakerski, Or Sattath, and Omri Shmueli. Pseudorandomness with proof of destruction and applications. Cryptology ePrint Archive, Paper 2023/543, 2023. URL: https://eprint.iacr.org/2023/543.

6 Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, and Zixin Zhou. Quantum pseudoentanglement. *arXiv preprint*, 2022. arXiv:2211.00747.

7 Adam Bouland, Bill Fefferman, and Umesh V. 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, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 63:1–63:2. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ITCS.2020.63.

8 Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 417–440. Springer, 2020. doi:10.1007/978-3-030-56880-1_15.

9 Zvika Brakerski and Henry Yuen. Quantum garbled circuits. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 804–817, 2022.

10 Persi Diaconis and David Freedman. A dozen de Finetti-style results in search of a theory. *Annales de l'I.H.P. Probabilités et statistiques*, 23(S2):397–423, 1987. URL: http://www.numdam.org/item/AIHPB_1987__23_S2_397_0/.

11 Yevgeniy Dodis, Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Security amplification for interactive cryptographic primitives. In *Theory of Cryptography: 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings 6*, pages 128–145. Springer, 2009.

12 Dmitry Gavinsky. Quantum money with classical verification. In *2012 IEEE 27th Conference on Computational Complexity*, pages 42–52. IEEE, 2012.

13 Oded Goldreich. A note on computational indistinguishability. *Information Processing Letters*, 34(6):277–281, 1990. doi:10.1016/0020-0190(90)90010-U.

14 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.

15 Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in miniqcrypt. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 531–561. Springer, 2021. doi:10.1007/978-3-030-77886-6_18.

16 Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1579–1588, 2023.

17 Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.

18 Russell Impagliazzo and Avi Wigderson. P= bpp if e requires exponential circuits: Derandomizing the xor lemma. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 220–229, 1997.

19 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018. doi:10.1007/978-3-319-96878-0_5.

20 William Kretschmer. Quantum pseudorandomness and classical complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, July 5-8, 2021, Virtual Conference*, volume 197 of *LIPIcs*, pages 2:1–2:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.TQC.2021.2.

21 Yanyi Liu and Rafael Pass. On one-way functions and kolmogorov complexity. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1243–1254. IEEE, 2020.

22 Ueli Maurer and Stefano Tessaro. Computational indistinguishability amplification: Tight product theorems for system composition. In Shai Halevi, editor, *Advances in Cryptology — CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 350–368. Springer-Verlag, August 2009.

23 Ueli Maurer and Stefano Tessaro. A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak prgs with optimal stretch. In Daniele Micciancio, editor, *Theory of Cryptography — TCC 2010*, volume 5978 of *Lecture Notes in Computer Science*, pages 237–254. Springer-Verlag, February 2010.

24 Elizabeth S Meckes. *The random matrix theory of the classical compact groups*, volume 218. Cambridge University Press, 2019.

25 Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions, 2021. doi:10.48550/ARXIV.2112.06369.

26 Mervin E. Muller. A note on a method for generating points uniformly on n-dimensional spheres. *Commun. ACM*, 2(4):19–20, April 1959. doi:10.1145/377939.377946.

27 Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991. doi:10.1007/BF00196774.

28 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. doi:10.1017/CBO9780511976667.

29 Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of computer and System Sciences*, 49(2):149–167, 1994.

30 Ryan ODonnell and David Witmer. Goldreich's prg: evidence for near-optimal polynomial stretch. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 1–12. IEEE, 2014.

31 Alexander A Razborov and Steven Rudich. Natural proofs. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 204–213, 1994.

32 S Kh Sirazhdinov and M Mamatov. On convergence in the mean for densities. *Theory of Probability & Its Applications*, 7(4):424–428, 1962.