



Commitments from Quantum One-Wayness

Dakshita Khurana

University of Illinois
Urbana-Champaign, USA
dakshita@illinois.edu

Kabir Tomer

University of Illinois
Urbana-Champaign, USA
ktomer2@illinois.edu

ABSTRACT

One-way functions are central to classical cryptography. They are necessary for the existence of non-trivial classical cryptosystems, and also sufficient to realize meaningful primitives including commitments, pseudorandom generators and digital signatures. At the same time, a mounting body of evidence suggests that assumptions *even weaker* than one-way functions may suffice for many cryptographic tasks of interest in a quantum world, including bit commitments and secure multi-party computation.

This work studies one-way state generators [Morimae-Yamakawa, CRYPTO 2022], a natural quantum relaxation of one-way functions. Given a secret key, a one-way state generator outputs a hard to invert quantum state. A fundamental question is whether this type of *quantum* one-wayness suffices to realize quantum cryptography. We obtain an affirmative answer to this question, by proving that one-way state generators with pure state outputs imply quantum bit commitments and secure multiparty computation.

Along the way, we use efficient shadow tomography [Huang et. al., Nature Physics 2020] to build an intermediate primitive with classical outputs, which we call a (quantum) one-way puzzle. Our main technical contribution is a proof that one-way puzzles imply quantum bit commitments. This proof develops new techniques for pseudoentropy generation [Hastad et. al., SICOMP 1999] from arbitrary distributions, which may be of independent interest.

CCS CONCEPTS

• Theory of computation → Cryptographic primitives.

KEYWORDS

Quantum cryptography, Commitments, One-way states, EFI, One-way functions, Pseudo-entropy

ACM Reference Format:

Dakshita Khurana and Kabir Tomer. 2024. Commitments from Quantum One-Wayness. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, June 24–28, 2024, Vancouver, BC, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3618260.3649654>

1 INTRODUCTION

A one-way function is a classically efficiently computable function that is hard to invert. This is a fundamental hardness assumption,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649654>

necessary for the existence of much of modern classical cryptography [20, 21, 29]. The classical crypto-complexity class “minicrypt” contains primitives like bit commitments, pseudorandom generators, pseudorandom functions and symmetric encryption, that are all equivalent to the existence of one-way functions. On the other hand, tasks like key exchange and secure multi-party computation classically require stronger, more structured assumptions [22].

The relationship between computational hardness and cryptography appears to be drastically different in a quantum world. Here, the seminal works of Wiesner [36] and Bennett and Brassard [9] first demonstrated the possibility of unconditional quantum key distribution (QKD) by exploiting the properties of quantum information. Unfortunately, it was also shown that other useful cryptographic primitives like bit commitments and secure computation *cannot* exist unconditionally [27, 30], and must necessarily rely on computational hardness, even in a quantum world. However, our understanding of computational hardness in a quantum world is still in its infancy. For instance, it was only recently understood [8, 15] that one-way functions suffice to enable secure multi-party computation in a quantum world, a task that is believed to be impossible classically.

Sources of Hardness in a Quantum World. Despite being necessary for classical cryptography, one-way functions may not be necessary for computational quantum cryptography.

Two recent concurrent works [6, 32] demonstrated that many cryptographic primitives including quantum bit commitments, (one-time secure) digital signatures, and multi-party secure computation can also be based on the existence of *pseudorandom state generators* (PRSGs), which were introduced in [23].

Given a secret key, a PRSG efficiently generates a quantum state, several copies of which are computationally indistinguishable from equally many copies of a Haar random state. There is some evidence that points to PRSGs being a weaker assumption than one-way functions. Specifically, PRSGs can exist even if $BQP = QMA$ (relative to a quantum oracle) [25] or if $P = NP$ (relative to a classical oracle) [26]. This indicates that PRSGs, and all the cryptographic primitives that they imply, can exist even if all quantum-secure (classical) cryptographic primitives, including one-way functions, are broken.

Can we base quantum cryptography on assumptions that are potentially even weaker than the existence of PRSGs? As pointed out in [31], PRSGs and bit commitments are “decision-type” primitives that rely on the hardness of distinguishing pseudorandom states from truly (Haar) random ones. On the other hand, there is a natural, simpler “search-type” assumption that significantly relaxes the pseudorandomness guarantee of a PRSG to one-wayness.

One-Way State Generators (OWSG). A OWSG [32] is an efficient quantum algorithm that given a secret key, generates a hard-to-invert quantum state.

Given that one-way functions enable a variety of classical cryptosystems, it is natural to ask whether one-way state generators play a similar role in quantum cryptography. Namely,

Can we obtain quantum cryptosystems including bit commitments and MPC, only assuming the existence of one-way state generators?

Our main theorem answers this question in the affirmative in the setting where OWSG outputs are pure states.

THEOREM 1.1. (Informal) *One-way state generators with pure state outputs imply quantum bit commitments.*

By combining with prior work that demonstrates conversions between various types of commitments [6, 10] and builds secure multi-party computation from commitments [6, 8, 15], we also obtain the following corollary.

COROLLARY 1.1. (Informal) *One-way state generators with pure state outputs imply secure multi-party computation for all quantum functionalities.*

We note that OWSGs were initially defined in [32] to only output pure states; but this definition was later generalized in [31] to also allow mixed states. OWSG with pure state outputs were also studied in [12], who also showed equivalences between variants (weak, distributional) of OWSGs. Outputs of random quantum circuits yield natural candidates for pure OWSG that do not rely on classical hardness; in fact the output states can even be conjectured to be pseudorandom [6]. However, only relying on one-wayness introduces the possibility of building cryptography from other natural candidates: for instance, the (pre-measurement) states generated by BosonSampling experiments are not indistinguishable from Haar random [3], but can plausibly be one-way.

Pure OWSGs are also implied by various cryptographic primitives such as digital signatures with pure verification keys and quantum money with pure banknotes [31]. This, combined with our theorem, shows that these other primitives also imply quantum bit commitments. In some sense, this establishes commitments as the leading candidate for a minimal/necessary assumption in quantum cryptography.

One-Way Puzzles. Enroute to our main theorem, we use efficient shadow tomography [19] to prove that OWSG imply an intermediate cryptographic primitive with entirely classical outputs, that we call a *one-way puzzle*. We find this implication from a OWSG with quantum outputs to a simple, cryptographically useful primitive with classical outputs, noteworthy.

THEOREM 1.2. (Informal) *One-way state generators with pure state outputs imply one-way puzzles.*

A (quantum) one-way puzzle consists of a pair (Samp, Ver) where Samp is a quantum polynomial time algorithm and Ver is Boolean function¹. Samp outputs a pair of classical strings – a key and puzzle (k, s) – satisfying $\text{Ver}(k, s) = 1$. The security guarantee is that given a “puzzle” s , it is (quantum) computationally infeasible to find a key k such that $\text{Ver}(k, s) = 1$, except with negligible probability.

¹In an earlier version, we defined Ver to be an unbounded time algorithm. However, Ver does not even need to be computable, so the definition may be relaxed to allow for arbitrary functions. This was independently noted in [13].

Unlike prior definitions of one-way puzzles in the literature, we do not require verification (Ver) to be efficiently implementable. As we will see later, only asking for inefficient verification turns out to be *necessary* for our implication from OWSG. Indeed, if verification were efficient, then a QMA oracle would be capable of breaking one-way puzzles, but such an oracle is unlikely to break OWSG [25]. Somewhat surprisingly, we show that inefficiently verifiable one-way puzzles are also *sufficient* to build quantum bit commitments.

The reader may have observed that one-way puzzles generalize one-way functions to allow joint, randomized sampling of keys and outputs. In a classical world, this generalization is unnecessary: one-way puzzles are equivalent to one-way functions. One direction of the implication is straightforward, since one-way functions imply one-way puzzles (almost) immediately by definition. In the other direction, a one-way function can be obtained from a classical one-way puzzle by “pulling out” the (uniform) randomness r used by Samp. The one-way function f on input r samples $(k, s) \leftarrow \text{Samp}(1^n, r)$ and outputs $f(r) = s$. It is easy to see that one-wayness of the puzzle implies one-wayness of f .

The conversion above is no longer applicable when Samp is quantum, because there may be no equivalent deterministic, efficient function that on input uniform randomness, outputs (k, s) distributed according to the output of Samp. Nevertheless, enroute to proving our main result, we show:

THEOREM 1.3. (Informal) *One-way puzzles imply quantum bit commitments.*

Theorem 1.3 is the most technically involved part of this work. In a nutshell, existing techniques for building commitments from classical one-way primitives (e.g., [17]) crucially only apply when the preimage distribution of every image of the function is flat (i.e., uniform over all preimages). This work develops a method to generate pseudorandomness from one-way puzzles with arbitrary preimage distributions, which we believe to be of independent interest.

Local/Hybrid Quantum Cryptography and One-Way Puzzles. As an aside, we observe that one-way puzzles are also implied by quantum cryptography with classical communication. In fact there is a large body of work that aims to understand the computational hardness yielding quantum cryptography with classical communication, including protocols for quantum advantage [11, 33, 34], quantum commitments with classical communication [5], and even black-box separations for key exchange [7]. Classical communication protocols are desirable as they can be used over the current infrastructure (e.g., the Internet). In this model, sometimes called the “local” or “hybrid” or quantum-computation classical-communication (QCCC) model [7], all the quantum computation is done locally by parties who exchange only classical messages.

We observe that natural cryptographic primitives such as public-key encryption and signatures in the QCCC model imply one-way puzzles. For example, given a public-key encryption scheme, a one-way puzzle can be defined as follows. The one-way puzzle sampler will output a puzzle consisting of a public key along with an encryption of a random message, and the corresponding solution will be the (plaintext) message. It is easy to see that an adversary that breaks one-wayness of the resulting puzzle can be used to break

CPA security of the encryption scheme. In fact, one can obtain a one-way puzzle even given any public key encryption with classical public and secret keys, but quantum ciphertexts. We formalize these ideas and also show how similar ideas prove that one-way puzzles are implied by digital signatures, natural bit commitments and symmetric encryption schemes in the QCCC model².

Finally, we note that [24] recently discussed a related but stronger primitive – *hard quantum planted problems for NP languages* – which is implied by cryptography with publicly verifiable deletion. A hard quantum planted problem for a language is specified by a QPT sampler that samples an instance-witness pair (x, w) for the language in a way that no adversary can find a witness for x with non-negligible probability. These are like one-way puzzles except that they admit efficient verification. By definition, hard quantum planted problems imply one-way puzzles (and therefore by our work, imply quantum bit commitments).

Conclusion and Future Directions. Prior to this work, bit commitments were known to be implied by pseudorandom state generators [6, 32] via a construction that roughly parallels the classical setting [35]. They were also known [31] from a restricted type of OWSG; namely one with injective, orthogonal outputs. However, as we discuss in the next section, building commitments from general-purpose OWSG requires methods that are quite different from known classical techniques, and which may be broadly applicable beyond this work.

We also hope that the one-way puzzle abstraction will enable a better understanding of quantum bit commitments. For example, some existing attempts to understand the complexity of quantum commitments [25] build oracles relative to which complexity classes collapse, but pseudorandom states exist (and thus, one-way puzzles exist). Directly establishing the existence of one-way puzzles relative to these oracles may be easier, and may enable even more general oracle separations. One-way puzzles may also help better understand the relationship between quantum cryptography and quantum notions of Kolmogorov complexity.

Finally, we discuss some open questions related to this work. An obvious one is whether our results extend to *mixed state* OWSG. One avenue towards proving this would be to build one-way puzzles from mixed-state OWSG, perhaps via better tomography. In addition, answering the following questions will shed some more light on the complexity of quantum cryptography.

- (1) Can quantum bit commitments with *classical communication* be based on the existence of OWSG or one-way puzzles? This is plausible because one-way puzzle outputs are classical after all. Moreover, many other intermediate primitives that we build in this work also have entirely classical outputs.
- (2) En route to building commitments, this work constructs pseudo-*entropy* generators from one-way state generators. Can other pseudorandom primitives, such as pseudorandom quantum states be obtained from OWSG or one-way puzzles? Techniques in this work may serve as a useful starting point towards addressing this question.

- (3) Do quantum bit commitments imply one-way puzzles? If not, is there a separation? It is easy to observe that one-way puzzles can be broken given (quantum) access to an oracle for a related boolean function f . Is this also true for every quantum bit commitment? This question appears to be connected with the unitary synthesis problem [1, 4], for which a recent work [28] gave a general one-query lower bound.
- (4) Is there a quantum analogue to the classical implication from one-way puzzles to one-way functions? In other words, does the existence of one-way puzzles with hardness over arbitrary distributions imply one-way primitives with hardness over uniform/flat preimage distributions?

2 TECHNICAL OVERVIEW

We begin this overview by outlining a well-known construction of classical commitments from any injective one-way function. This construction relies on hardcore predicates: roughly, a hardcore predicate for a one-way function f is a bit that is easy to compute given a preimage k but hard to compute given $f(k)$. The Goldreich-Levin theorem [14] shows that the bit $\langle k, r \rangle$ is hardcore for the function $f(k) || r$. When f is injective, the hardcore bit is uniquely determined for every element in the image, and gives rise to a simple commitment scheme, as follows.

A commitment to bit b is $f(k), r, \langle k, r \rangle \oplus b$ for randomly sampled k and r . This commitment is binding because of the injectivity of f , and computationally hides the bit b due to $\langle k, r \rangle$ being hardcore. This construction does not work when f is not injective. In this case, for an image y , there may exist two preimages $k_1, k_2 \in \{f^{-1}(y)\}$ such that $\langle k_1, r \rangle \neq \langle k_2, r \rangle$, which will allow the committer to break binding.

The celebrated work of Hastad et. al. [17] showed how to overcome the binding issue, and base *classical* commitments on *general* (not necessarily injective) one-way functions. We outline (some relevant parts of) their technique next.

Pairwise Independent Hashing Reduces the Number of Preimages. The starting point of the technique in [17] is to append to the image $f(k)$ a *pairwise-independent* hash $h(k)$, thereby reducing the total number of preimages of $f(k), h(k)$. This makes $f(k), h(k)$ behave somewhat like an injective function for carefully chosen output sizes of $h(k)$.

In more detail, let N_k denote the number of preimages of $f(k)$. When the output size $h(k)$ is set to (slightly larger than) $\log N_k$, then [17] (roughly) show that:

- $h(k)$ is computationally indistinguishable from uniform given $f(k)$, and
- $h(k)$ is statistically (somewhat) distinguishable from uniform given $f(k)$.

The fact that $f(k), h(k)$ *appears* to a computationally bounded adversary to have more entropy than it actually does is formalized by building an object called a *weak pseudoentropy generator* (WPEG) [16, 17]. We will now describe this object in more detail.

²One may ask whether computational cryptographic primitives in the QCCC model also imply one-way functions. But it is unclear if this is true; and at the very least this is challenging to prove, for the same reason as above – namely, we cannot explicitly pull out the sampling randomness from an arbitrary quantum algorithm.

2.1 Weak Pseudoentropy Generators (WPEG)

A distribution G_0 is a weak pseudoentropy generator (WPEG) if there exists another, possibly inefficient *simulated* distribution G_1 whose output is computationally indistinguishable from, and yet has more Shannon entropy than G_0 .

For a one-way function f and pairwise independent hash h , we can consider distributions

$$G_0(1^n) := f(k), h, i, h(k)_i \text{ and}$$

$$G_1(1^n) := \begin{cases} f(k), h, i, h(k)_{i-1}, u_1 & \text{if } i = \lceil \log N_k \rceil + 1 \\ f(k), h, i, h(k)_i & \text{otherwise} \end{cases}$$

where k and the hash key h are sampled uniformly in $\{0, 1\}^n$, $i \leftarrow [n]$, N_k denotes the number of preimages of $f(k)$, $h(k)_i$ denotes $h(k)$ truncated to the first i bits, and u_1 denotes a uniformly random bit.

Prior works [16, 17] show that the distributions G_0 and G_1 are computationally indistinguishable, but G_1 has more entropy than G_0 . This can be understood as follows.

- **Entropy Gap.** Roughly, the pairwise independence of h implies that with probability at least $\frac{1}{2}$, $f(k), h(k)_{\lceil \log N_k \rceil}$ has a single preimage, i.e., k . Thus, with probability at least $\frac{1}{2}$ the last bit in G_0 is a deterministic function of the remaining bits, and has less entropy than the corresponding (uniform) bit in G_1 .
- **Computational Indistinguishability.** By the Leftover Hash Lemma, for x sampled from any distribution X with min-entropy ℓ , the first $\ell - 2c \log n$ bits of $h(x)$ are $\frac{1}{n^c}$ statistically close to uniform, even given h . By setting X to be the (uniform) distribution over preimages of $f(k)$, this implies that the first $\ell - O(\log N_k)$ bits of $h(k)$ are statistically close to uniform given h , for $\ell = \log N_k$. Then applying the Goldreich-Levin theorem while guessing the last $O(\log n)$ bits of $h(k)$ converts a distinguisher between G_0 and G_1 to an inverter for f^3 .

Next, we discuss barriers in extending these ideas to quantum one-way state generators.

A Preliminary Approach that Does Not Work. A natural first approach to building commitments from OWSG could be to replace the classical string $f(k)$ in the distributions above, with the quantum state $|\psi_k\rangle$ output by the OWSG.

Then the two WPEG distributions G_0 and G_1 are replaced by the following mixed states.

$$\rho_0(1^n) := \sum_{k, h, i} |\psi_k\rangle \langle \psi_k|, h, i, h(k)_i$$

and

$$\rho_1(1^n) := \begin{cases} \sum_{k, h, i, u} |\psi_k\rangle \langle \psi_k|, h, i, h(k)_{i-1}, u_1 & \text{if } i = \lceil \log N_k \rceil + 1 \\ \sum_{k, h, i, u} |\psi_k\rangle \langle \psi_k|, h, i, h(k)_i & \text{otherwise} \end{cases}$$

Unfortunately, classical arguments demonstrating statistical entropy gap and computational indistinguishability do not extend to the mixed states above, because of barriers that we describe next.

³This step requires the hash to be a specific inner-product based function which is compatible with the Goldreich-Levin technique.

Barrier 1: Non-orthogonality of Outputs, or, What is a Preimage Anyway? The mixed state ρ_1 above is well-defined only when N_k is as well. In the classical setting, N_k denotes the number of pre-images of $f(k)$. But it is unclear how to define “preimages” of a quantum state under a OWSG. For two keys x and x' , the corresponding OWSG output states $|\psi_x\rangle$ and $|\psi_{x'}\rangle$ could have arbitrary overlap. What overlaps qualify x' to be a pre-image of $|\psi_x\rangle$? One could consider fixing some inverse polynomial function (say $\frac{1}{n}$) and say that x' is a pre-image of $|\psi_x\rangle$ whenever $\langle \psi_x | \psi_{x'} \rangle \geq \frac{1}{n}$. Unfortunately, setting an arbitrary threshold does not accurately capture the adversary’s uncertainty about k , given $|\psi_k\rangle$. In fact, such an approach is fundamentally doomed for the following reason.

It is possible to build one-way state generators that are *unconditionally* statistically uninvertible given only a single copy of the output state $|\psi_k\rangle$. A simple example is the following construction based on Weisner encodings/BB84 states. On input classical key $k = (\theta, x)$ where $\theta, x \leftarrow \{0, 1\}^n$, the OWSG outputs pure state $|x\rangle_\theta$. This OWSG is *statistically* single-copy secure, because $|x\rangle_\theta$ hides the string θ (over the randomness of the choice of x).

Since quantum bit commitments cannot be secure against unbounded adversaries, this would rule out any possible constructions of commitments (including the one above) that rely only on the existence of single-copy (pure) OWSG. Instead, we will crucially rely on multi-copy security of the OWSG to obtain an intermediate primitive where for every pair of keys (k_1, k_2) , their images are either orthogonal or parallel.

Resolving Barrier 1: From Quantum to Classical Outputs via Shadow Tomography. Shadow tomography, introduced in [2], allows one to estimate a large number of observables by obtaining classical information from relatively few copies of an unknown quantum state. In more detail, shadow tomography is a procedure that applied to $t = \text{poly}(n, \frac{1}{\epsilon})$ copies of an unknown state $|\psi\rangle$ yields a *classical* string, the shadow S . Given S , it is possible to simultaneously estimate $\langle \psi | O_j | \psi \rangle$ upto ϵ error for an exponentially large number of observables $\{O_j\}_{j \in [2^n]}$.

Applying shadow tomography to a OWSG with pure outputs yields (at least) a statistical inverter for the OWSG. Given $t = \text{poly}(n)$ copies of some state $|\psi_k\rangle$, an inverter can use shadow tomography on $|\psi_k\rangle^{\otimes t}$ to (inefficiently) find a k' such that $\langle \psi_k | \psi_{k'} \rangle > 1 - \frac{1}{n}$.

Given a OWSG output state $|\psi_k\rangle$, it may even be tempting to define its “preimages” as the set of possible keys k' returned by this statistical inverter, and try to apply arguments similar to the classical argument above. Unfortunately this approach breaks down too. The statistical inverter given $|\psi_k\rangle^{\otimes t}$ only finds a key k' where $|\psi_{k'}\rangle$ has nontrivial *single-copy* overlap with $|\psi_k\rangle$. It is possible that for such k' , $\langle \psi_k | \psi_{k'} \rangle^{\otimes t}$ is close to 0. Thus k' is *not even close to being* a preimage of $|\psi_k\rangle^{\otimes t}$, at least for the purposes of arguing computational indistinguishability.

Thus instead of trying to define preimages of quantum states, we will crucially use the fact that certain shadow tomography methods [19] have efficiently computable classical shadows. We now outline how this fact turns out to be useful.

Our Main Insight. On input key k , instead of having ρ_0 (and ρ_1) contain one or more copies of the OWSG state $|\psi_k\rangle$, they will only

contain a classical shadow S_k of $|\psi_k\rangle$. Whenever these classical shadows can be efficiently computed, the WPEG distribution ρ_0 remains efficiently sampleable, and even becomes entirely classical!

While OWSG are defined to be secure given an arbitrary (unbounded) polynomial copies of $|\psi_k\rangle$, computing the shadow S_k will require only a fixed linear number of copies. Indeed, our proof shows that commitments are implied by a weaker variant of OWSG, where security only holds given a fixed linear number of copies of $|\psi_k\rangle$.

We point out that the shadow S_k is a *randomized* (i.e., not deterministic) function of the key k . Moreover, given a shadow S_k obtained from $|\psi_k\rangle$, it is computationally infeasible to find *any* key k' such that $|\psi_k\rangle$ and $|\psi_{k'}\rangle$ have non-negligible overlap, as otherwise this would break the OWSG. Indeed, this means that the (randomized) classical map $k \rightarrow S_k$ is efficiently computable but computationally uninvertible, assuming OWSG security. However, given a shadow S and a candidate key k , it is not possible to efficiently verify whether S was generated as a shadow of $|\psi_k\rangle$. Indeed, as discussed before, the resulting primitive necessitates inefficient verification. This approach allows us to build a *one-way puzzle from any pure-state OWSG*.

This allows us to reduce our problem to building commitments from one-way puzzles. The latter may at first appear to be easy, given the HILL technique. But the quantum nature of one-way puzzles leads to a major technical barrier, that we describe next.

Barrier 2: No Flatness in a Quantum World. Recall that a one-way puzzle sampler outputs classical (k, s) pairs which satisfy the following: (1) $\text{Ver}(k, s) = 1$ and (2) given s , it is computationally hard to find a preimage k such that $\text{Ver}(k, s) = 1$. Here, observe that the distribution on preimage keys k induced by fixing a puzzle output s is not a “flat” distribution, i.e., it *does not necessarily assign equal probability mass to each preimage key*. Why does this matter?

For the following discussion, given any puzzle output string s , we let K_s denote the distribution on keys induced by s , ℓ_s denote the min-entropy of K_s and $N_s = |\text{Supp}(K_s)|$. Since K_s is an arbitrary distribution, it can always be the case that $\ell_s \ll \lceil \log N_s \rceil$.

The construction of weak PEGs from one-way functions, discussed at the beginning of the overview, may seem to extend naturally to one-way puzzles as follows. Consider distributions

$$\begin{aligned} G_0(1^n) &:= s, h, i, h(k)_i \text{ and} \\ G_1(1^n) &:= \begin{cases} s, h, i, h(k)_{i-1}, u_1 & \text{if } i = \lceil \log N_s \rceil + 1 \\ s, h, i, h(k)_i & \text{otherwise} \end{cases} \end{aligned}$$

where $(k, s) \leftarrow \text{OWPuzzle.Samp}(1^n)$, $h \leftarrow \{0, 1\}^n$, $i \leftarrow [n]$.

These distributions do differ in entropy, but they may not be computationally indistinguishable. The leftover hash lemma (LHL) would imply that for any s , the first $\ell_s - 2c \log n$ bits of $h(k)$ are $\frac{1}{n^c}$ -statistically close to uniform given s , where ℓ_s is the min-entropy of K_s . Any subsequent bits may leak information about the preimage k . However, we note that G_0 and G_1 differ on the i^{th} bit of $h(k)$ for $i = \lceil \log N_s \rceil + 1$ and $\lceil \log N_s \rceil \gg \ell_s$. But it is possible that all remaining bits of $h(k)$, i.e., $h(k)_{\lceil \log N_s \rceil}$, computationally leak the entire key k . This would make the distributions G_0 and G_1 easily computationally distinguishable.

The argument above describes why modifying G_1 on $i = \lceil \log N_s \rceil + 1$ doesn't work. What if we instead modified G_1 on $i = \ell_s + 1$ instead,

where ℓ_s is the min-entropy in K_s . That is, consider changing G_1 to the following distribution

$$G_1'(1^n) := \begin{cases} s, h, i, h(k)_{i-1}, u_1 & \text{if } i = \ell_s + 1 \\ s, h, i, h(k)_i & \text{otherwise} \end{cases}$$

In this case, the distributions G_0 and G_1 become computationally indistinguishable, but the last bit in G_0 corresponding to $i = \ell_s + 1$ could also be statistically close to uniform. As a result, G_0 and G_1' could end up being (almost) identical, with no entropy gap at all!

This problem does not arise in the classical setting, because flat preimages can be assumed without loss of generality by “pulling out” the (uniform) randomness from any classical algorithm. Letting r denote the randomness used to sample k , one can always define a (one-way) function that uses its uniform input r to sample k and finally outputs $y = f(k)$. This ensures that $\ell_s = \log N_s$ above, enabling simultaneous arguments for both computational indistinguishability and statistical entropy gap. Unfortunately, this type of flattening is no longer possible when the sampler is a *quantum circuit*, because the randomness comes from a quantum process and we do not know how to explicitly pull it out.

At this point, it is natural to wonder whether there is some index i for every key k such that changing the i^{th} bit of $h(k)$ in G_1 yields a distribution that is computationally indistinguishable from G_0 but has a statistical entropy gap. For example, perhaps one could consider modifying G_1 at the first i for which the statistical distance between G_0 and G_1 *jumps* from a value that is negligible at $i - 1$ to a value that is not negligible at i . But there exist distributions for which there never is a clear cut “jump”; for example, if the statistical distance between G_0 and G_1 increases proportionally to $2^{-(n-i)}$.

To overcome this issue, we will further modify G_1 . Our starting idea will be to fix for every puzzle s , a “good set” $\mathbb{G}_s$ of preimage keys which is almost flat. We set the distribution G_1 to differ from G_0 *only when the key k that is output by Samp belongs to the set $\mathbb{G}_s$* . Making this approach work requires several additional ideas, and we provide a detailed overview of these below.

Resolving Barrier 2: Pseudoentropy or “HILL” for Quantumly Sampled Distributions. Our goal is to prove that the distribution

$$G_0(1^n) := s, h, i, h(k)_i$$

is a weak pseudoentropy generator, where $(k, s) \leftarrow \text{Samp}(1^n)$, $h \leftarrow \{0, 1\}^n$ and $i \leftarrow [n]$. Recall that this means we must demonstrate the existence of a different distribution which is computationally close to but has more entropy than G_0 .

We already outlined why setting G_1 identically to the case of one-way functions creates issues with proving either computational indistinguishability or statistical entropy gap. Instead, we observe that for every puzzle s and corresponding preimage distribution K_s , there is a “good” subset $\mathbb{G}_s$ of preimages k such that⁴:

- (1) The set $\mathbb{G}_s$ is dense enough in K_s , that is, for every s ,

$$\Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s] \geq \frac{1}{n}, \text{ and}$$

- (2) For every pair of preimages $(k_1, k_2) \in \mathbb{G}_s$,

$$0.5 \Pr_{K_s} [k_2] \leq \Pr_{K_s} [k_1] \leq 2 \Pr_{K_s} [k_2]$$

⁴We use numbers like $\frac{1}{n}$ below for simplicity. In our main section, we use slightly different fractions than the ones depicted here, for various technical reasons.

The observation above essentially follows from a pigeonhole argument over the preimages of s .

We can now consider a *different* simulated distribution G_1 as follows.

$$G_1(1^n) := \begin{cases} s, h, i, h(k)_{i-1}, u_1 & \text{if } i = \log |\mathbb{G}_s| + 1 \text{ and } k \in \mathbb{G}_s \\ s, h, i, h(k)_i & \text{otherwise} \end{cases}$$

where $(k, s) \leftarrow \text{Samp}(1^n)$, $h \leftarrow \{0, 1\}^n$, $i \leftarrow [n]$, and $\mathbb{G}_s$ is the good slice of preimages defined above. Computational indistinguishability between G_0 and G_1 follows by noting that any distinguishing advantage can only exist when $k \in \mathbb{G}_s$. Because we are reducing to a search problem, it is still possible to apply the Leftover Hash Lemma and the Goldreich-Levin theorem to convert any distinguisher into an inverter for the one-way puzzle.

Moreover, conditioned on $k \in \mathbb{G}_s$, G_1 obviously has more entropy than G_0 (and when $k \notin \mathbb{G}_s$, the two distributions are identical)⁵. This unfortunately *does not* imply that G_1 has more entropy than G_0 overall. The reason can best be explained with the following toy examples.

Example 1: There is a (hidden) event B that occurs with probability $\frac{2}{3}$, and distributions (A_0, A_1) such that

- Distribution A_0 outputs 0 when B occurs, and 1 when B doesn't occur.
- Distribution A_1 outputs a uniform bit when B occurs, and 1 when B doesn't.

A_1 has more entropy than A_0 conditioned on B , and the distributions are identical when B doesn't occur – but the overall entropy in A_0 is equal to that of A_1 ! Similarly, while G_1 has more entropy than G_0 when $k \in \mathbb{G}_s$, and the distributions are identical otherwise, the overall entropy in G_0 could end up being equal to that in G_1 .

Consider, however, the following example where both distributions A_0 and A_1 are uniform when B doesn't occur, i.e.,

Example 2: There is a (hidden) event B that occurs with probability $\frac{2}{3}$, and distributions (A_0, A_1) such that

- Distribution A_0 outputs 0 when B occurs, and a uniform bit when B doesn't occur.
- Distribution A_1 outputs a uniform bit when B occurs, and a uniform bit when B doesn't.

A_1 has more entropy than A_0 conditioned on B . Moreover, since A_0 and A_1 are *uniform* when B doesn't occur, then A_1 having higher entropy than A_0 conditioned on B does imply that A_1 has higher entropy overall.

We could hope to apply a similar argument to G_0 and G_1 if somehow it were the case that for $k \notin \mathbb{G}_s$, the last bit of G_0 (and also G_1) is close to uniform given the remaining bits $s, h, i, h(k)_{i-1}$. But why would this even be the case?

⁵It may appear that we are close: we seem to have a pair of distributions that are statistically far but computationally close. Unfortunately, to obtain a commitment, we also need these distributions to be efficiently sampleable, which is something we will address in a later subsection. At this point, sampling from G_1 requires knowing i_s and $\mathbb{G}_s$ which are not necessarily efficiently computable functions of (k, s) . But for now, we only aim to prove that G_0 is a weak PEG, for which we only need G_0 to be efficiently sampleable, and to prove that G_1 has *more entropy* than G_0 .

Establishing an Entropy Gap. Inspired by the insight above, we will modify G_1 to provably obtain an entropy gap. For every s , the two distributions are identical when $i \neq \log |\mathbb{G}_s| + 1$, therefore, we only focus on the case where $i = \log |\mathbb{G}_s| + 1$.

In this case, ideally we want the bias in the i^{th} bit of $h(k)$ when $k \notin \mathbb{G}_s$ to not cancel out the bias that arises when $k \in \mathbb{G}_s$, in the distribution G_0 . This would hold if the i^{th} bit of $h(k)$ when $k \notin \mathbb{G}_s$ were uniform (even given the remaining bits output by G_0). But we do not know if this is the case, or even what the distribution of keys when $k \notin \mathbb{G}_s$ looks like.

To resolve this, let us first try to ensure that for *most* $k \in \mathbb{G}_s$, all preimages (besides k) of $(s, h(k)_{i-1})$ have *extremely* low sampling probability in K_s . (Recall that K_s is the distribution induced on preimages of s .) This can be achieved by modifying G_1 as follows, where $i_s^* = (\log |\mathbb{G}_s| + 600 \log n)$:

$$G_1(1^n) := \begin{cases} s, h, i, h(k)_{i-1}, u_1 & \text{if } i = i_s^* \text{ and } k \in \mathbb{G}_s \\ s, h, i, h(k)_i & \text{otherwise} \end{cases}$$

In the modified distribution G_1 , pairwise independence of h ensures that for most $k \in \mathbb{G}_s$, all preimages (besides k) of $(s, h(k)_{i-1})$ are sampled with probability less than $\frac{1}{n^{600}}$ in K_s , for $i = (\log |\mathbb{G}_s| + 600 \log n)$. For this overview, we assume that this holds for *all* $k \in \mathbb{G}_s$ ⁶.

We will now consider the following two cases.

- (1) *Preimages (besides k) of $(s, h(k)_{i-1})$ are sampled with total probability $\leq \frac{1}{n}$ in K_s .*

In this case, since the unique $k \in \mathbb{G}_s$ has probability mass at least $1 - \frac{1}{n}$, the bias in the i^{th} bit of $h(k)$ from keys outside $\mathbb{G}_s$ barely stacks up against the bias that arises from $k \in \mathbb{G}_s$.

- (2) *Preimages (besides k) of $s, h(k)_{i-1}$ are sampled with total probability $> \frac{1}{n}$ in K_s .*

In this case, the individual probability mass of every preimage (besides k) is very low, but their total probability mass is high. This means that the overall distribution of pre-images (besides k) of $(s, h(k)_{i-1})$ necessarily has high entropy. Then by the Leftover Hash Lemma, the i^{th} bit of $h(k)$ when $k \notin \mathbb{G}_s$ will be close to uniform, which is what we desired.

In both cases, we conclude that the *overall* Shannon entropy in G_1 is larger than that in G_0 by a (fixed) inverse polynomial value: which means that G_0 is indeed a WPEG.

Furthermore, the proof of computational indistinguishability between G_0 and G_1 goes through as before, with the leftover hash lemma guaranteeing that all but the last $O(\log n)$ bits of $h(k)$ are statistically indistinguishable from uniform. With this guarantee, the Goldreich-Levin inverter simply needs to guess the last $O(\log n)$ bits of $h(k)$, which it can do with inverse polynomial probability.

At this point, we have a WPEG. However, because G_1 cannot be efficiently sampled, we cannot use it directly to build commitments. Indeed, obtaining a full-fledged commitment requires additional steps, which we outline next.

⁶In the technical sections, we further modify G_1 to account for the fact that for a few choices of h and a few $k \in \mathbb{G}_s$, there are multiple preimages of $(h, h(k))$ that are each sampled with probability much higher than $\frac{1}{n^{600}}$.

2.2 Pseudoentropy Generators (PEG)

Our next step follows a similar technique as [17] to (1) amplify the entropy gap between real and simulated distributions and (2) bring the min-entropy of the real distribution close to its Shannon entropy, all while maintaining computational indistinguishability.

This is done by taking a product distribution of the outputs of the weak PEG. In more detail, we sample $q(n) = \text{poly}(n)$ (for a large enough polynomial $\text{poly}(\cdot)$) random keys $k_1, \dots, k_q$ along with $q(n)$ independent (h, i) values. We use these to generate $q(n)$ samples from distribution G_0 , and we append these samples together as our PEG output. This also has the effect of “concentrating” the entropy to an expected value independent of the choice of k (whereas in weak PEG this entropy would necessarily depend on k via i_s). In the PEG, we have that for every choice of security parameter n , there is a *single* value $\hat{h}_n$ that corresponds to the Shannon-entropy in the output of the PEG, and this value is smaller than the min-entropy in the corresponding simulated product distribution by n^c , for some $c > 1$.

2.3 Imbalanced EFI Pairs

An EFI is a pair of *efficiently sampleable* distributions that are statistically far but computationally indistinguishable. Such distributions have been shown [10, 37] to be equivalent to quantum bit commitments.

Let us consider hashing the output of our PEG to approximately $\hat{h}_n + n$ bits. That is, the size of hash outcome is larger than the actual Shannon entropy in the PEG output, making the resulting distribution statistically distinguishable from uniform. At the same, since the PEG outputs are *computationally indistinguishable* from a distribution with (much) more than $\hat{h}_n + n$ bits of min-entropy, the resulting hash output is still *computationally indistinguishable* from uniform. It may now seem like we have an EFI pair: consider distributions

- $h(\text{PEG}(n))$ truncated to $\hat{h}(n) + n$ bits
- $U_{\hat{h}(n)+n}$ which is uniform over $\hat{h}(n) + n$ bits

While these distributions are computationally close but statistically far, they cannot be sampled efficiently without non-uniform advice, i.e., the value $\hat{h}(n)$ for every n .

In fact, observe that truncating the hash output to any less than $\hat{h}_n + n$ bits would still preserve computational indistinguishability, and truncating to any more would still ensure statistical distance, but the two can *simultaneously* be guaranteed only when truncating to exactly $\hat{h}(n) + n$ bits. This is why we call the resulting object an *imbalanced EFI*.

Due to the equivalence between EFI and commitment, we can equivalently claim to have statistically binding, computationally hiding quantum bit commitments [10, 37] – albeit dependent on non-uniform advice $z(n)$. When $z(n) \leq \hat{h}(n) + n$, the commitments are hiding, and when $z(n) \geq \hat{h}(n) + n$, these commitments are binding. We call this an *imbalanced commitment scheme*. The next few steps discuss how to remove this imbalanced drawback by appropriately combining variants of these commitments⁷.

⁷This upcoming part diverges from techniques in [17] which build uniform pseudorandom generators by appropriately stretching the output of a nonuniform PRG. These techniques break down in our setting because there is no clear way to run a puzzle on

2.4 Always Binding, Non-Uniform Hiding Commitments

In the next step, we rely on prior work in *flavor conversion* of quantum commitments [18] to convert our statistically hiding, computationally binding EFI pairs/commitments to commitments with the reverse property: namely, where for $z(n) \geq \hat{h}(n) + n$, the commitments are hiding, and when $z(n) \leq \hat{h}(n) + n$, these commitments are binding. Next, given these two types of complementary commitments, we *combine them* by using both to commit to the *same* bit b : note that for every choice of advice $z(n)$ (i.e. length to which we truncate the hash outcome), at least one of the two commitments is necessarily binding. This allows us to show that the resulting combined commitment is always binding (for every choice of $z(n)$) and hiding whenever $z(n) = \hat{h}(n) + n$. We call this a non-uniform hiding commitment.

2.5 Standard (Uniform) Commitments

Finally, we observe that for each n , the number of possible $\hat{h}(n)$ values is bounded by a fixed polynomial $t(n)$. Thus, we can repeat the above construction for every possible value of $\hat{h}(n)$, obtaining a sequence of commitments where for every n , at least one commitment in the sequence is hiding (and all are binding). By secret sharing the committed bit between various commitments, we can show that the overall commitment scheme satisfies both hiding and binding. Thus, we have removed dependence on the advice string $z(n)$, yielding a uniform construction of commitments.

Quick Detour: An Alternative Template. We briefly note an alternative technique [32] using quantum information to sidestep the use of the hardcore bit. While this was developed to build commitments from a strong “injective” variant of OWSG, for simplicity, we describe it as applied to injective one-way functions. Very roughly (and ignoring some garbage registers), a commitment to 0 is $\sum_k |k\rangle_C |k, f(k)\rangle_D$ and a commitment to 1 is $\sum_k |k\rangle_C |0, f(k)\rangle_D$; where C is the commit register and D the decommit register. Statistical hiding follows because tracing out the D register leaves identical mixtures on C in both cases. Computational binding follows by the hardness of finding k given $f(k)$ for a random k . Our methods, including hashing preimages to appropriate lengths and slicing, will also similarly apply to this template. We do not find any one of these templates to be simpler than the other, but we focus on the Goldreich-Levin template because it yields interesting intermediate primitives with entirely classical outputs.

QCCC Cryptography implies One-Way Puzzles. Finally, we provide evidence that one-way puzzles are a necessary assumption for quantum cryptography with classical communication by showing that central QCCC cryptographic primitives imply one-way puzzles.

3 PURE OWSG IMPLY ONE-WAY PUZZLES

In this section, we show how to use shadow tomography to build one-way puzzles (with inefficient verification) from any OWSG with pure state outputs.

its own output, and thus to achieve significant “stretch” in a puzzle-based PRG-type object.

DEFINITION 3.1 ($t(n)$ -COPY SECURE ONE-WAY STATE GENERATORS). [32] A one-way state generator (OWSG) is a set of QPT algorithms (KeyGen, StateGen, Ver) where:

- KeyGen(1^n): On input the security parameter n , output a classical key string $k \in \{0, 1\}^n$.
- StateGen(k): On input key $k \in \{0, 1\}^n$, output an m -qubit quantum state $|\psi_k\rangle$.
- Ver($k, |\psi\rangle$): On input key $k \in \{0, 1\}^n$ and m -qubit quantum state $|\psi\rangle$, output $\top$ or $\perp$.

These algorithms satisfy the following properties.

- **Correctness.** For every $n \in \mathbb{N}$,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ |\psi_k\rangle \leftarrow \text{StateGen}(k)}} [\top \leftarrow \text{Ver}(k, |\psi_k\rangle)] \geq 1 - \text{negl}(n)$$

- **$t(n)$ -Copy Security.** For every quantum polynomial-sized adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and $n \in \mathbb{N}$,

$$\Pr \left[\top \leftarrow \text{Ver}(\mathcal{A}_n((|\psi_k\rangle)^{\otimes t(n)}, |\psi_k\rangle)) \right] \leq \text{negl}(n)$$

where $k \leftarrow \text{KeyGen}(1^n)$ and $|\psi_k\rangle \leftarrow \text{StateGen}(k)$

This definition was later generalized in [31] to allow StateGen outputs to be mixed states. Furthermore, existing definitions of OWSG [32] require $t(n)$ -copy security for every (a-priori unbounded) polynomial $t(\cdot)$. In this work, we only need to consider cn -copy security for a large enough, a-priori fixed, constant c . We will show that even this weaker variant implies commitments, thus obtaining a stronger result.

DEFINITION 3.2 (ONE-WAY PUZZLES). A one-way puzzle consists of a pair (Samp, Ver) with the following syntax.

- Samp(1^n) $\rightarrow (k, s)$, outputs a pair of classical strings (k, s) . We refer to s as the puzzle and k as its key. Without loss of generality we may assume that $k \in \{0, 1\}^n$.
- Ver(k, s) $\rightarrow \top$ or $\perp$, is a Boolean function that maps any pair of classical strings (k, s) to either $\top$ or $\perp$.

These satisfy the following properties.

- **Correctness.** Outputs of the sampler pass verification with overwhelming probability, i.e.,

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [\text{Ver}(k, s) = \top] = 1 - \text{negl}(n)$$

- **Security.** Given s , it is (quantum) computationally infeasible to find k satisfying $\text{Ver}(k, s) = \top$, i.e., for every quantum polynomial-sized adversary $\mathcal{A}$,

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [\text{Ver}(\mathcal{A}(s), s) = \top] = \text{negl}(n)$$

Note that since puzzles are efficiently sampleable, there exists a polynomial $p(\cdot)$ such that $|s| \leq p(n)$.

We will rely on the following theorem on shadow tomography from [19].

THEOREM 3.1. [19] (Rephrased, following [38]) Fix any $\epsilon, \delta > 0$. There exists a polynomial $p(\cdot)$ and QPT algorithm ShadowGen that, given $T = O(\log(1/\delta)/\epsilon^2)$ copies of an unknown state $|\psi\rangle$ generates a classical string (called the “shadow”) S of size $p(n)$ with the following property:

For some $t \in \mathbb{N}$, let $\{M_i\}_{i \in [t]}$ be a set of observables such that $\text{Tr}(M_i^2) \leq 1$. Then there exists an “estimator” function E such that:

$$\Pr \left[\forall i \in [t], |E(S, M_i) - \langle \psi | M_i | \psi \rangle| \leq \epsilon \right] \geq 1 - t\delta$$

where $S \leftarrow \text{ShadowGen}(|\psi\rangle^{\otimes T})$

We now proceed to state our main theorem for this section.

THEOREM 3.2. There exists a constant $c > 0$ such that any cn -copy secure one-way state generator with pure state outputs (Definition 3.1) implies a one-way puzzle (Definition 3.2).

PROOF. (of Theorem 3.2) Let (KeyGen, StateGen) be a one-way state generator (OWSG) with pure state outputs and let $|\psi_k\rangle$ represent the output of StateGen(k).

To build a puzzle from this OWSG, we will apply shadow tomography to the output states of the OWSG. In fact, the one-way puzzle will simply sample a OWSG key k , compute $|\psi_k\rangle \leftarrow \text{StateGen}(k)$, and finally compute s_k as a classical shadow of $|\psi_k\rangle$. It will output s_k as the puzzle, with solution k . In what follows, we formalize this construction and define an (inefficient) verification algorithm for the one-way puzzle.

Defining Preimage Keys of a Classical Shadow. First, it will be useful to define an (inefficient) algorithm $\mathcal{L}$ that obtains a classical shadow and outputs a list of keys, roughly corresponding to possible preimages of the shadow.

Set $\epsilon = 1/10$ and for $n \in \mathbb{N}$, set $\delta = \delta(n) = 2^{-2n}$. For all $k \in \text{Supp}(\text{KeyGen}(1^n))$, define $M_k := |\psi_k\rangle \langle \psi_k|$. Note that these satisfy $\text{Tr}(M_k^2) = 1$. Let ShadowGen and E be algorithms as defined by Theorem 3.1 applied to δ, ϵ and $\{M_k\}_{k \in \text{Supp}(\text{KeyGen}(1^n))}$. Let $T = T(n) = O(n)$ be the required number of copies, and let $t = t(n) := |\{M_k\}_{k \in \text{Supp}(\text{KeyGen}(1^n))}| \leq |\text{Supp}(\text{KeyGen}(1^n))| \leq 2^n$.

Define the (inefficient) deterministic algorithm $\mathcal{L}$ that takes a shadow s as input and outputs a list of keys such that the estimated overlap of the shadow with each key in the list is at least $1 - \epsilon$, i.e.

$$\mathcal{L}(s) = \left\{ k : \left(k \in \text{Supp}(\text{KeyGen}(1^n)) \right) \wedge \left(E(s, M_k) \geq 1 - \epsilon \right) \right\}$$

The following claim about the algorithm $\mathcal{L}$ states that for any key k , with high probability over sampling a corresponding shadow s_k of $|\psi_k\rangle$, (1) the key k appears in $\mathcal{L}(s_k)$ and (2) for all $j \in \mathcal{L}(s_k)$, the (pure) states $|\psi_k\rangle$ and $|\psi_j\rangle$ have high overlap. The proof of this claim follows from the correctness of shadow tomography (Theorem 3.1).

CLAIM 3.1. For large enough $n \in \mathbb{N}$, for all $k \in \text{Supp}(\text{KeyGen}(1^n))$:

- (1) $\Pr_{s_k \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})} [k \in \mathcal{L}(s_k)] \geq 1 - 2^{-n}$
- (2) $\Pr_{s_k \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})} [\forall j \in \mathcal{L}(s_k), |\langle \psi_k | \psi_j \rangle|^2 \geq 1 - 2\epsilon] \geq 1 - 2^{-n}$

PROOF. For any large enough $n \in \mathbb{N}$ and any $k \in \text{Supp}(\text{KeyGen}(1^n))$, applying Theorem 3.1 on δ, t, ϵ set as above, we have:

$$\Pr[\forall j \in \text{Supp}(\text{KeyGen}(1^n)), |E(s_k, M_j) - |\langle \psi_k | \psi_j \rangle||^2 \leq \epsilon] \geq 1 - t\delta \geq 1 - 2^{-n} \quad (1)$$

where $s_k \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})$. Setting $j = k$, we have:

$$\Pr[E(s_k, M_k) \geq 1 - \epsilon] \geq 1 - 2^{-n}$$

By definition of $\mathcal{L}$, this implies

$$\Pr[k \in \mathcal{L}(s_k)] \geq 1 - 2^{-n}$$

which is the first part of the claim.

Again, fix any $k \in \text{Supp}(\text{KeyGen}(1^n))$. If we restrict j to $\mathcal{L}(s_k)$, then by equation (1), we have

$$\Pr[\forall j \in \mathcal{L}(s_k), |E(s_k, M_j) - |\langle \psi_k | \psi_j \rangle|^2| \leq \epsilon] \geq 1 - 2^{-n}$$

But $j \in \mathcal{L}(s_k) \iff E(s_k, M_j) \geq 1 - \epsilon$. Substituting in the above equation gives:

$$\Pr[\forall j \in \mathcal{L}(s_k), |\langle \psi_k | \psi_j \rangle|^2 \geq 1 - 2\epsilon] \geq 1 - 2^{-n}$$

which is the second part of the claim. $\square$

Before describing our puzzle, we will define the set $\mathbb{C}$ of keys that have low correctness error, as follows:

$$\mathbb{C} := \{k \in \{0, 1\}^n \text{ such that } \Pr[\text{Ver}(k, |\phi_k\rangle) = \top] \geq 1 - 1/100\}$$

Observe that with overwhelming probability, the OWSG KeyGen algorithm outputs keys in the set $\mathbb{C}$ (otherwise, this would contradict correctness of the one-way puzzle). Looking ahead, our puzzle verification algorithm will reject keys that are not in $\mathbb{C}$. We can now formally describe the puzzle.

Constructing the One-Way Puzzle. Define a one-way puzzle as follows.

- $\text{Puzz.Samp}(1^n)$:
 - Sample $k \leftarrow \text{KeyGen}(1^n)$.
 - Compute $s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})$
 - Return (k, s)
- $\text{Puzz.Ver}(k, s)$:
 - If $k \in \mathcal{L}(s)$ and $k \in \mathbb{C}$, return $\top$
 - Else return $\perp$

CLAIM 3.2. ($\text{Puzz.Samp}, \text{Puzz.Ver}$) satisfies Definition 3.2.

Correctness. By Claim 3.1 part 1, for large enough $n \in \mathbb{N}$ and all $k \in \text{Supp}(\text{KeyGen}(1^n))$

$$\Pr_{s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})} [k \in \mathcal{L}(s)] \geq 1 - 2^{-n}$$

Since the OWSG must have negligible correctness error, a Markov argument applied to Definition 3.1 shows that

$$\Pr_{k \leftarrow \text{KeyGen}(1^n)} [k \notin \mathbb{C}] \leq \text{negl}(n)$$

Putting these together,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})}} [k \in \mathcal{L}(s) \text{ and } k \in \mathbb{C}] \geq 1 - 2^{-n} - \text{negl}(n)$$

which by the definition of Puzz.Ver implies

$$\Pr_{(k,s) \leftarrow \text{Puzz.Samp}(1^n)} [\top \leftarrow \text{Puzz.Ver}(k, s)] \geq 1 - 2^{-n} - \text{negl}(n)$$

Security. We prove one-wayness by contradiction. Suppose there exists a quantum polynomial-sized adversary $\mathcal{A}$ that breaks the

one-wayness of the puzzle, i.e. there exists a polynomial $q(\cdot)$ such that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{(k,s) \leftarrow \text{Puzz.Samp}(1^n)} [\top \leftarrow \text{Puzz.Ver}(\mathcal{A}(s), s)] \geq 1/q(n)$$

We build a reduction that breaks the one-wayness of the OWSG. First, by the definition of Puzz.Ver and Puzz.Samp , for infinitely many $n \in \mathbb{N}$,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})}} [\mathcal{A}(s) \in (\mathcal{L}(s) \cap \mathbb{C})] \geq 1/q(n) \quad (2)$$

By Claim 3.1, for all $n \in \mathbb{N}$ and all $k \in \text{Supp}(\text{KeyGen}(1^n))$,

$$\Pr_{s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})} [\forall k' \in \mathcal{L}(s), |\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5] \geq 1 - 2^{-n} \quad (3)$$

For any events A and B , $\Pr[A \wedge B] \geq \Pr[A] - \Pr[\neg B]$. Therefore, from equations (2) and (3), for infinitely many $n \in \mathbb{N}$,

$$\Pr \left[(\mathcal{A}(s) \in (\mathcal{L}(s) \cap \mathbb{C})) \wedge (\forall k' \in \mathcal{L}(s), |\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5) \right] \geq 1/q(n) - 2^{-n} \quad (4)$$

where $k \leftarrow \text{KeyGen}(1^n)$ and $s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})$, which can be simplified to say that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{k' \leftarrow \mathcal{A}(s)} \left[(k' \in \mathbb{C}) \wedge (|\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5) \right] \geq 1/q(n) - 2^{-n}$$

where $k \leftarrow \text{KeyGen}(1^n)$ and $s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})$.

If $|\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5$ then the success probabilities of $\text{Ver}(k', |\psi_{k'}\rangle)$ and $\text{Ver}(k', |\psi_k\rangle)$ differ by at most $\frac{1}{\sqrt{5}}$. Since $\forall k' \in \mathbb{C}$, $\text{Ver}(k', |\psi_{k'}\rangle)$ succeeds with probability at least $1 - \frac{1}{100}$, for infinitely many $n \in \mathbb{N}$,

$$\begin{aligned} & \Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T}) \\ k' \leftarrow \mathcal{A}(s)}} [\top \leftarrow \text{Ver}(k', |\psi_k\rangle)] \\ & \geq (1 - 1/100 - 1/\sqrt{5}) \cdot (1/q(n) - 2^{-n}) > \frac{1}{2q(n)} \end{aligned}$$

Then, letting $\mathcal{B}$ be the algorithm that on input $|\psi_k\rangle^{\otimes T}$ outputs $\mathcal{A}(\text{ShadowGen}(|\psi_k\rangle^{\otimes T}))$, we have that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ k' \leftarrow \mathcal{B}(|\psi_k\rangle^{\otimes T})}} [\top \leftarrow \text{Ver}(k', |\psi_k\rangle)] > \frac{3}{5q(n)}$$

Since ShadowGen and $\mathcal{A}$ are quantum polynomial-sized circuits, this contradicts $O(n)$ -copy security of the OWSG. $\square$

4 ONE-WAY PUZZLES IMPLY QUANTUM COMMITMENTS

Here, we show that (inefficiently verifiable) one-way puzzles imply quantum weak pseudoentropy generators, defined below.

DEFINITION 4.1 (QUANTUM WEAK PSEUDOENTROPY GENERATOR). A Quantum Weak Pseudoentropy Generator consists of an ensemble of distributions $\{G_0(n), G_1(n)\}_{n \in \mathbb{N}}$ over classical strings :

- *Efficiency.* There exists a QPT algorithm $\mathcal{G}$ where for all $n \in \mathbb{N}$, $\mathcal{G}(1^n)$ returns a sample from $G_0(n)$.
- *Bounded Length.* There exists a polynomial $p(\cdot)$ such that for all $n \in \mathbb{N}$, for all $z_0 \in \text{Supp}(G_0(n))$, for all $z_1 \in \text{Supp}(G_1(n))$, $|z_0| = |z_1| \leq p(n)$.

- (Shannon) Entropy Gap. There exists an explicit constant $c > 0$ such that for all sufficiently large $n \in \mathbb{N}$,

$$H(G_1(n)) - H(G_0(n)) \geq \frac{1}{n^c}$$

- Indistinguishability. There exists a negligible function μ such that for all quantum polynomial-sized adversaries $\mathcal{A}$, for all large enough $n \in \mathbb{N}$:

$$\left| \Pr_{z \leftarrow \overline{G}_0(n)} [\mathcal{A}(z) = 1] - \Pr_{z \leftarrow \overline{G}_1(n)} [\mathcal{A}(z) = 1] \right| \leq \mu(n)$$

THEOREM 4.1. *One-way puzzles (Definition 3.2) imply quantum weak pseudoentropy generators (Definition 4.1).*

In this section, we show that a parallel repetition of quantum weak PEGs yields a strong pseudoentropy property, which we formalize into a quantum PEG, defined below.

DEFINITION 4.2 (QUANTUM PSEUDOENTROPY GENERATOR). A Quantum Pseudoentropy Generator consists of an ensemble of distributions $\{\overline{G}_0(n), \overline{G}_1(n)\}_{n \in \mathbb{N}}$ over classical strings such that:

- Bounded Length. There exists a polynomial $p(\cdot)$ such that for all $n \in \mathbb{N}$, for all $z_0 \in \text{Supp}(\overline{G}_0(n))$, for all $z_1 \in \text{Supp}(\overline{G}_1(n))$, $|z_0| = |z_1| \leq p(n)$.
- Efficiency. There exists some QPT algorithm that for all $n \in \mathbb{N}$, on input 1^n , returns a sample from $\overline{G}_0(n)$.
- Indistinguishability. There exists a negligible function μ such that for all quantum polynomial-sized adversaries $\mathcal{A}$, for all large enough $n \in \mathbb{N}$,

$$\left| \Pr_{z \leftarrow \overline{G}_0(n)} [\mathcal{A}(z) = 1] - \Pr_{z \leftarrow \overline{G}_1(n)} [\mathcal{A}(z) = 1] \right| \leq \mu(n)$$

- Entropy Gap. Here, we work with min and max entropies, as opposed to Shannon entropy. We require the min-entropy of $\overline{G}_1$ to be higher than the max-entropy of $\overline{G}_0$. Formally, there is some explicit constant $c > 0$ and some negligible function ϵ such that for all sufficiently large $n \in \mathbb{N}$,

$$H_{\min}^{\epsilon(n)}(\overline{G}_1(n)) - H_{\max}^{\epsilon(n)}(\overline{G}_0(n)) \geq n^c$$

THEOREM 4.2. *Quantum weak pseudoentropy generators (Definition 4.1) imply quantum pseudoentropy generators (Definition 4.2).*

Next, we show that quantum pseudoentropy generators imply a (non-uniform) variant of EFI, that we define below. This definition modifies the standard definition (Definition 4.4) of EFI to allow the algorithm to depend on a nonuniform parameter s , and requires the existence of a function $s^*(n)$ such that computational indistinguishability (resp. statistical distance) holds when $s \leq s^*(n)$ (resp. $s \geq s^*(n)$).

DEFINITION 4.3 (s^* -IMBALANCED EFI). Let $s^*(\cdot)$ denote a positive integer-valued function. An s^* -non-uniform EFI is a QPT algorithm $\text{EFI}_s(1^n, b) \rightarrow \rho_b$ that obtains classical parameter-dependent advice string s , and on input $b \in \{0, 1\}$ and security parameter n , outputs a (potentially mixed) quantum state such that:

- (1) Computational Indistinguishability. There exists a negligible function $\mu(\cdot)$ such that for all quantum polynomial-sized circuits $\mathcal{A}$, for large enough $n \in \mathbb{N}$ and every $s \leq s^*(n)$,

$$|\Pr[1 \leftarrow \mathcal{A}(\text{EFI}_s(1^n, 0))] - \Pr[1 \leftarrow \mathcal{A}(\text{EFI}_s(1^n, 1))]| \leq \mu(n)$$

- (2) Statistical Distance. There exists a negligible function $\delta(\cdot)$ such for large enough $n \in \mathbb{N}$ and every $s \geq s^*(n)$,

$$\text{TD}(\text{EFI}_s(1^n, 0), \text{EFI}_s(1^n, 1)) \geq 1 - \delta(n)$$

THEOREM 4.3. *There exists a positive integer-valued function $s^*(\cdot)$ such that*

- There exists a polynomial p such that for all $n \in \mathbb{N}$, $|s^*(n)| \leq \log p(n)$ and
- Quantum pseudoentropy generators (Definition 4.2) imply s^* -imbalanced EFI (Definition 4.3).

Finally, we show that s^* -imbalanced EFI (Definition 4.3) imply EFI pairs, defined below.

DEFINITION 4.4 (EFI PAIRS). [10] An EFI pair is a QPT algorithm $\text{EFI}(1^n, b) \rightarrow \rho_b$ that on input $b \in \{0, 1\}$ and the security parameter n , outputs a (potentially mixed) quantum state ρ_b such that the following hold:

- (1) Computational Indistinguishability. There exists a negligible function $\mu(\cdot)$ such that for every quantum polynomial-sized adversary $\mathcal{A}$, for large enough $n \in \mathbb{N}$,

$$|\Pr[1 \leftarrow \mathcal{A}(\text{EFI}(1^n, 0))] - \Pr[1 \leftarrow \mathcal{A}(\text{EFI}(1^n, 1))]| \leq \mu(n)$$

- (2) Statistical Binding. There exists a negligible function $\delta(\cdot)$ such that for large enough $n \in \mathbb{N}$,

$$\text{TD}(\text{EFI}(1^n, 0), \text{EFI}(1^n, 1)) \geq 1 - \delta(n)$$

THEOREM 4.4. *s^* -imbalanced EFI (Definition 4.3) imply EFI pairs (Definition 4.4).*

Putting them all together we obtain the following theorem.

THEOREM 4.5. *There exists a constant $c > 0$ such that cn copy secure one-way state generators with pure state outputs (Definition 3.1) imply EFI pairs (Definition 4.4)*

EFI pairs are known to imply secure computation for all classical and quantum functionalities [6, 8, 10, 15]. We therefore also have the following corollary.

COROLLARY 4.1. *There exists a constant $c > 0$ such that cn copy secure one-way state generators with pure state outputs imply secure computation for all quantum functionalities.*

QCCC Cryptography. Finally, we show that one-way puzzles are essential for QCCC cryptography.

THEOREM 4.6. *In the QCCC setting, the existence of public key encryption, digital signatures, bit commitments, or symmetric encryption with classical keys implies one-way puzzles.*

Complete proofs can be found in the full version of the paper.

ACKNOWLEDGMENTS

We thank James Bartusek, Yanyi Liu, Amit Sahai, and Taiga Hiroka for illuminating discussions and useful comments. The authors were supported in part by AFOSR, NSF 2112890 and NSF CNS-2247727. This material is based upon work supported by the Air Force Office of Scientific Research under award FA9550-23-1-0543.

REFERENCES

- [1] Scott Aaronson. 2016. The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes. arXiv:1607.05256 [quant-ph]
- [2] Scott Aaronson. 2020. Shadow Tomography of Quantum States. *SIAM J. Comput.* 49, 5 (2020), STOC18–368–STOC18–394. <https://doi.org/10.1137/18M120275X> arXiv:https://doi.org/10.1137/18M120275X
- [3] Scott Aaronson and Alex Arkhipov. 2014. Bosonsampling is far from uniform. *Quantum Inf. Comput.* 14, 15–16 (2014), 1383–1423. <https://doi.org/10.26421/QIC14.15-16-7>
- [4] Scott Aaronson and Greg Kuperberg. 2007. Quantum Versus Classical Proofs and Advice. *Theory Comput.* 3, 1 (2007), 129–157. <https://doi.org/10.4086/toc.2007.v003a007>
- [5] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. 2022. Pseudorandom (Function-Like) Quantum State Generators: New Definitions and Applications. In *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7–10, 2022, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13747)*, Eike Kiltz and Vinod Vaikuntanathan (Eds.). Springer, 237–265. https://doi.org/10.1007/978-3-031-22318-1_9
- [6] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. 2022. Cryptography from Pseudorandom Quantum States. In *Advances in Cryptology – CRYPTO 2022, Part I (Lecture Notes in Computer Science)*. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 208–236. https://doi.org/10.1007/978-3-031-15802-5_8
- [7] Per Austrin, Hao Chung, Kai-Min Chung, Shiuai Fu, Yao-Ting Lin, and Mohammad Mahmoody. 2022. On the Impossibility of Key Agreements from Quantum Random Oracles. In *Advances in Cryptology – CRYPTO 2022, Part II (Lecture Notes in Computer Science)*. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 165–194. https://doi.org/10.1007/978-3-031-15979-4_6
- [8] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. 2021. One-Way Functions Imply Secure Computation in a Quantum World. In *Advances in Cryptology – CRYPTO 2021, Part I (Lecture Notes in Computer Science, Vol. 12825)*, Tal Malkin and Chris Peikert (Eds.). Springer, Heidelberg, Germany, Virtual Event, 467–496. https://doi.org/10.1007/978-3-030-84242-0_17
- [9] Charles H Bennett and Gilles Brassard. 1984. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*. 175–179.
- [10] Zvika Brakerski, Ran Canetti, and Luowen Qian. 2023. On the Computational Hardness Needed for Quantum Cryptography. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10–13, 2023, MIT, Cambridge, Massachusetts, USA (LIPIcs, Vol. 251)*, Yael Tauman Kalai (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 24:1–24:21. <https://doi.org/10.4230/LIPIcs.ITCS.2023.24>
- [11] Zvika Brakerski, Paul F. Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. 2018. A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7–9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 320–331. <https://doi.org/10.1109/FOCS.2018.00038>
- [12] Shujiao Cao and Rui Xue. 2022. On Constructing One-Way Quantum State Generators, and More. *IACR Cryptol. ePrint Arch.* (2022), 1323. <https://eprint.iacr.org/2022/1323>
- [13] Kai-Min Chung, Eli Goldin, and Matthew Gray. 2024. On Central Primitives for Quantum Cryptography with Classical Communication. *IACR Cryptol. ePrint Arch.* (2024), 356. <https://eprint.iacr.org/2024/356>
- [14] Oded Goldreich and Leonid A. Levin. 1989. A Hard-Core Predicate for all One-Way Functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14–17, 1989, Seattle, Washington, USA*, David S. Johnson (Ed.). ACM, 25–32. <https://doi.org/10.1145/73007.73010>
- [15] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. 2021. Oblivious Transfer Is in MiniQCrypt. In *Advances in Cryptology – EUROCRYPT 2021, Part II (Lecture Notes in Computer Science, Vol. 12697)*, Anne Canteaut and François-Xavier Standaert (Eds.). Springer, Heidelberg, Germany, Zagreb, Croatia, 531–561. https://doi.org/10.1007/978-3-030-77886-6_18
- [16] Iftach Haitner, Omer Reingold, and Salil P. Vadhan. 2013. Efficiency Improvements in Constructing Pseudorandom Generators from One-Way Functions. *SIAM J. Comput.* 42, 3 (2013), 1405–1430. <https://doi.org/10.1137/100814421>
- [17] Johan Hästad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. 1999. A Pseudorandom Generator from any One-way Function. *SIAM J. Comput.* 28, 4 (1999), 1364–1396. <https://doi.org/10.1137/S0097539793244708>
- [18] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. 2023. From the Hardness of Detecting Superpositions to Cryptography: Quantum Public Key Encryption and Commitments. In *Advances in Cryptology – EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23–27, 2023, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 14004)*, Carmi Hazay and Martijn Stam (Eds.). Springer, 639–667. https://doi.org/10.1007/978-3-031-30545-0_22
- [19] Hsin-Yuan Huang, Richard Kueng, and John Preskill. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics* 16, 10 (01 Oct 2020), 1050–1057. <https://doi.org/10.1038/s41567-020-0932-7>
- [20] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. 1989. Pseudo-random Generation from one-way functions (Extended Abstracts). In *21st Annual ACM Symposium on Theory of Computing*. ACM Press, Seattle, WA, USA, 12–24. <https://doi.org/10.1145/73007.73009>
- [21] Russell Impagliazzo and Michael Luby. 1989. One-way Functions are Essential for Complexity Based Cryptography (Extended Abstract). In *30th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Research Triangle Park, NC, USA, 230–235. <https://doi.org/10.1109/SFCS.1989.63483>
- [22] Russell Impagliazzo and Steven Rudich. 1990. Limits on the Provable Consequences of One-way Permutations. In *Advances in Cryptology – CRYPTO’88 (Lecture Notes in Computer Science, Vol. 403)*, Shafi Goldwasser (Ed.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 8–26. https://doi.org/10.1007/0-387-34799-2_2
- [23] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. 2018. Pseudorandom Quantum States. In *Advances in Cryptology – CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 10993)*, Hovav Shacham and Alexandra Boldyreva (Eds.). Springer, 126–152. https://doi.org/10.1007/978-3-319-96878-0_5
- [24] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. 2023. Publicly Verifiable Deletion from Minimal Assumptions. In *Theory of Cryptography*, Guy Rothblum and Hoeteck Wee (Eds.). Springer Nature Switzerland, Cham, 228–245.
- [25] William Kretschmer. 2021. Quantum Pseudorandomness and Classical Complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, July 5–8, 2021, Virtual Conference (LIPIcs, Vol. 197)*, Min-Hsiu Hsieh (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2:1–2:20. <https://doi.org/10.4230/LIPIcs.TQC.2021.2>
- [26] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. 2023. Quantum Cryptography in Algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20–23, 2023*, Barna Saha and Rocco A. Servedio (Eds.). ACM, 1589–1602. <https://doi.org/10.1145/3564246.3585225>
- [27] Hoi-Kwong Lo and Hoi Fung Chau. 1997. Is quantum bit commitment really possible? *Physical Review Letters* 78, 17 (1997), 3410.
- [28] Alex Lombardi, Fermi Ma, and John Wright. 2023. A one-query lower bound for unitary synthesis and breaking quantum cryptography. *Cryptology ePrint Archive*, Paper 2023/1602. <https://eprint.iacr.org/2023/1602> <https://eprint.iacr.org/2023/1602>
- [29] Michael Luby and Charles Rackoff. 1986. Pseudo-random Permutation Generators and Cryptographic Composition. In *18th Annual ACM Symposium on Theory of Computing*. ACM Press, Berkeley, CA, USA, 356–363. <https://doi.org/10.1145/12130.12167>
- [30] Dominic Mayers. 1997. Unconditionally secure quantum bit commitment is impossible. *Physical review letters* 78, 17 (1997), 3414.
- [31] Tomoyuki Morimae and Takashi Yamakawa. 2022. One-Wayness in Quantum Cryptography. *Cryptology ePrint Archive*, Paper 2022/1336. <https://eprint.iacr.org/2022/1336> <https://eprint.iacr.org/2022/1336>
- [32] Tomoyuki Morimae and Takashi Yamakawa. 2022. Quantum Commitments and Signatures Without One-Way Functions. In *Advances in Cryptology – CRYPTO 2022, Part I (Lecture Notes in Computer Science)*. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 269–295. https://doi.org/10.1007/978-3-031-15802-5_10
- [33] Tomoyuki Morimae and Takashi Yamakawa. 2023. Proofs of Quantumness from Trapdoor Permutations. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10–13, 2023, MIT, Cambridge, Massachusetts, USA (LIPIcs, Vol. 251)*, Yael Tauman Kalai (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 87:1–87:14. <https://doi.org/10.4230/LIPIcs.ITCS.2023.87>
- [34] Tomoyuki Morimae and Takashi Yamakawa. 2023. Quantum Advantage from One-Way Functions. *CoRR abs/2302.04749* (2023). <https://doi.org/10.48550/arXiv.2302.04749> arXiv:2302.04749
- [35] Moni Naor. 1989. Bit Commitment Using Pseudo-Randomness. In *Advances in Cryptology – CRYPTO ’89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20–24, 1989, Proceedings (Lecture Notes in Computer Science, Vol. 435)*, Gilles Brassard (Ed.). Springer, 128–136. https://doi.org/10.1007/0-387-34805-0_13
- [36] Stephen Wiesner. 1983. Conjugate coding. *SIGACT News* 15 (1983), 78–88.
- [37] Jun Yan. 2022. General Properties of Quantum Bit Commitments (Extended Abstract). In *Advances in Cryptology – ASIACRYPT 2022, Part IV (Lecture Notes in Computer Science)*. Springer, Heidelberg, Germany, 628–657. https://doi.org/10.1007/978-3-031-22972-5_22
- [38] Henry Yuen and Asa Kosto. [n. d.]. Scribe notes on Shadow Tomography. <https://www.henryyuen.net/spring2022/lec5-shadow.pdf>

Received 13-NOV-2023; accepted 2024-02-11