

Commitments from Quantum One-Wayness

Dakshita Khurana*

Kabir Tomer*

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.

*UIUC. {dakshita, ktomer2}@illinois.edu

Contents

1	Introduction	1
2	Technical Overview	4
2.1	Weak Pseudoentropy Generators (WPEG)	5
2.2	Pseudoentropy Generators (PEG)	10
2.3	Imbalanced EFI Pairs	10
2.4	Always binding, Non-uniform hiding Commitments	11
2.5	Standard (Uniform) Commitments	11
3	Preliminaries	12
3.1	Notation and Conventions	12
3.2	Quantum Cryptographic Primitives	13
3.3	Entropy and Randomness Extractors	14
4	Pure OWSG imply One-Way Puzzles	15
5	One-way Puzzles imply Quantum Weak PEGs	18
5.1	Defining G_0 and G_1	19
5.2	Establishing an Entropy Gap	22
5.3	Establishing Computational Indistinguishability	26
6	Quantum Weak PEGs imply Quantum PEGs	30
7	Quantum PEGs Imply Imbalanced EFI	32
8	Imbalanced EFI Imply Commitments	35
8.1	Defining Commitments	35
8.1.1	Quantum Bit Commitments	35
8.1.2	Imbalanced Commitments	36
8.2	Imbalanced EFI imply Imbalanced Commitments	36
8.3	Obtaining a Non-Uniform Hiding Commitment	37
8.4	Obtaining a (Uniform) Commitment	42
9	Acknowledgments	46
A	One-Way Puzzles are Necessary for QCCC Cryptography	50
A.1	Public Key Encryption	50
A.2	Digital Signatures	52
A.3	Bit Commitments	53
A.4	Symmetric Encryption with Classical Keys and Ciphertexts	55
B	Proving Lemma 5.2	58
B.1	Preliminary Claims	58
B.2	Proof of the Lemma	60

1 Introduction

A one-way function is a classically efficiently computable function that is hard to invert. This is a fundamental hardness assumption, necessary for the existence of much of modern classical cryptography [LR86, IL89, ILL89]. 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, there are tasks like key exchange and secure multi-party computation that classically require stronger, more structured assumptions [IR90].

The relationship between computational hardness and cryptography appears to be drastically different in a quantum world. Here, the seminal works of Wiesner [Wie83] and Bennett and Brassard [BB84] 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 [LC97, May97], 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 [BCKM21, GLSV21] 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 [AQY22, MY22b] 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 [JLS18].

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) [Kre21] or if $P = NP$ (relative to a classical oracle) [KQST23]. 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 [MY22a], 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.

A one-way state generator (OWSG) [MY22b] is a quantum algorithm that given a secret key, generates a hard-to-invert quantum state. This is a natural quantum analogue of a one-way function, and appears to be weaker, as a definition, than most known quantum cryptographic primitives. It is in particular known to be implied by several quantum cryptographic primitives including quantum signature and encryption schemes, as well as quantum money¹.

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?

¹See, for example, [MY22a], the figure at <https://sattath.github.io/qcrypto-graph/> and references therein.

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 [AQY22, BCQ23] and builds secure multi-party computation from commitments [GLSV21, BCKM21, AQY22], 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 [MY22b] to only output pure states; but this definition was later generalized in [MY22a] to also allow mixed states. OWSG with pure state outputs were also studied in [CX22], 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 [AQY22]. 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 [AA14], 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 [MY22a]. 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 [HKP20] 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 arbitrary 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** is a pair of algorithms (Samp, Ver) where Samp is quantum polynomial time and 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.

Unlike prior definitions of one-way puzzles in the literature, we do not require the verification (Ver) algorithm to be efficient. 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 [Kre21]. 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 .

However, 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., [HILL99]) 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 [BCM⁺18, MY23a, MY23b], quantum commitments with classical communication [AGQY22], and even black-box separations for key exchange [ACC⁺22]. 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 [ACC⁺22], 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. In Appendix A, 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 [KNY23] 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, deterministic 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 [AQY22, MY22b] via a construction that roughly parallels the classical setting [Nao89]. They were also known [MY22a] 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 [Kre21] 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

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.

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 [AK07, Aar16], for which a recent work [LMW23] 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 inputs?

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 [GL89] shows that the bit $\langle k, r \rangle$ is hard-core 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. [HILL99] 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 [HILL99] 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 [HILL99] (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)* [HILL99, HRV13]. We will now describe this object in some more detail.

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

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

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

where k, 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 [HILL99, HRV13] 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., x . 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 \mathcal{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 \mathcal{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

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

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

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.

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. 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 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 [Aar20], 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 [HKP20] have efficiently computable classical shadows. We now outline how this fact turns out to be useful.

Our Main Insight is the following: 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. In Section 4 we formalize this approach 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 $\ell_s = N_s$ 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: Pseudentropy 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 pseudentropy 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 ,

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

2. For every pair of preimages $(x_1, x_2) \in \mathbb{G}_s$,

$$0.5\Pr_{K_s}[x_2] \leq \Pr_{K_s}[x_1] \leq 2\Pr_{K_s}[x_2]$$

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.

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.

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. That is,

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?

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 .) We show that this is achieved by changing G_1 as follows:

$$G_1(1^n) := \begin{cases} s, h, i, h(k)_{i-1}, u_1 & \text{if } i = (\log |\mathbb{G}_s| + 600 \log n) \text{ and } k \in \mathbb{G}_s \\ s, h, i, h(k)_i & \text{otherwise} \end{cases}$$

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 .

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. We prove this formally in Section 5.

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.

2.2 Pseudoentropy Generators (PEG)

Our next step follows a similar technique as [HILL99] 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 [BCQ23, Yan22] 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

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 $\frac{1}{n^{600}}$.

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 *be computationally indistinguishable* from uniform. It may now seem like we have an EFI pair: consider distributions

- $h(\text{PEG}(n))$ where the output of h is truncated to $\hat{h}(n) + n$ bits
- $U_{\hat{h}(n)+n}$ which is the uniform distribution 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 [Yan22, BCQ23] – 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.

2.4 Always binding, Non-uniform hiding Commitments

In the next step, we rely on prior work in *flavor conversion* of quantum commitments [HMY23] 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 [MY22b] 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

⁷This upcoming part diverges from techniques in [HILL99] 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 its own output, and thus to achieve significant “stretch” in a puzzle-based PRG-type object.

$\sum_k |k\rangle_C |k, f(k)\rangle_D$ 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.

Roadmap. In the following section, we recall useful definitions and facts about quantum cryptography and various notions of entropy. In Section 4, we show that OWSG with pure states imply one-way puzzles. Our construction of commitments from one-way puzzles is detailed in Sections 5-8. In Section 5, we show how one-way puzzles imply a weak pseudoentropy generator (WPEG). Then, in Section 6 we describe how a product distribution of quantum weak PEGs gives rise to a pseudoentropy generator (PEG). Next in Section 7 we obtain imbalanced EFI by appropriately hashing the output of the PEG. Finally, in Section 8 we apply flavor swap and other combiners to the imbalanced EFI to obtain a uniform construction of quantum bit commitments.

In Appendix A, we provide evidence that one-way puzzles are a necessary assumption for quantum cryptography with classical communication.

3 Preliminaries

In this section, we discuss some notation and preliminary information, including definitions, that will be useful in the rest of the exposition.

3.1 Notation and Conventions

We write $\text{negl}(\cdot)$ to denote any *negligible* function, which is a function f such that for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. For any k , we will denote by U_k the uniform distribution supported on k bits. We will use $\text{SD}(A, B)$ to denote the statistical distance between (classical) distributions A and B .

Quantum conventions. A register X is a named Hilbert space \mathbb{C}^{2^n} . A pure state on register X is a unit vector $|\psi\rangle \in \mathbb{C}^{2^n}$, and we say that $|\psi\rangle$ consists of n qubits. A mixed state on register X is described by a density matrix $\rho \in \mathbb{C}^{2^n \times 2^n}$, which is a positive semi-definite Hermitian operator with trace 1.

A *quantum operation* F is a completely-positive trace-preserving (CPTP) map from a register X to a register Y , which in general may have different dimensions. That is, on input a density matrix ρ , the operation F produces $F(\rho) = \tau$ a mixed state on register Y . A *unitary* $U : X \rightarrow X$ is a special case of a quantum operation that satisfies $U^\dagger U = U U^\dagger = \mathbb{I}^X$, where \mathbb{I}^X is the identity matrix on register X . A *projector* Π is a Hermitian operator such that $\Pi^2 = \Pi$, and a *projective measurement* is a collection of projectors $\{\Pi_i\}_i$ such that $\sum_i \Pi_i = \mathbb{I}$.

The security of all our constructions holds against adversaries that receive non-uniform quantum advice. More specifically, we refer to security against all quantum polynomial-sized adversaries in our definitions. By this we mean every family of (polynomial-sized) quantum circuits $\{C_n\}_{n \in \mathbb{N}}$ along with a family of states $\{\rho_n\}_{n \in \mathbb{N}}$ such that for every $n \in \mathbb{N}$, C_n obtains ancilla registers set to ρ_n . We rely on a non-uniform reduction in Section 6, which obtains (classical) non-uniform advice for every value of the security parameter n .

Finally, we often define security of primitives by bounding an adversary's advantage in a search/distinguishing game by $\text{negl}(n)$. Depending on the order of quantifiers, this is typically considered to mean one of two

things – (1) there exists a negligible function $\mu(\cdot)$ such that the probability that any quantum polynomial-sized adversary \mathcal{A} wins a search/distinguishing game is at most $\mu(n)$ (for large enough n), and (2) for every quantum polynomial-sized adversary \mathcal{A} , there is a negligible function $\mu_{\mathcal{A}}(\cdot)$ such that the probability that \mathcal{A} wins a search/distinguishing game is at most $\mu_{\mathcal{A}}(n)$. The two definitions/interpretations are equivalent (proved in [Bel02] for classical circuits but the proof also applies to quantum circuits), and we use both interchangeably.

3.2 Quantum Cryptographic Primitives

Definition 3.1 ($t(n)$ -Copy Secure One-Way State Generators). [MY22b] *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 $n \in \mathbb{N}$,*

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ |\psi_k\rangle \leftarrow \text{StateGen}(k)}} \left[\top \leftarrow \text{Ver} \left(\mathcal{A}_n(|\psi_k\rangle^{\otimes t(n)}, |\psi_k\rangle) \right) \right] \leq \text{negl}(n)$$

This definition was later generalized in [MY22a] to allow StateGen outputs to be mixed states. Furthermore, existing definitions of OWSG [MY22b] 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 is a pair of sampling and verification algorithms (Samp, Ver) with the following syntax.*

- **Samp**(1^n) $\rightarrow (k, s)$, *is a QPT algorithm that 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 an unbounded algorithm that on input any pair of classical strings (k, s) halts and outputs 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 $|s| \leq p(n)$.

Definition 3.3 (EFI pairs). [BCQ23] 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 $n \in \mathbb{N}$,

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

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

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

3.3 Entropy and Randomness Extractors

We will use $H(X)$ to denote the Shannon entropy, $H_{\min}(X)$ to denote the min-entropy and $H_{\max}(X)$ to denote the max-entropy of distribution X . For an arbitrary classical random variable X , the sample entropy [HRV13] of any $x \in \text{Supp}(X)$ is defined as

$$H_X(x) := \log(1/\Pr[X = x]).$$

The min-entropy of distribution X is then

$$H_{\min}(X) := \min_{x \in \text{Supp}(X)} H_X(x)$$

and the max-entropy is

$$H_{\max}(X) := \max_{x \in \text{Supp}(X)} H_X(x)$$

For discrete random variables X and Y , we define conditional entropy of X given Y as

$$H(X|Y) := \sum_{y \in \text{Supp}(Y)} \Pr[Y = y] \cdot H(X|Y = y)$$

We also note the chain rule for conditional entropy

$$H(X, Y) = H(X|Y) + H(Y)$$

Definition 3.4. (Smooth Min-Entropy) Let X be a discrete random variables. The ϵ -Smooth Min Entropy of X is defined as:

$$H_{\min}^\epsilon(X) := \sup_{X'} H_{\min}(X')$$

where X' is an arbitrary random variable at most ϵ statistically far from X .

We also have the following well-known lemmas and theorems about entropy and randomness extraction.

Lemma 3.1 (Leftover Hash Lemma). Fix $\epsilon, \epsilon' > 0$. Let X be a discrete random variable distributed over $\{0, 1\}^n$, and $H_{\min}^\epsilon(X) \geq k$. Let H be a universal hash family with output length $k - 2 \log(1/\epsilon')$. Then

$$\text{SD}\left((H(X), H), (U, H)\right) \leq \epsilon + \epsilon'$$

where U is uniformly distributed in $\{0, 1\}^{k-2 \log(1/\epsilon')}$.

Theorem 3.1 (Entropy Concentration in Product Distributions). [Ren08] Let X be a discrete random variable taking values in a universe \mathcal{U} , let $t \in \mathbb{N}$ and let $\epsilon > 0$. Then,

$$H_{\min}^\epsilon(X^t) \geq t \cdot H(X) - \sqrt{2t \cdot \log(1/\epsilon)} \cdot \log(3 + |\mathcal{U}|)$$

Finally, we will use the following version of the Goldreich-Levin theorem (in the presence of quantum advice).

Theorem 3.2 (Goldreich-Levin with Quantum Advice). [AC02] There exists a QPT algorithm \mathcal{E} such that if for any QPT algorithm \mathcal{D} , any $\epsilon > 0$, any mixed state ρ , and any $a \in \{0, 1\}^n$,

$$\Pr_{x \leftarrow \{0, 1\}^n} [\mathcal{D}(\rho, x) = \langle a, x \rangle] \geq \frac{1}{2} + \epsilon$$

then,

$$\Pr[\mathcal{E}^{\mathcal{D}}(\rho) = a] \geq 4\epsilon^2$$

4 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.

We will rely on the following theorem on shadow tomography from [HKP20].

Theorem 4.1. [HKP20] (Rephrased, following [YK]) 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:

$$\Pr_{S \leftarrow \text{ShadowGen}(|\psi\rangle^{\otimes T})} \left[\forall i \in [t], |E(S, M_i) - \langle \psi | M_i | \psi \rangle| \leq \epsilon \right] \geq 1 - t\delta$$

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

Theorem 4.2. 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 4.2) Let $(\text{KeyGen}, \text{StateGen})$ be a one-way state generator (OWSG) with pure state outputs and let $|\psi_k\rangle$ represent the output of $\text{StateGen}(k)$.

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 4.1 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 $1 - \epsilon$. That is,

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

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 4.1).

Claim 4.1. *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 4.1 on δ, t, ϵ set as above, we have:

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

Setting $j = k$, we have:

$$\Pr_{s_k \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})} [E(s_k, M_k) \geq 1 - \epsilon] \geq 1 - 2^{-n}$$

By definition of \mathcal{L} , this implies

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

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

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

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

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

which is the second part of the claim. □

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

Correctness. By Claim 4.1 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

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

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 $n \in \mathbb{N}$,

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

By Claim 4.1, 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 $n \in \mathbb{N}$,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T})}} \left[(\mathcal{A}(s) \in (\mathcal{L}(s) \cap \mathbb{C})) \wedge (\forall k' \in \mathcal{L}(s), |\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5) \right] \geq 1/q(n) - 2^{-n} \quad (4)$$

which can be simplified to say that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ s \leftarrow \text{ShadowGen}(|\psi_k\rangle^{\otimes T}) \\ k' \leftarrow \mathcal{A}(s)}} \left[(k' \in \mathbb{C}) \wedge (|\langle \psi_k | \psi_{k'} \rangle|^2 \geq 4/5) \right] \geq 1/q(n) - 2^{-n}$$

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

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

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

5 One-way Puzzles imply Quantum Weak PEGs

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

Definition 5.1 (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 $n \in \mathbb{N}$:

$$\left| \Pr_{z \leftarrow G_0(n)} [\mathcal{A}(z) = 1] - \Pr_{z \leftarrow G_1(n)} [\mathcal{A}(z) = 1] \right| \leq \mu(n)$$

Theorem 5.1. *One-way puzzles (Definition 3.2) imply quantum weak pseudoentropy generators (Definition 5.1).*

The rest of this section is devoted to the proof of this theorem. For the following discussion, fix some sufficiently large $n \in \mathbb{N}$. In what follows, we will sometimes drop explicit parameterization on n when it is clear from the context.

5.1 Defining G_0 and G_1

In order to build QWPEG, we first need to define the two distributions, G_0 and G_1 .

The Distribution G_0 . The distribution G_0 is defined as:

$$G_0(n) := s, h, i, r, h(k)_i, \langle k, r \rangle$$

where $(k, s) \leftarrow \text{Samp}(1^n)$, $h \in \{0, 1\}^\ell$ is a (uniform) seed for a pairwise independent hash function mapping $\{0, 1\}^n$ bits to $\{0, 1\}^{3n}$ bits, $i \leftarrow [3n]$, and $r \leftarrow \{0, 1\}^n$. Moreover, $h(k)_i$ denotes the output of $h(k)$ truncated to the first i bits.

The Distribution G_1 . The distribution G_1 is defined as:

$$G_1(n) := \begin{cases} s, h, i, r, h(k)_i, u & \text{if } i = i_s, k \in \mathbb{G}_s, (h, h(k)_i) \in \mathbb{F}_s \\ s, h, i, r, h(k)_i, \langle k, r \rangle & \text{otherwise} \end{cases}$$

where $(k, s) \leftarrow \text{Samp}(1^n)$, $h \in \{0, 1\}^\ell$ is a (uniform) seed for a pairwise independent hash function mapping n bits to $3n$ bits, $i \leftarrow [3n]$, u is a uniform bit, and $r \leftarrow \{0, 1\}^n$.

It remains to define i_s and the sets $\mathbb{G}_s, \mathbb{F}_s$, which we do next. First, we define some notation.

- Let $\mathbb{K} = \text{Supp}(K)$, where K is the marginal distribution on keys induced by $\text{Samp}(1^n)$.
- Let $\mathbb{S} = \text{Supp}(S)$, where S is the marginal distribution on puzzles induced by $\text{Samp}(1^n)$.
- For any $s \in \mathbb{S}$, denote by K_s the distribution on keys output by Samp conditioned on the puzzle being s , and let $\mathbb{K}_s = \text{Supp}(K_s)$.
- For any $k \in \mathbb{K}$, denote by S_k the distribution on puzzles output by Samp conditioned on the key being k , and let $\mathbb{S}_k = \text{Supp}(S_k)$.

Defining \mathbb{G}_s and i_s . We claim the existence of a dense enough set \mathbb{G}_s of preimage keys of every puzzle s , where all keys in \mathbb{G}_s have (roughly) the same sample entropy in distribution K_s . We call this sample entropy j_s , and we will later use it to define i_s .

Claim 5.1. *For all $n \in \mathbb{N}$ and all $s \in \mathbb{S}$, there exists a set $\mathbb{G}_s \subseteq \mathbb{K}_s$ and $j_s \in [0, 2n - 1]$ such that*

1. $\Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s] \geq 1/3n$
2. For all $k \in \mathbb{G}_s$, we have $j_s \leq H_{K_s}(k) \leq j_s + 1$.

Proof. We first show that most of the probability mass in K_s is on keys with sample entropies less than $2n$. We fix some arbitrary $s \in \mathbb{S}$ for the following discussion.

$$\begin{aligned}
\Pr_{k \leftarrow K_s} [H_{K_s}(k) \geq 2n] &= \sum_{k: H_{K_s}(k) \geq 2n} \Pr_{K_s}[k] \\
&= \sum_{k: \Pr_{K_s}(k) \leq 2^{-2n}} \Pr_{K_s}[k] \\
&\leq \sum_{k: \Pr_{K_s}(k) \leq 2^{-2n}} 2^{-2n} \\
&\leq 2^n \cdot 2^{-2n} \leq 2^{-n}
\end{aligned}$$

Now, we can divide these keys into sets with similar sample entropies. For $j \in [0, 2n - 1]$, let $\mathbb{C}_j := \{k \in K_s : j \leq H_{K_s}(k) \leq j + 1\}$. We now show that with overwhelming probability, k sampled from K_s will belong to some \mathbb{C}_j .

$$\begin{aligned}
\Pr_{k \leftarrow K_s} [\exists j \in [0, 2n - 1] \text{ s.t. } k \in \mathbb{C}_j] &= \Pr_{k \leftarrow K_s} [0 \leq H_{K_s}(k) \leq 2n] \\
&\geq 1 - 2^{-n}
\end{aligned}$$

Then the Pigeonhole Principle says there must exist some $j_s \in [0, 2n - 1]$ such that

$$\Pr_{k \leftarrow K_s} [k \in \mathbb{C}_{j_s}] \geq \frac{1 - 2^{-n}}{2n} \geq 1/3n$$

Set $\mathbb{G}_s := \mathbb{C}_{j_s}$. This completes the proof. \square

For every $s \in \mathbb{S}$, fix \mathbb{G}_s, j_s to be an arbitrary set and index satisfying the above claim. Set $i_s := j_s + 600 \log(n)$. Note that since $j_s \leq 2n$, $i_s \in [3n]$ for sufficiently large n . Furthermore, let

$$\mathbb{A}_s := \{k \in K_s : H_{K_s}(k) \leq j_s + 500 \log(n)\}$$

We will now proceed to defining \mathbb{F}_s .

Defining \mathbb{F}_s . We now claim the existence of a dense enough set \mathbb{F}_s of h, y values such that (1) the preimage distribution of h, y in K_s (i.e. K_s conditioned on $h(K_s)_{i_s} = y$) has reasonably high probability mass in \mathbb{G}_s and (2) there is at most one key $k \in \mathbb{A}_s$ that hashes (under h) to y . Formally, we have the following claim.

Claim 5.2. *For all $n \in \mathbb{N}$ and all $s \in \mathbb{S}$, there exists a set \mathbb{F}_s such that*

1. $\Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ k \leftarrow K_s}} [h, h(k)_{i_s} \in \mathbb{F}_s] \geq 1/6n - 1/n^{100}$
2. *For all $h, y \in \mathbb{F}_s$, the following holds:*
 $\Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s | h(k)_{i_s} = y] \geq 1/6n$ and there exists at most one $k \in \mathbb{A}_s$ such that $h(k)_{i_s} = y$

Proof. We prove the claim in two parts. The following claim shows that with high probability over the randomness of sampling (h, k) , there exists at most one $k' \in \mathbb{A}_s$ such that $h(k')_{i_s} = h(k)_{i_s}$.

SubClaim 5.1. *For all $s \in \mathbb{S}$,*

$$\Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ k \leftarrow K_s}} [\exists k' \in \mathbb{A}_s \text{ s.t. } k' \neq k \text{ and } h(k')_{i_s} = h(k)_{i_s}] \leq 1/n^{100}$$

Proof. We first show a bound on the size of \mathbb{A}_s . By definition of \mathbb{A}_s , for all $k \in \mathbb{A}_s$, we have that $H_{K_s}(k) \leq j_s + 500 \log n$. This implies

$$\begin{aligned} \Pr_{K_s}[k] &\geq \frac{1}{2^{j_s+500 \log n}} \\ \implies \Pr_{k' \leftarrow K_s}[k' \in \mathbb{A}_s] &\geq \frac{|\mathbb{A}_s|}{2^{j_s+500 \log n}} \\ \implies |\mathbb{A}_s| &\leq 2^{j_s+500 \log n} \end{aligned}$$

By the properties of universal hash functions, for all $s \in \mathbb{S}$, for all k, k' such that $k \neq k'$

$$\Pr_{h \leftarrow \{0,1\}^\ell}[h(k')_{i_s} = h(k)_{i_s}] = \frac{1}{2^{i_s}}$$

Applying a union bound to the set $\mathbb{A}_s \setminus \{k\}$

$$\Pr_{h \leftarrow \{0,1\}^\ell}[\exists k' \in \mathbb{A}_s \setminus \{k\} \text{ s.t. } h(k')_{i_s} = h(k)_{i_s}] \leq \frac{|\mathbb{A}_s|}{2^{i_s}}$$

Recall that i_s is defined as $j_s + 600 \log n$ and $|\mathbb{A}_s| \leq 2^{j_s+500 \log n}$. Substituting these above gives

$$\Pr_{h \leftarrow \{0,1\}^\ell}[\exists k' \in \mathbb{A}_s \setminus \{k\} \text{ s.t. } h(k')_{i_s} = h(k)_{i_s}] \leq \frac{1}{n^{100}}$$

which proves this subclaim. \square

In the next subclaim we show that for some noticeable fraction of $(h, h(k)_{i_s})$, the induced preimage distribution on keys has noticeable probability mass in \mathbb{G}_s .

SubClaim 5.2. For all $s \in \mathbb{S}$,

$$\Pr_{h \leftarrow \{0,1\}^\ell} \left[\Pr_{k' \leftarrow K_s}[k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] \geq 1/6n \right] \geq 1/6n$$

Proof. Consider the following inequality: For all $s \in \mathbb{S}$,

$$\Pr_{\substack{k \leftarrow K_s \\ h \leftarrow \{0,1\}^\ell \\ k' \leftarrow K_s}}[k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] \geq 1/3n$$

It suffices to prove the above inequality, since the statement of our subclaim follows from it by a Markov argument. Note that

$$\Pr[k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] = \frac{\Pr[(k' \in \mathbb{G}_s) \wedge (h(k')_{i_s} = h(k)_{i_s})]}{\Pr[h(k')_{i_s} = h(k)_{i_s}]}$$

where the probabilities are over $k \leftarrow K_s, h \leftarrow \{0,1\}^\ell, k' \leftarrow K_s$.

By applying the properties of universal hashing to the RHS,

$$\begin{aligned} \Pr[k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] &= \Pr_{k' \leftarrow K_s}[k' \in \mathbb{G}_s] \cdot \frac{1/2^{i_s}}{1/2^{i_s}} \\ &= \Pr_{k' \leftarrow K_s}[k' \in \mathbb{G}_s] \end{aligned}$$

By Claim 5.1 for all $s \in \mathbb{S}$,

$$\Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s] \geq 1/3n$$

which gives

$$\Pr[k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] \geq 1/3n$$

□

We now combine these subclaims to define \mathbb{F}_s . For any events A and B , $\Pr[A \wedge B] \geq \Pr[A] - \Pr[\neg B]$. Applying this to the above two subclaims, for all $s \in \mathbb{S}$,

$$\Pr_{\substack{k \leftarrow K_s \\ h \leftarrow \{0,1\}^\ell}} \left[\left(\Pr_{k' \leftarrow K_s} [k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] \geq 1/6n \right) \wedge (\forall k' \in \mathbb{A}_s \setminus \{k\}, h(k')_{i_s} \neq h(k)_{i_s}) \right] \geq 1/6n - 1/n^{100}$$

which implies

$$\Pr_{\substack{k \leftarrow K_s \\ h \leftarrow \{0,1\}^\ell}} \left[\Pr_{k' \leftarrow K_s} [k' \in \mathbb{G}_s | h(k')_{i_s} = h(k)_{i_s}] \geq 1/6n \wedge (\exists \text{ at most one } k' \in \mathbb{A}_s \text{ s.t. } h(k')_{i_s} = h(k)_{i_s}) \right] \geq 1/6n - 1/n^{100}$$

Define \mathbb{F}_s as follows:

$$\mathbb{F}_s := \left\{ h, y \mid \begin{array}{l} h \in \{0,1\}^\ell \wedge y \in \{0,1\}^{i_s} \wedge \\ \Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s | h(k)_{i_s} = y] \geq 1/6n \\ \wedge (\exists \text{ at most one } k \in \mathbb{A}_s \text{ s.t. } h(k)_{i_s} = y) \end{array} \right\}$$

which implies

$$\Pr_{\substack{k \leftarrow K_s \\ h \leftarrow \{0,1\}^\ell}} [h, h(k)_{i_s} \in \mathbb{F}_s] \geq 1/6n - 1/n^{100}$$

which concludes the proof of our claim. □

This completes a description of G_0 and G_1 . It is easy to see that $G_0(n)$ is efficiently sampleable and that G_0 and G_1 are polynomially bounded in length, i.e. there exists some polynomial $p(\cdot)$ such that for all $z_0 \in \text{Supp}(G_0(n))$ and for all $z_1 \in \text{Supp}(G_1(n))$, $|z_0| = |z_1| \leq p(n)$. Next, we prove that $G_1(n)$ has inverse-polynomially higher Shannon entropy than $G_0(n)$.

5.2 Establishing an Entropy Gap

Lemma 5.1. *For all sufficiently large $n \in \mathbb{N}$,*

$$H(G_1(n)) - H(G_0(n)) \geq 1/n^5$$

The rest of this subsection is dedicated to a proof of the above lemma. We will make use of some intermediate lemmas and claims, that we state below.

First, we state a standalone lemma described below, the proof of which is in Appendix B. Intuitively, the lemma defines a special distribution X over $\{0,1\}^n$ that samples a single element x^* with much higher probability than all other elements $x \neq x^*$. It then finds a lower bound on the difference in Shannon entropy between distributions A_0 and A_1 , where $A_0 = R, \langle X, R \rangle$ and A_1 is identical to A_0 except replacing $\langle X, R \rangle$ with a uniform random bit when $X = x^*$ (and where R is the uniform distribution on $\{0,1\}^n$). The proof makes use of the leftover hash lemma to show that the two distributions are close to uniform (and hence have high entropy) conditioned on $X \neq x^*$. Thus, a difference in entropies arises from the case of $X = x^*$.

Lemma 5.2. For sufficiently large $n \in \mathbb{N}$, let X be a distribution over $\{0, 1\}^n$ such that there exists $x^* \in \{0, 1\}^n$ s.t.

1. $\Pr_{x \leftarrow X}[x = x^*] \geq 1/6n$
2. $\forall x' \neq x^*, \Pr_{x \leftarrow X}[x = x'] \leq 2/n^{600}$

Let R be uniformly distributed on $\{0, 1\}^n$, and U_1 denote the uniform distribution over $\{0, 1\}$. Define α_0 and α_1 as follows.

$$\alpha_0(x, r) := \langle x, r \rangle$$

$$\alpha_1(x, r) := \begin{cases} U_1 & \text{if } x = x^* \\ \langle x, r \rangle & \text{otherwise} \end{cases}$$

For $b \in \{0, 1\}$, define distribution $A_b := R, \alpha_b(X, R)$. Then

$$H(A_1) - H(A_0) \geq 1/100n^2$$

To prove lemma 5.1, we will also do the following. For each fixing of $s \in \mathbb{S}$, we will establish an entropy gap between the distributions G_0 and G_1 when conditioned on sampling (k, h, i) such that $i = i_s$ and $h, h(k)_{i_s} \in \mathbb{F}_s$. To do this, we will use the following claim, which helps establish a difference in entropy in the last bit of G_0 and G_1 for every fixing of $s \in \mathbb{S}$ and $(h, y) \in \mathbb{F}_s$.

Claim 5.3. Fix any $s \in \mathbb{S}$ and any $(h, y) \in \mathbb{F}_s$. Let R be uniformly distributed on $\{0, 1\}^n$.

For any $r \in \{0, 1\}^n$, define distributions P_r, Q_r as follows: First sample $k \leftarrow K_s$ conditioned on $h(k)_{i_s} = y$. Set

$$P_r = \begin{cases} U_1 & \text{if } k \in \mathbb{G}_s \\ \langle k, r \rangle & \text{otherwise} \end{cases}$$

and

$$Q_r = \langle k, r \rangle$$

Then,

$$H(R, P_R) - H(R, Q_R) \geq 1/100n^2$$

Before proceeding, we provide a complete proof of Claim 5.3.

Proof. (of Claim 5.3) We first require some claims about the preimage distribution on keys in \mathbb{K}_s such that $h(k)_{i_s} = y$. These will set us up to use Lemma 5.2.

First, we show that there is a unique key $k^* \in \mathbb{G}_s$ that hashes to y .

SubClaim 5.3. There exists a unique $k^* \in \mathbb{G}_s$ such that $h(k^*)_{i_s} = y$.

Proof. By definition of \mathbb{F}_s ,

$$\Pr_{k \leftarrow K_s}[k \in \mathbb{G}_s | h(k)_{i_s} = y] \geq 1/6n$$

Thus, there must exist at least one $k^* \in \mathbb{G}_s$ such that $h(k^*)_{i_s} = y$.

By definition, $\mathbb{G}_s \subseteq \mathbb{A}_s$. By Claim 5.2, there exists atmost one $k \in \mathbb{A}_s$ such that $h(k)_{i_s} = y$. Therefore, k^* is unique. \square

Next we show that since all other preimage keys are outside of \mathbb{A}_s , they have significantly lower sampling probabilities than k^* .

SubClaim 5.4. For all $k' \in \{0, 1\}^n$ such that $k' \neq k^*$

$$\Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y] \leq 2/n^{600}$$

Proof. By Claim 5.2, there is at most one $k \in \mathbb{A}_s$ such that $h(k)_{i_s} = y$. Note that $k^* \in \mathbb{G}_s \subseteq \mathbb{A}_s$.

Fix any $k' \neq k^*$ such that $h(k')_{i_s} = y$ (the subclaim follows trivially if $h(k')_{i_s} \neq y$). It must therefore be the case that $k' \notin \mathbb{A}_s$. By definition of \mathbb{A}_s ,

$$H_{K_s}(k') \geq j_s + 600 \log n$$

or equivalently,

$$\Pr_{K_s}[k'] \leq \frac{1}{n^{600} \cdot 2^{j_s}}$$

We aim to bound $\Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y]$. We do so by noting that the ratio of probabilities of k^* and k' is unaffected by conditioning on $h(k)_{i_s} = y$.

$$\begin{aligned} \Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y] &= \frac{\Pr_{k \leftarrow K_s} [k = k']}{\Pr_{k \leftarrow K_s} [h(k)_{i_s} = y]} \\ \Pr_{k \leftarrow K_s} [k = k^* | h(k)_{i_s} = y] &= \frac{\Pr_{k \leftarrow K_s} [k = k^*]}{\Pr_{k \leftarrow K_s} [h(k)_{i_s} = y]} \end{aligned}$$

which implies

$$\frac{\Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y]}{\Pr_{k \leftarrow K_s} [k = k^* | h(k)_{i_s} = y]} = \frac{\Pr_{k \leftarrow K_s} [k = k']}{\Pr_{k \leftarrow K_s} [k = k^*]}$$

Since $k^* \in \mathbb{G}_s$, by Claim 5.1

$$H_{K_s}(k^*) \leq j_s + 1 \implies \Pr_{K_s}[k^*] \geq 1/2^{j_s+1}$$

Together with the previous bound on the probability of k' this implies

$$\frac{\Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y]}{\Pr_{k \leftarrow K_s} [k = k^* | h(k)_{i_s} = y]} = \frac{\Pr_{k \leftarrow K_s} [k = k']}{\Pr_{k \leftarrow K_s} [k = k^*]} \leq \frac{2^{j_s+1}}{n^{600} \cdot 2^{j_s}} \leq \frac{2}{n^{600}}$$

Since probabilities are upper bounded by one, this proves the subclaim. \square

We may now apply Lemma 5.2 with $X := (K_s | h(K_s)_{i_s} = y)$ and $x^* := k^*$. By Claim 5.2

$$\Pr_{k \leftarrow K_s} [k \in \mathbb{G}_s | h(k)_{i_s} = y] \geq 1/6n$$

Since k^* is the only inverse in \mathbb{G}_s , this implies

$$\Pr_{k \leftarrow K_s} [k = k^* | h(k)_{i_s} = y] \geq 1/6n$$

Moreover, for all $k' \in \{0, 1\}^n$ such that $k' \neq k^*$, $\Pr_{k \leftarrow K_s} [k = k' | h(k)_{i_s} = y] \leq 2/n^{600}$. Thus, Lemma 5.2 implies that

$$H(R, P_R) - H(R, Q_R) \geq 1/100n^2$$

which proves Claim 5.3. \square

The next claim relates the entropy gap between arbitrary joint distributions (A, B_0) and (A, B_1) to their entropy gap conditioned on an event.

Claim 5.4. Let $A, B_0,$ and B_1 be random variables and let $\mathbb{A} := \text{Supp}(A)$. Additionally, let $\mathbb{A}^* \subseteq \mathbb{A}$ be a set with the following properties:

1. There exists some $d \geq 0$ such that for all $a \in \mathbb{A}^*$, $H(B_1|A = a) - H(B_0|A = a) \geq d$.
2. For all $a \in \mathbb{A} \setminus \mathbb{A}^*$, $H(B_1|A = a) = H(B_0|A = a)$

Then the following holds

$$H(A, B_1) - H(A, B_0) \geq d \cdot \Pr[A \in \mathbb{A}^*]$$

Proof. By the chain rule, for $b \in \{0, 1\}$

$$H(A, B_b) = H(B_b|A) + H(A)$$

Applying this to the difference in entropies

$$\begin{aligned} H(A, B_1) - H(A, B_0) &= H(B_1|A) + H(A) - H(B_0|A) - H(A) \\ &= H(B_1|A) - H(B_0|A) \\ &= \sum_{a \in \mathbb{A}} \Pr_A[a] \left(H(B_1|A = a) - H(B_0|A = a) \right) \end{aligned}$$

We may now split the sum into terms where $a \in \mathbb{A}^*$ and terms where $a \in \mathbb{A} \setminus \mathbb{A}^*$, and apply the properties of \mathbb{A}^* .

$$\begin{aligned} H(A, B_1) - H(A, B_0) &= \sum_{a \in \mathbb{A}^*} \Pr_A[a] \left(H(B_1|A = a) - H(B_0|A = a) \right) \\ &\quad + \sum_{a \in \mathbb{A} \setminus \mathbb{A}^*} \Pr_A[a] \left(H(B_1|A = a) - H(B_0|A = a) \right) \\ &= \sum_{a \in \mathbb{A}^*} \Pr_A[a] \left(H(B_1|A = a) - H(B_0|A = a) \right) \\ &\geq \sum_{a \in \mathbb{A}^*} \Pr_A[a] \cdot d \\ &= d \cdot \Pr[A \in \mathbb{A}^*] \end{aligned}$$

which concludes the proof. □

We are now ready to prove the main lemma of this subsection.

Proof. (of Lemma 5.1) Let H be uniformly distributed over $\{0, 1\}^\ell$, I be uniformly distributed over $[3n]$, and R be uniformly distributed over $\{0, 1\}^n$. Recall that

$$G_0 = (S_K, H, I, R, H(K)_I, Q_{K,R})$$

and

$$G_1 = (S_K, H, I, R, H(K)_I, P_{K,R})$$

where

$$Q_{K,R} := \langle K, R \rangle$$

and

$$P_{K,R} = \begin{cases} U_1 & \text{if } K \in \mathbb{G}_{S_K} \text{ and } I = i_{S_K} \text{ and } h, h(K)_I \in \mathbb{F}_{S_K} \\ \langle K, R \rangle & \text{otherwise} \end{cases}$$

First, we note that for any $s \in \mathbb{S}$ and any $(h, y) \in \mathbb{F}_s$, the distribution $(R, P_{K,R} | S_K = s, H = h, H(K)_I = y, I = i_s)$ is identical to the distribution (R, P_R) as defined in Claim 5.3. Similarly, $(R, Q_{K,R} | S_K = s, H = h, H(K)_I = y, I = i_s)$ is identical to the distribution (R, Q_R) as defined in Claim 5.3. Applying the claim therefore gives

$$\begin{aligned} & \mathbb{H}(R, P_{K,R} | S_K = s, H = h, H(K)_I = y, I = i_s) \\ & \quad - \mathbb{H}(R, Q_{K,R} | S_K = s, H = h, H(K)_I = y, I = i_s) \geq 1/100n^2 \end{aligned}$$

Additionally, note that if $H, H(K)_I \notin \mathbb{F}_S$ or $I \neq i_S$, the distributions are identical. We now apply Claim 5.4. Setting A to be $(S_K, H, H(K)_I, I)$, B_1 to $(R, P_{K,R})$, B_0 to $(R, Q_{K,R})$, and \mathbb{A}^* to the set $\{(s, h, y, i_s) : s \in \mathbb{S} \wedge (h, y) \in \mathbb{F}_s\}$, we get

$$\mathbb{H}(S_K, H, I, R, H(K)_I, P_{K,R}) - \mathbb{H}(S_K, H, I, R, H(K)_I, Q_{K,R}) \geq \frac{\Pr[I = i_S \wedge (H, H(K)_I) \in \mathbb{F}_S]}{100n^2}$$

which may be written as

$$\mathbb{H}(G_0) - \mathbb{H}(G_1) \geq \Pr[I = i_S \wedge (H, H(K)_I) \in \mathbb{F}_S] / 100n^2$$

Since I is uniform over $[3n]$ and $i_S \in [3n]$,

$$\mathbb{H}(G_0) - \mathbb{H}(G_1) \geq \Pr[(H, H(K)_{i_S}) \in \mathbb{F}_S] / 300n^3$$

By Claim 5.2, for every $s \in \mathbb{S}$, $\Pr[H, H(K)_{i_s} \in \mathbb{F}_s] \geq 1/6n - 1/n^{100}$. Therefore for large enough $n \in \mathbb{N}$

$$\begin{aligned} \mathbb{H}(G_0) - \mathbb{H}(G_1) & \geq (1/6n - 1/n^{100}) \cdot (1/300n^3) \\ & \geq (1/7n) \cdot (1/300n^3) \\ & \geq \frac{1}{2100n^4} > \frac{1}{n^5} \end{aligned}$$

which concludes the proof. \square

5.3 Establishing Computational Indistinguishability

Lemma 5.3. *There exists a negligible function μ such that for every quantum polynomial sized adversary \mathcal{A} , for all $n \in \mathbb{N}$:*

$$\left| \Pr_{z \leftarrow G_0(n)}[\mathcal{A}(z) = 1] - \Pr_{z \leftarrow G_1(n)}[\mathcal{A}(z) = 1] \right| \leq \mu(n)$$

Proof. We prove this lemma by contradiction. Suppose there exists some quantum polynomial-sized \mathcal{A} and an inverse polynomial function ε such that for infinitely many $n \in \mathbb{N}$,

$$\left| \Pr_{z \leftarrow G_0(n)}[\mathcal{A}(z) = 1] - \Pr_{z \leftarrow G_1(n)}[\mathcal{A}(z) = 1] \right| \geq \varepsilon(n)$$

For every s , G_0 and G_1 are identical whenever $i \neq i_s$ or $k \notin \mathbb{G}_s$. As a result, the case where this constraint is met must give rise to all the distinguishing advantage. By definition of G_0 and G_1 , this means that

$$\begin{aligned} & \left| \Pr[\mathcal{A}(s, h, i_s, r, h(k)_{i_s}, \langle k, r \rangle) = 1 \mid k \in \mathbb{G}_s] \right. \\ & \quad \left. - \Pr[\mathcal{A}(s, h, i_s, r, h(k)_{i_s}, b) = 1 \mid k \in \mathbb{G}_s] \right| \\ & \geq \varepsilon(n) \end{aligned} \tag{5}$$

where the probability is over $k, s \leftarrow (K, S_K), b \leftarrow \{0, 1\}, h \leftarrow \{0, 1\}^\ell, r \leftarrow \{0, 1\}^n$.

Let \mathcal{A}' be a *predictor* that on input (s, h, i, r, y) samples a uniform bit b and outputs $b \oplus \mathcal{A}(s, h, i, r, y, b)$. Then the previous equation implies:

$$\Pr [\mathcal{A}'(s, h, i_s, r, h(k)_{i_s}) = \langle k, r \rangle \mid k \in \mathbb{G}_s] \geq 1/2 + \varepsilon(n) \quad (6)$$

Note that \mathcal{A}' is a quantum polynomial-sized circuit family that guesses $\langle k, r \rangle$ with noticeable probability, given $(s, h, i_s, h(k)_{i_s})$ as an auxiliary function of the key k , along with r , and when conditioned on $k \in \mathbb{G}_s$. Therefore, by the Goldreich-Levin theorem (Theorem 3.2) applied to \mathcal{A}' , there exists a quantum polynomial-sized circuit family \mathcal{B} and an inverse polynomial function ε' such that:

$$\Pr [\mathcal{B}(s, h, i_s, h(k)_{i_s}) = k \mid k \in \mathbb{G}_s] \geq \varepsilon'(n) \quad (7)$$

Our goal is to prove the existence of a quantum polynomial-sized algorithm that outputs inverses k given a puzzle s . In the above equation, however, \mathcal{B} requires $(h, i_s, h(k)_{i_s})$ as input in addition to the puzzle s . The next claim proves that \mathcal{B} will output an inverse of the puzzle with non-negligible probability even when i_s and $h(k)_{i_s}$ are replaced with uniform strings.

Claim 5.5. *There exists an inverse polynomial function ε'' such that for infinitely many $n \in \mathbb{N}$:*

$$\Pr [\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i, u), s)] \geq \varepsilon''(n)$$

where the probability is over $s \leftarrow S, i \leftarrow [3n], h \leftarrow \{0, 1\}^\ell, u \leftarrow \{0, 1\}^i$.

Proof. To prove this claim, we will first define a statistical test using the input-output behavior of \mathcal{B} .

SubClaim 5.5. *Define test \mathcal{T} as follows.*

$$\mathcal{T}(s, h, y) := \begin{cases} 1 & \text{if } \Pr[\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i_s, y), s)] \geq \varepsilon'(n)/4 \\ 0 & \text{otherwise} \end{cases}$$

where the probability is over the randomness of \mathcal{B} and Ver . Then,

$$\Pr_{(k, s) \leftarrow (K, S), h \leftarrow \{0, 1\}^\ell} [\mathcal{T}(s, h, h(k)_{i_s}) = 1 \mid k \in \mathbb{G}_s] \geq \varepsilon'(n)/4$$

Proof. (of SubClaim 5.5) For any events A and B , $\Pr[A \wedge B] \geq \Pr[A] - \Pr[\neg B]$. Applying this to the correctness of the puzzle

$$\begin{aligned} & \Pr_{\substack{k, s \leftarrow (K, S_K) \\ h \leftarrow \{0, 1\}^\ell}} [\mathcal{B}(s, h, i_s, h(k)_{i_s}) = k \wedge \top \leftarrow \text{Ver}(k, s) \mid k \in \mathbb{G}_s] \\ & \geq \varepsilon'(n) - \Pr_{k, s \leftarrow (K, S_K)} [\perp \leftarrow \text{Ver}(k, s) \mid k \in \mathbb{G}_s] \\ & \geq \varepsilon'(n) - \frac{\Pr_{k, s \leftarrow (K, S_K)} [\perp \leftarrow \text{Ver}(k, s)]}{\Pr_{k, s \leftarrow (K, S_K)} [k \in \mathbb{G}_s]} \\ & = \varepsilon'(n) - 3n \cdot \text{negl}(n) > \varepsilon'(n)/2 \end{aligned}$$

where the last step follows due to Claim 5.1, and due to the correctness of the one-way puzzle. Thus,

$$\Pr [\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i_s, h(k)_{i_s}), s) \mid k \in \mathbb{G}_s] \geq \varepsilon'(n)/2$$

and by a Markov argument

$$\Pr [\Pr[\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i_s, h(k)_{i_s}), s)] \geq \varepsilon'(n)/4 \mid k \in \mathbb{G}_s] \geq \varepsilon'(n)/4$$

which proves this subclaim. \square

Let $\ell^* := 600 \log n + \log 3n + 2 \log(8/\varepsilon'(n))$. Fix any $s \in \mathbb{S}$, and define $\ell_s^* = \min(i_s, \ell^*)$. Let K'_s denote the distribution of the random variable $(K_s | K_s \in \mathbb{G}_s)$. In what follows, we will bound the distance between $h(k)_{i_s - \ell_s^*}$ and uniform.

SubClaim 5.6. For every $s \in \mathbb{S}$,

$$\text{SD}\left((H, H(K'_s)_{i_s - \ell_s^*}), (H, U_{i_s - \ell_s^*})\right) \leq \varepsilon'(n)/8$$

Proof. For all $s \in \mathbb{S}$, for all $k^* \in \mathbb{G}_s$,

$$\begin{aligned} \Pr_{k \leftarrow K'_s}[k = k^*] &= \Pr_{k \leftarrow K_s}[k = k^* | k \in \mathbb{G}_s] \\ &= \Pr_{k \leftarrow K_s}[k = k^*] / \Pr_{k \leftarrow K_s}[k \in \mathbb{G}_s] \end{aligned}$$

By Claim 5.1, for all $s \in \mathbb{S}$, $\Pr_{k \leftarrow K_s}[k \in \mathbb{G}_s] \geq 1/3n$. Additionally, for all $k^* \in \mathbb{G}_s$, $H_{K_s}(k^*) \geq j_s \implies \Pr_{K_s}[k^*] \leq 2^{-j_s}$.

Therefore for all $s \in \mathbb{S}$, for all $k^* \in \mathbb{G}_s$,

$$\begin{aligned} \Pr_{k \leftarrow K'_s}[k = k^*] &= \Pr_{k \leftarrow K_s}[k = k^*] / \Pr_{k \leftarrow K_s}[k \in \mathbb{G}_s] \\ &\leq 2^{-j_s} \cdot 3n \end{aligned}$$

Thus for $s \in \mathbb{S}$, for $k^* \in \mathbb{G}_s$, $H_{K'_s}(k^*) \geq j_s - \log 3n$. In other words, $H_{\min}(K'_s) \geq j_s - \log 3n$.

By the Leftover Hash Lemma (Theorem 3.1),

$$\begin{aligned} \text{SD}\left((H, H(K'_s)_{i_s - \ell_s^*}), (H, U_{i_s - \ell_s^*})\right) &\leq 2^{-0.5(H_{\min}(K'_s) - i_s + \ell_s^*)} \\ &= 2^{-0.5(j_s - \log 3n - i_s + 600 \log n + \log 3n + 2 \log(8/\varepsilon'(n)))} \\ &= 2^{-\log(8/\varepsilon'(n))} = \varepsilon'(n)/8 \end{aligned}$$

which proves the claim. □

Next, relaxing the probability bound from SubClaim 5.5 gives us

$$\Pr \left[\exists \gamma \in \{0, 1\}^{\ell_s^*} \text{ s.t. } \mathcal{T}(s, h, (h(k)_{i_s - \ell_s^*} \| \gamma)) = 1 \mid k \in \mathbb{G}_s \right] \geq \varepsilon'(n)/4$$

Defining \mathcal{T}' as:

$$\mathcal{T}'(s, h, y) := \begin{cases} 1 & \text{if } \exists \gamma \in \{0, 1\}^{\ell_s^*} \text{ s.t. } \mathcal{T}(s, h, (y \| \gamma)) = 1 \\ 0 & \text{otherwise} \end{cases}$$

we have

$$\Pr_{(k,s) \leftarrow (K, S_K), h \leftarrow \{0,1\}^\ell} [\mathcal{T}'_{\ell^*}(s, h, h(k)_{i_s - \ell_s^*}) = 1 \mid k \in \mathbb{G}_s] \geq \varepsilon'(n)/4$$

Since no statistical test can distinguish between two distributions with advantage better than their statistical distance, replacing the hash with a uniform string in the equation above, and invoking SubClaim 5.6 gives us

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s - \ell_s^*}}} [\mathcal{T}'(s, h, u) = 1 \mid k \in \mathbb{G}_s] \geq \frac{\varepsilon'(n)}{4} - \frac{\varepsilon'(n)}{8} = \frac{\varepsilon'(n)}{8}$$

which by definition of \mathcal{T}' implies

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s - \ell^*}}} \left[\exists \gamma \in \{0,1\}^{\ell^*} \text{ s.t } \mathcal{T}(s, h, u \| \gamma) = 1 \mid k \in \mathbb{G}_s \right] \geq \frac{\varepsilon'(n)}{8}$$

If we sample γ randomly as well, we obtain

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s}}} [\mathcal{T}(s, h, u) = 1 \mid k \in \mathbb{G}_s] \geq \frac{\varepsilon'(n)}{8} \cdot 2^{-\ell^*}$$

We may now relax the constraint that $k \in \mathbb{G}_s$ as follows:

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s}}} [\mathcal{T}(s, h, u) = 1] \geq \frac{\varepsilon'(n)}{8} \cdot 2^{-\ell^*} \cdot \Pr_{k,s \leftarrow K, S_K} [k \in \mathbb{G}_s] \geq \frac{\varepsilon'(n)}{24n} \cdot 2^{-\ell^*}$$

where the last inequality follows from Claim 5.1.

Substituting $\ell^* := 600 \log n + \log 3n + 2 \log(8/\varepsilon'(n))$, we obtain

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s}}} [\mathcal{T}(s, h, u) = 1] \geq \frac{(\varepsilon'(n))^3}{24n \cdot n^{600} \cdot 3n \cdot 64} \geq \frac{(\varepsilon'(n))^3}{4608n^{602}} := \varepsilon(n)$$

By the definition of \mathcal{T} , this means

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s}}} [\Pr[\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i_s, u), s)] \geq \varepsilon'(n)/4] \geq \varepsilon(n)$$

which after undoing the Markov argument implies

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell, u \leftarrow \{0,1\}^{i_s}}} [\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i_s, u), s)] \geq \frac{\varepsilon'(n)\varepsilon(n)}{4}$$

When i is sampled uniformly from $[3n]$, $i = i_s$ occurs with probability $1/3n$. Therefore

$$\Pr_{\substack{k,s \leftarrow (K, S_K) \\ h \leftarrow \{0,1\}^\ell \\ i \leftarrow [3n], u \leftarrow \{0,1\}^i}} [\top \leftarrow \text{Ver}(\mathcal{B}(s, h, i, u), s)] \geq \frac{\varepsilon'(n)\varepsilon(n)}{12n} := \varepsilon''(n)$$

which completes the proof of the claim. □

To complete the proof of the lemma, let \mathcal{B}' denote an algorithm that on input s does the following.

- Sample $h \leftarrow \{0,1\}^\ell$
- Sample $i \leftarrow [3n]$
- Sample $u \leftarrow \{0,1\}^i$
- Output $\mathcal{B}(s, h, i, u)$

Then Claim 5.5 implies

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [\top \leftarrow \text{Ver}(\mathcal{B}'(s), s)] \geq \varepsilon''(n)$$

which contradicts security of the puzzle, as desired. □

6 Quantum Weak PEGs imply Quantum PEGs

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 6.1 (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 6.1. *Quantum weak pseudoentropy generators (Definition 5.1) imply quantum pseudoentropy generators (Definition 6.1).*

Let $\{G_0(n), G_1(n)\}_{n \in \mathbb{N}}$ be a Quantum Weak Pseudoentropy Generator with entropy gap greater than $1/n^c$ for some constant $c > 0$. By Definition 5.1, there exists a polynomial $L(\cdot)$ such that for all $z_0 \in \text{Supp}(G_0(n))$ and $z_1 \in \text{Supp}(G_1(n))$, $|z_0| = |z_1| \leq L(n)$. Let $q(n) := n^{c+3}L(n)^2$.

We construct a pseudoentropy generator $\{\overline{G}_0(n), \overline{G}_1(n)\}_{n \in \mathbb{N}}$ by simply generating $q(n)$ independent samples of $G_0(n)$ and $G_1(n)$ respectively. That is,

- $\overline{G}_0(n) := G_0(n)^{q(n)}$
- $\overline{G}_1(n) := G_1(n)^{q(n)}$

Claim 6.1. *The distribution ensembles $\{\overline{G}_0(n)\}_{n \in \mathbb{N}}$ and $\{\overline{G}_1(n)\}_{n \in \mathbb{N}}$ satisfy Definition 6.1.*

Note that for all $z_0 \in \text{Supp}(\overline{G}_0(n))$ and $z_1 \in \text{Supp}(\overline{G}_1(n))$, $|z_0| = |z_1| \leq q(n) \cdot L(n)$ and $\overline{G}_0(n)$ is efficiently sampleable since $G_0(n)$ is efficiently sampleable. Only indistinguishability and the entropy gap remain to be shown.

Computational Indistinguishability. There exists a negligible function μ_2 such that for all quantum polynomial-sized adversaries \mathcal{A} , $n \in \mathbb{N}$:

$$\left| \Pr_{z \leftarrow \overline{G}_0(n)} [1 \leftarrow \mathcal{A}(z)] - \Pr_{z \leftarrow \overline{G}_1(n)} [1 \leftarrow \mathcal{A}(z)] \right| \leq \mu_2(n)$$

Proof. We prove this by contradiction. Suppose there exists an adversary \mathcal{A} that distinguishes the two distributions with non-negligible advantage ϵ . We use this adversary to build a (non-uniform) reduction to computational indistinguishability of G_0 and G_1 , contradicting the security of the QWPEG. We proceed through a series of hybrids $H_0, H_1, \dots, H_{q(n)}$. In hybrid H_j , the adversary is run on inputs sampled from $D_j := G_0(n)^j \times G_1(n)^{q(n)-j}$. By definition, $D_0 = \overline{G_0}$ and $D_{q(n)} = \overline{G_1}$. Suppose the adversary distinguishes adjacent hybrids H_{j-1} and H_j with advantage ϵ_j . Note that by assumption, the adversary distinguishes H_0 and H_q with advantage ϵ which implies that $\epsilon \leq \sum_j \epsilon_j$. Therefore, for each sufficiently large n , there exists a j such that $\epsilon_j \geq \epsilon/q(n)$, i.e:

$$\left| \Pr_{z \leftarrow D_{j-1}} [1 \leftarrow \mathcal{A}(z)] - \Pr_{z \leftarrow D_j} [1 \leftarrow \mathcal{A}(z)] \right| \geq \epsilon/q$$

This adversary therefore distinguishes between samples from G_0 and G_1 at the j th position, given appropriate samples at the other positions. While samples can be computed from G_0 efficiently, computing G_1 is not necessarily efficient. To resolve this we show that for each sufficiently large n there must exist fixed samples for which the adversary distinguishes in the j th position, and which the reduction may receive as non-uniform advice. If we parse z as $z_1, \dots, z_{q(n)}$, by a Markov argument:

$$\Pr_{z_1^*, \dots, z_{q(n)}^* \leftarrow D_j} \left[\left| \Pr_{z_j \leftarrow G_0(n)} [1 \leftarrow \mathcal{A}(z_1^*, \dots, z_{j-1}^*, z_j, z_{j+1}^*, \dots, z_{q(n)}^*)] - \Pr_{z_j \leftarrow G_1(n)} [1 \leftarrow \mathcal{A}(z_1^*, \dots, z_{j-1}^*, z_j, z_{j+1}^*, \dots, z_{q(n)}^*)] \right| \geq \epsilon/2q \right] \geq \epsilon/2q$$

Therefore, for each sufficiently large n there exist $z_1^*, \dots, z_{j-1}^*, z_{j+1}^*, \dots, z_q^* \in \text{Supp}(G_0(n)) \cup \text{Supp}(G_1(n))$ such that:

$$\left| \Pr_{z_j \leftarrow G_0(n)} [1 \leftarrow \mathcal{A}(z_1^*, \dots, z_{j-1}^*, z_j, z_{j+1}^*, \dots, z_{q(n)}^*)] - \Pr_{z_j \leftarrow G_1(n)} [1 \leftarrow \mathcal{A}(z_1^*, \dots, z_{j-1}^*, z_j, z_{j+1}^*, \dots, z_{q(n)}^*)] \right| \geq \epsilon/2q$$

We may now build a reduction R that receives as input z and takes as non-uniform advice τ defined as:

$$\tau := j, z_1^*, \dots, z_{j-1}^*, z_{j+1}^*, \dots, z_q^*$$

The reduction R on input z, τ outputs $\mathcal{A}(z_1^*, \dots, z_{j-1}^*, z, z_{j+1}^*, \dots, z_{q(n)}^*)$ and achieves a distinguishing advantage of at least $\epsilon/2q$, contradicting the security of QWPEG. \square

Entropy Gap. There exists a constant c' such that for $\epsilon(n) = 1/2^n$ and sufficiently large n :

$$H_{\min}^\epsilon(\overline{G_1}(n)) - H_{\max}^\epsilon(\overline{G_0}(n)) \geq n^{c'}$$

Proof. Setting $t = q(n), \epsilon = 2^{-n}$ and $\mathcal{U} = 2^{|\mathcal{G}_1|} \leq 2^{L(n)}$, we may apply Theorem 3.1 to obtain:

$$H_{\min}^\epsilon(\overline{G_1}(n)) \geq q(n) \cdot H(G_1(n)) - O(\sqrt{qnL^2})$$

Since $H(G_1(n)) - H(G_0(n)) \geq 1/n^c$

$$H_{\min}^\epsilon(\overline{G_1}(n)) \geq q(n) \cdot H(G_0(n)) + q(n)/n^c - O(\sqrt{qnL^2})$$

Similarly setting $t = q(n), \epsilon = 2^{-n}$ and $\mathcal{U} = 2^{|\mathcal{G}_0|} \leq 2^{L(n)}$, we may apply Theorem 3.1 to obtain:

$$H_{\max}^\epsilon(\overline{G_0}(n)) \leq q(n) \cdot H(G_0(n)) + O(\sqrt{qnL^2})$$

Taking the difference gives us

$$H_{\min}^{\epsilon}(\overline{G}_1(n)) - H_{\max}^{\epsilon}(\overline{G}_0(n)) \geq q(n)/n^c - O(\sqrt{qnL^2})$$

Substituting $q(n) = n^{2c+3}L(n)^2$

$$\begin{aligned} H_{\min}^{\epsilon}(\overline{G}_1(n)) - H_{\max}^{\epsilon}(\overline{G}_0(n)) &\geq n^{c+3}L(n)^2 - O(n^{c+2}L(n)^2) \\ &\geq n^{c+3}L(n)^2/2 \end{aligned}$$

which implies there exists some constant $c' > 0$ such that for sufficiently large n

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

which concludes the proof. \square

7 Quantum PEGs Imply Imbalanced EFI

In this section, we prove that quantum pseudoentropy generators imply a (non-uniform) variant of EFI, that we define below. This definition modifies the standard definition (Definition 3.3) 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 7.1 (s^* -Imbalanced EFI). *Let $s^*(\cdot)$ denote a 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 7.1. *There exists a 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 6.1) imply s^* -imbalanced EFI (Definition 7.1).*

Let $\{\overline{G}_0(n), \overline{G}_1(n)\}_{n \in \mathbb{N}}$ be a QPEG (Definition 6.1). As a result, there exists a constant c and some negligible smoothing parameter ϵ such that the entropy gap is greater than n^c . Additionally, there exists a polynomial $L(\cdot)$ such that the length of the generator output is upper bounded by $L(n)$. Finally, let ℓ be the key length of a universal hash function from $\{0, 1\}^{L(n)}$ to $\{0, 1\}^{L(n)}$. Let h be sampled uniformly from $\{0, 1\}^{\ell}$. We define the EFI distributions as follows.

- $\text{EFI}_s(1^n, 0) := h, h(\overline{G}_0)_s$
- $\text{EFI}_s(1^n, 1) := h, U_s$

We define $s^*(n) := H_{\max}(\bar{G}_0) + \frac{n^c}{2}$. Note that $s^*(n) \leq p(n) := L(n) + \frac{n^c}{2}$.

We will now show that EFl_s is an s^* -imbalanced EFI (Definition 7.1). It is easy to see that $\text{EFl}_s(1^n, 0)$ and $\text{EFl}_s(1^n, 1)$ are efficiently sampleable for all $s, n \in \mathbb{N}$. The proof of the theorem follows from the two lemmas below.

Lemma 7.1. (Computational Indistinguishability.) *There exists a negligible function $\mu(\cdot)$ such that for all quantum polynomial-sized adversaries \mathcal{A} , for every $n \in \mathbb{N}$ and every $s \leq s^*(n)$,*

$$\left| \Pr_{z \leftarrow \text{EFl}_s(1^n, 0)} [\mathcal{A}(z) = 1] - \Pr_{z \leftarrow \text{EFl}_s(1^n, 1)} [\mathcal{A}(z) = 1] \right| \leq \mu(n)$$

Proof. Fix some quantum polynomial-sized adversary \mathcal{A} . Consider the distribution $\{h, h(\bar{G}_1(n))_s\}$. Since \bar{G}_0 and \bar{G}_1 are computationally indistinguishable, there exists some negligible function $\mu'(\cdot)$ such that

$$\left| \Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ x \leftarrow \bar{G}_0(n)}} [\mathcal{A}(h, h(x)_s) = 1] - \Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ x \leftarrow \bar{G}_1(n)}} [\mathcal{A}(h, h(x)_s) = 1] \right| \leq \mu(n)$$

Additionally, since for sufficiently large n

$$H_{\min}^\epsilon(\bar{G}_1(n)) \geq s^*(n) + n^c/2 \geq s + n^c/2$$

by the Leftover Hash Lemma (Theorem 3.1)

$$\text{SD}(\{h, h(\bar{G}_1(n))_s\}, \{h, U_s\}) \leq \frac{1}{2^{n^c/4}} + \epsilon(n)$$

Putting both together,

$$\left| \Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ x \leftarrow \bar{G}_0(n)}} [\mathcal{A}(h, h(x)_s) = 1] - \Pr_{\substack{h \leftarrow \{0,1\}^\ell \\ x \leftarrow \{0,1\}^s}} [\mathcal{A}(h, U_s) = 1] \right| \leq \mu(n) + \frac{1}{2^{n^c/4}} + \epsilon(n)$$

which is negligible. □

Lemma 7.2. (Statistical Distance.) *There exists a negligible function $\mu(\cdot)$ such for every $n \in \mathbb{N}$ and every $s \geq s^*(n)$,*

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

Proof. Since the output of EFI is classical, trace distance is equivalent to statistical distance. We use the following sub-claim.

SubClaim 7.1. *Let X be a random variable such that $H_{\max}^\epsilon(X) \leq v$, and let d be the keylength of a universal hash function from $\{0, 1\}^{|X|}$ to $\{0, 1\}^{v+t}$ for some $t \geq 0$. Let h be sampled uniformly from $\{0, 1\}^d$. Then statistical distance between the distributions*

$$\{h, h(X)\} \text{ and } \{h, U_{v+t}\}$$

is at least $1 - 2^{-t} - \epsilon$.

Proof. Since $H_{\max}^\epsilon(X) \leq v$, there must exist a random variable X' such that $\text{SD}(X, X') \leq \epsilon$ and $H_{\max}(X') \leq v$. This means that for all $x \in X'$

$$\begin{aligned} H_{X'}(x) &\leq v \\ \implies \Pr_{X'}(x) &\geq 1/2^v \\ \implies |\text{Supp}(X')| &\leq 2^v \end{aligned}$$

Now,

$$\text{SD}(\{h, h(X')\}, \{h, U_{v+t}\}) = \mathbb{E}_h[\text{SD}(h(X'), U_{v+t})]$$

For any h , $\text{SD}(h(X'), U_{v+t})$ may be written as

$$\begin{aligned} \text{SD}(h(X'), U_{v+t}) &= \sum_{z \in \{0,1\}^{v+t}} \left| \Pr_{x \leftarrow X'}[h(x) = z] - \frac{1}{2^{v+t}} \right| \\ &= \sum_{z \in \text{Supp}(h(X'))} \left| \Pr_{x \leftarrow X'}[h(x) = z] - \frac{1}{2^{s+n}} \right| \\ &\quad + \sum_{z \in \{0,1\}^{v+t} \setminus \text{Supp}(h(X'))} \left| \Pr_{x \leftarrow X'}[h(x) = z] - \frac{1}{2^{v+t}} \right| \\ &\geq \sum_{z \in \{0,1\}^{v+t} \setminus \text{Supp}(h(X'))} \left| \Pr_{x \leftarrow X'}[h(x) = z] - \frac{1}{2^{v+t}} \right| \\ &\geq \sum_{z \in \{0,1\}^{s+n} \setminus \text{Supp}(h(X'))} \frac{1}{2^{v+t}} \\ &\geq \frac{2^{v+t} - |\text{Supp}(h(X'))|}{2^{v+t}} \\ &\geq \frac{2^{v+t} - |\text{Supp}(X')|}{2^{v+t}} \\ &\geq 1 - 2^{-t} \end{aligned}$$

Therefore

$$\text{SD}(\{h, h(X')\}, \{h, U_{v+t}\}) \geq 1 - 2^{-t}$$

By the triangle inequality

$$\text{SD}(\{h, h(X)\}, \{h, U_{v+t}\}) \geq 1 - 2^{-t} - \epsilon$$

which concludes the proof of the subclaim. \square

Note $s - n^c/2 \geq s^*(n) - n^c/2 \geq H_{\max}^\epsilon(\overline{G}_0(n))$. Therefore, setting $\overline{G}_0(n)$ as X , $s(n) - n^c/2$ as v , and $n^c/2$ as t in the above claim, the statistical distance between $\{h, h(\overline{G}_0(n))_s\}$ and $\{h, U_s\}$ is at least $1 - \epsilon(n) - 2^{-n^c/2}$ which concludes the proof of the claim. \square

Corollary 7.1. *Assume the existence of pure one-way state generators according Definition 3.1. Then there exists a function $s^*(n)$ such that*

- *There exists a polynomial p such that for all $n \in \mathbb{N}$, $|s^*(n)| \leq \log p(n)$ and*
- *There exists a s^* -imbalanced EFI (Definition 7.1).*

Proof. This follows immediately from Theorems 4.2, 5.1, 6.1 and 7.1. \square

8 Imbalanced EFI Imply Commitments

In this section, we prove that imbalanced EFI imply non-uniform commitments. First, we state the definitions of quantum commitments that we work with.

8.1 Defining Commitments

8.1.1 Quantum Bit Commitments

A non-interactive quantum bit commitment is defined by a commitment unitary Com , and has the following syntax.

- **Commit Phase.** On input bit b , the committer initializes a single bit register M to b and $\ell(n)$ -qubit auxiliary register W to $|0\rangle^{\ell(n)}$. The committer then applies a unitary Com to registers M, W and writes the resulting state on registers (C, D) , where C denotes the “commit” register, that is sent to the receiver.
- **Decommit Phase.** The committer decommits by sending the D register. The receiver applies Com^\dagger to the pair (C, D) , obtaining (M, W) . It verifies the decommitment by checking that W is in the state $|0\rangle^{\ell(n)}$ via a projective measurement. If the measurement does not succeed, it outputs \perp . Otherwise, it outputs the classical bit obtained by measuring M .

We will require these bit commitments to satisfy the following properties.

Completeness. The (honest) receiver’s output at the end of the decommit phase equals the (honest) committer’s input.

Collapse Binding for Bit Commitments. For the case of commitments to classical bits, it will sometimes be convenient to work with collapse binding [Unr14, GJMZ23], which intuitively says that an adversary cannot detect whether or not their opened message was measured. Formally, we consider the following experiment parameterized by a (malicious) committer \mathcal{A} and security parameter n .

Expmt-Binding $_{\mathcal{A},n}^{\text{Com}}$:

- \mathcal{A} (arbitrarily) prepares and outputs registers (C, D) to a challenger.
- **Validity Check.** The challenger applies Com^\dagger to (C, D) to obtain (M, W) , then projects W onto $|0\rangle^{\ell(n)}$, i.e. measures $\{|0\rangle^{\ell(n)}\langle 0|^{\ell(n)}|, \mathbb{I} - |0\rangle^{\ell(n)}\langle 0|^{\ell(n)}|\}$ on W . If the measurement rejects, the challenger sets $b' \leftarrow \{0, 1\}$ and ends the experiment. If not, the experiment continues.
- The challenger samples $b \leftarrow \{0, 1\}$ and:
 - If $b = 0$, the challenger does nothing.
 - If $b = 1$, the challenger measures M in the standard basis.
- Finally, the challenger applies Com to (M, W) to obtain (C, D) , and returns D to \mathcal{A} .
- Set b' to equal the bit output by \mathcal{A} .

We define $\text{Adv}_{\text{Com}, \mathcal{A}, n}^{\text{Expmt-Binding}^{\text{Com}}}$ as the probability that $b' = b$ in $\text{Expmt-Binding}_{\mathcal{A}, n}^{\text{Com}}$.

Definition 8.1 (Binding for Quantum Bit Commitments [Unr14, GJMZ23]). *A non-interactive quantum bit commitment satisfies statistical (resp., computational) binding if for every unbounded (resp., quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that $\text{Adv}_{\text{Com}, \mathcal{A}, n}^{\text{Expmt-Binding}} \leq \frac{1}{2} + \nu(n)$.*

Hiding for Quantum Bit Commitments. We consider the following experiment parameterized by a (malicious) receiver \mathcal{A} and security parameter n . Intuitively, this says that an adversary cannot tell whether a commitment is to zero or one.

Expmt-Hiding_{Com, A, n}:

- The challenger samples $b \leftarrow \{0, 1\}$, then sets register M to the bit b . It initializes W to $|0^{\ell(n)}\rangle$. It then applies Com to registers (M, W) writing the result on registers (C, D). It sends C to \mathcal{A} .
- Denote the output of \mathcal{A} by b' .

We define $\text{Adv}_{\text{Com}, \mathcal{A}, n}^{\text{Expmt-Hiding}}$ as the probability that $b' = b$ in Expmt-Hiding_{Com, A, n}.

Definition 8.2 (Hiding for Non-interactive Quantum Bit Commitments). *A non-interactive quantum bit commitment satisfies statistical (resp., computational) hiding if for every unbounded (resp., quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that $\text{Adv}_{\text{Com}, \mathcal{A}, n}^{\text{Expmt-Hiding}} \leq \frac{1}{2} + \nu(n)$.*

8.1.2 Imbalanced Commitments

We will also consider constructions of commitments parameterized by a classical security-parameter dependent advice string s , with the corresponding unitary denoted by Com_s . We will require these to satisfy completeness for all values of the advice string s . Furthermore, for a function s^* , we define notions of left- s^* -binding, right- s^* -binding, left- s^* -hiding and right- s^* -hiding.

Intuitively, left (resp. right) s^* -binding implies that binding holds for all values of advice $s \leq s^*$ (resp. $s \geq s^*$); and similarly for hiding.

Definition 8.3 (Imbalanced Binding for Quantum Bit Commitments). *Let $s^*(\cdot)$ denote a function. A non-interactive quantum bit commitment satisfies **left- s^* -statistical** (respectively, computational) collapse binding if for every unbounded (respectively, quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that for every $n \in \mathbb{N}$ and every $s \leq s^*(n)$, $\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Binding}} \leq \frac{1}{2} + \nu(n)$.*

*It satisfies **right- s^* -statistical** (respectively, computational) collapse binding if for every unbounded (respectively, quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that for every $n \in \mathbb{N}$ and every $s \geq s^*(n)$, $\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Binding}} \leq \frac{1}{2} + \nu(n)$.*

Definition 8.4 (Imbalanced Hiding for Quantum Bit Commitments). *Let $s^*(\cdot)$ denote a function. A non-interactive quantum bit commitment satisfies **left- s^* -statistical** (respectively, computational) hiding if for every unbounded (respectively, quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that for every $n \in \mathbb{N}$ and every $s \leq s^*(n)$, $\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Hiding}} \leq \frac{1}{2} + \nu(n)$.*

*It satisfies **right- s^* -statistical** (respectively, computational) hiding if for every unbounded (respectively, quantum polynomial-sized) adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that for every $n \in \mathbb{N}$ and every $s \geq s^*(n)$, $\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Hiding}} \leq \frac{1}{2} + \nu(n)$.*

8.2 Imbalanced EFI imply Imbalanced Commitments

We import the following lemmas from prior work: while these were proven in the uniform setting, we observe that they also carry over to the imbalanced setting with similar proofs.

Lemma 8.1. (Imported, rephrased) [Yan21, Yan22] For every function $s^*(\cdot)$, an s^* -imbalanced EFI satisfying Definition 7.1 implies an imbalanced quantum bit commitment satisfying right- s^* -statistical collapse binding according to Definition 8.3 and left- s^* -computational hiding according to Definition 8.4.

Proof. (Sketch.) Consider the purification EFI'_s of the procedure from Definition 7.1. When $s \leq s^*(n)$, the unitary $\text{Com}_s(1^n, b) := \text{EFI}'_s(1^n, b)$ satisfies computational hiding according to Definition 8.3 by inheriting the hiding from Definition 7.1.

When $s \geq s^*(n)$, then by Uhlmann's theorem there are no purification registers that can result in $\text{EFI}'_s(1^n, b)$ and $\text{EFI}'_s(1^n, b')$ having noticeable overlap. Thus the corresponding Com_s is statistical honest binding [Yan21] for $s \geq s^*(n)$, which is known [Yan22] to imply sum-binding, which in turn implies collapse-binding [DS23, GJMZ23] for the same parameters. \square

The lemma below follows from the generic transformation in [HMY23, GJMZ23] (referred to as a "flavour swap"), which can be applied to any commitment Com_s to obtain a commitment $\widetilde{\text{Com}}_s$ that is hiding whenever Com_s is binding and binding whenever Com_s is hiding.

Lemma 8.2. (Imported, rephrased) [HMY23, GJMZ23] For every function $s^*(\cdot)$, an imbalanced quantum bit commitment satisfying right- s^* -statistical collapse binding according to Definition 8.3 and left- s^* -computational hiding according to Definition 8.4 implies an imbalanced quantum bit commitment satisfying right- s^* -statistical hiding according to Definition 8.4 and left- s^* -computational collapse binding according to Definition 8.3.

We obtain the following corollary by combining the results in previous sections with the two lemmas above.

Corollary 8.1. Assume the existence of pure one-way state generators according to Definition 3.1. Then there exists a function $s^*(\cdot)$ for which there exists

1. An imbalanced quantum bit commitment satisfying right- s^* -statistical collapse binding according to Definition 8.3 and left- s^* -computational hiding according to Definition 8.4.
2. An imbalanced quantum bit commitment satisfying left- s^* -computational collapse binding according to Definition 8.3 and right- s^* -statistical hiding according to Definition 8.4.

Furthermore, there exists an explicit constant c such $s^*(\cdot)$ satisfies that for large enough $n \in \mathbb{N}$, $|s^*(n)| \leq n^c$.

Proof. This follows immediately from Corollary 7.1, Lemma 8.1 and Lemma 8.2. \square

8.3 Obtaining a Non-Uniform Hiding Commitment

Next, we obtain a commitment that is non-uniformly hiding (i.e., hiding for a given choice of non-uniform advice s^*) by combining a sequence of imbalanced commitments.

Definition 8.5 (s^* -Non-Uniform Hiding Commitment). Let $s^*(\cdot)$ denote a function. An s^* -non-uniform hiding commitment is a quantum polynomial time algorithm $\text{Com}_s(b, 1^n)$ that obtains a classical parameter-dependent advice string s , and on input $b \in \{0, 1\}$, outputs a (potentially mixed) quantum state such that:

1. **Non-Uniform Computational Hiding.** For every quantum polynomial-sized circuit $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that for every $n \in \mathbb{N}$, $\text{Adv}_{\text{Com}_{s^*(n)}, \mathcal{A}, n}^{\text{Expmt-QBinding}} \leq \frac{1}{2} + \nu(n)$.
2. **(Always) Computational Binding.** There exists a negligible function $\mu(\cdot)$ such that for every quantum polynomial-sized circuit \mathcal{A}_s , all $n, s \in \mathbb{N}$,

$$\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Binding}} \leq \frac{1}{2} + \mu(n)$$

where $\text{Adv}_{\text{Com}_s, \mathcal{A}, n}^{\text{Expmt-Binding}}$ is the probability that $b' = b$ in $\text{Expmt-Binding}_{\mathcal{A}, n}^{\text{Com}_s}$, defined above.

Theorem 8.1. Assume that there exists a function $s^*(\cdot)$ for which there exists

1. An imbalanced quantum bit commitment satisfying right- s^* -statistical collapse binding according to Definition 8.3 and left- s^* -computational hiding according to Definition 8.4.
2. An imbalanced quantum bit commitment satisfying left- s^* -computational collapse binding according to Definition 8.3 and right- s^* -statistical hiding according to Definition 8.4.

Then there exists an s^* -non-uniform hiding commitment according to Definition 8.5.

Proof. (of Theorem 8.1). Fix s^* as in the statement of the Theorem.

Let $\text{Com}_{1,s}$, $\text{Com}_{2,s}$ denote commitment unitaries for the two quantum bit commitments in bullets 1 and 2 of the theorem statement.

Construction. Define unitary $\widetilde{\text{Com}}_s$ that given advice string s commits to its input (classical) string twice simultaneously, via $\text{Com}_{1,s}$ and $\text{Com}_{2,s}$, as follows:

1. The input register is a $p = p(n)$ qubit register \widetilde{M} , and there is a $p(n) + 2\ell(n)$ qubit auxiliary register \widetilde{W} , divided into sub-registers W_1, M_2, W_2 of $\ell(n), p(n)$ and $\ell(n)$ qubits respectively.
2. First apply a CNOT operation from the \widetilde{M} register onto the M_2 sub-register. (This serves to copy the input string onto M_2 .)
3. Next, apply the $\text{Com}_{1,s}$ unitary to registers \widetilde{M}, W_1 to obtain (C_1, D_1) and the $\text{Com}_{2,s}$ unitary to registers M_2, W_2 to obtain (C_2, D_2) .
4. The output registers are $C = (C_1, C_2), D = (D_1, D_2)$.

Lemma 8.3. $\widetilde{\text{Com}}_s$ satisfies computational binding according to Definition 8.5.

Proof. Suppose, towards a contradiction, that the statement of the claim is false. Then, there exists a (malicious) committer \mathcal{A}_s and a polynomial $q(\cdot)$ such that in $\text{Expmt-Binding}_{\mathcal{A},n,s}^{\widetilde{\text{Com}}}$,

1. \mathcal{A} (arbitrarily) prepares and outputs registers $(\widetilde{C}, \widetilde{D})$ to a challenger.
2. **Validity Check.** The challenger applies $(\widetilde{\text{Com}}_s)^\dagger$ to $(\widetilde{C}, \widetilde{D})$ to obtain $(\widetilde{M}, \widetilde{W})$, then projects \widetilde{W} onto $|0\rangle^{p(n)+2\ell(n)}$, i.e. measures $\{|0\rangle^{p(n)+2\ell(n)}\langle 0|^{p(n)+2\ell(n)}, \mathbb{I} - |0\rangle^{p(n)+2\ell(n)}\langle 0|^{p(n)+2\ell(n)}\}$ on \widetilde{W} . If the measurement rejects, the challenger returns uniform b' , and ends the experiment (this corresponds to the case where decommitment fails to verify).
3. Otherwise, the challenger samples $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, the challenger does nothing.
 - (b) If $b = 1$, the challenger measures \widetilde{M} in the standard basis.
4. Finally, the challenger applies $\widetilde{\text{Com}}_s$ to $(\widetilde{M}, \widetilde{W})$ to obtain $(\widetilde{C}, \widetilde{D})$, and returns \widetilde{D} to \mathcal{A}_s .
5. Denote the output of \mathcal{A}_s by b' , then it holds that for infinitely many n , there exists s such that

$$\Pr[b' = b] \geq \frac{1}{2} + \frac{1}{q(n)}$$

Fix such \mathcal{A}_s and $q(\cdot)$. We will build a reduction \mathcal{B} that uses \mathcal{A}_s to obtain a contradiction as follows.

Reduction \mathcal{B}_s :

1. $\mathcal{B}_s(1^n)$ obtains $s^*(n)$ as non-uniform advice. If $s \geq s^*(n)$, it sets $e = 1$, otherwise $e = 2$. It initializes a challenger for Com_e .
2. It obtains registers (\tilde{C}, \tilde{D}) from \mathcal{A}_s .
3. **Validity Check.** \mathcal{B}_s applies $(\widetilde{\text{Com}}_s)^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then projects \tilde{W} onto $|0\rangle^{p(n)+2\ell(n)}$, i.e. measures $\{|0\rangle^{p(n)+2\ell(n)}\langle 0\rangle^{p(n)+2\ell(n)}|, \mathbb{I} - |0\rangle^{p(n)+2\ell(n)}\langle 0\rangle^{p(n)+2\ell(n)}\}$ on \tilde{W} . If the measurement rejects, \mathcal{B}_s returns uniform $b' \leftarrow \{0, 1\}$, and ends the game.
4. Otherwise, \mathcal{B}_s applies a CNOT operation from the \tilde{M} register onto the M_2 sub-register⁸. It then applies the $\text{Com}_{1,s}$ unitary to registers \tilde{M}, W_1 obtaining C_1, D_1 , and the $\text{Com}_{2,s}$ unitary to registers M_2, W_2 obtaining C_2, D_2 .
5. \mathcal{B}_s finally sends C_e, D_e to its external challenger for $\text{Com}_{e,s}$, then obtains D_e from the external challenger, and returns $\tilde{D} = (D_1, D_2)$ to \mathcal{A}_s .
6. \mathcal{B}_s returns the output bit b' of \mathcal{A}_s as its own output.

Note that $\mathcal{B}_{s,n}$ attempts to break $\text{Com}_{1,s}$ if $s \geq s^*(n)$, and tries to break $\text{Com}_{2,s}$ otherwise. Then the statement of the lemma follows by the following claim (Claim 8.1) together with the fact that \mathcal{B} only incurs a polynomial overhead above the size of \mathcal{A}_s .

Claim 8.1. *At least one of the following is true.*

- *Either are infinitely many $n \in \mathbb{N}$ for which there exists $s \geq s^*(n)$ such that*

$$\text{Adv}_{\text{Com}_{1,s}, \mathcal{B}, n}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}$$

- *Or there are infinitely many $n \in \mathbb{N}$ for which there exists $s \leq s^*(n)$ such that*

$$\text{Adv}_{\text{Com}_{2,s}, \mathcal{B}, n}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}$$

Proof. In what follows, we will show that \mathcal{A} 's state at every step in an interaction with the challenger of $\widetilde{\text{Com}}_s$ when the $\widetilde{\text{Com}}_s$ challenger picks bit b , is identical to its state at every step in an interaction with \mathcal{B} when the $\text{Com}_{e,s}$ challenger picks the same bit b .

Because \mathcal{B} mirrors the output of \mathcal{A} , this will imply that

$$\text{Adv}_{\text{Com}_{e,s}, \mathcal{B}, n}^{\text{Expmt-Binding}} = \text{Adv}_{\widetilde{\text{Com}}_s, \mathcal{A}, n}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}.$$

Fix any $n, s \in \mathbb{N}$. The following experiment describes an interaction of the reduction \mathcal{B}_s above and external challenger for $\text{Com}_{e,s}$, with adversary \mathcal{A}_s against $\widetilde{\text{Com}}_s$.

Expmt₀:

⁸Recall that \tilde{W} is divided into sub-registers W_1, M_2, W_2 of $\ell(n), p(n)$ and $\ell(n)$ qubits respectively

1. \mathcal{B}_s obtains registers (\tilde{C}, \tilde{D}) from \mathcal{A}_s .
2. **Validity Check.** \mathcal{B}_s applies $(\widetilde{\text{Com}}_s)^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then applies measurement $\{|0^{2\ell(n)+p(n)}\rangle \langle 0^{2\ell(n)+p(n)}|, |0^{2\ell(n)+p(n)}\rangle \langle 0^{2\ell(n)+p(n)}|\}$ on \tilde{W} . If the measurement rejects, \mathcal{B}_s returns uniform $b' \leftarrow \{0, 1\}$, and ends the game.
 Otherwise, \mathcal{B}_s applies $\widetilde{\text{Com}}_s$, which entails the following. First, apply a CNOT operation from the \tilde{M} register onto the M_2 sub-register. Next, apply the $\text{Com}_{1,s}$ unitary to registers \tilde{M}, W_1 obtaining C_1, D_1 , and the $\text{Com}_{2,s}$ unitary to registers M_2, W_2 obtaining C_2, D_2 . \mathcal{B}_s then sends C_e, D_e to its external challenger.
3. **Challenger's Validity Check.** The external challenger obtains (C_e, D_e) . It applies $\text{Com}_{e,s}^\dagger$ to (C_e, D_e) to obtain (M_e, W_e) , then applies measurement $\{|0^{\ell(n)}\rangle \langle 0^{\ell(n)}|, \mathbb{I} - |0^{\ell(n)}\rangle \langle 0^{\ell(n)}|\}$ on W_e . If the measurement rejects, it returns uniform $b' \leftarrow \{0, 1\}$, and ends the game.
4. Otherwise, the challenger samples $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, the challenger does nothing.
 - (b) If $b = 1$, the challenger measures M_e in the standard basis.
5. The challenger applies $\text{Com}_{e,s}$ to (M_e, W_e) to obtain (C_e, D_e) and returns D_e to \mathcal{B} .
6. \mathcal{B}_s obtains D_e from the challenger, and returns $\tilde{D} = (D_1, D_2)$ to \mathcal{A}_s .
7. \mathcal{B}_s returns the output bit b' of \mathcal{A}_s .

Observe that the unitary $\text{Com}_{e,s}$ applied at the end of Step 2 in the experiment above is immediately reversed by the challenger at the beginning of Step 3. Since $\text{Com}_{e,s}^\dagger \text{Com}_{e,s} = \mathbb{I}$ for any unitary $\text{Com}_{e,s}$, we can remove both operations from the experiment. At this point, observe that measurement $\{|0^{\ell(n)}\rangle \langle 0^{\ell(n)}|, \mathbb{I} - |0^{\ell(n)}\rangle \langle 0^{\ell(n)}|\}$ on register W_e is performed twice in succession, and since $MM = M$ for any measurement M , we can remove one application of the measurement. Finally, also noting that operations on disjoint subsystems commute, it follows that the above experiment is equivalent to the following simpler experiment.

Expmt₁:

1. \mathcal{B}_s obtains registers (\tilde{C}, \tilde{D}) from \mathcal{A}_s .
2. \mathcal{B}_s 's **Validity Check.** Apply $(\widetilde{\text{Com}}_s)^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then apply measurement $\{|0^{\ell(n)}\rangle \langle 0^{\ell(n)}|, \mathbb{I} - |0^{\ell(n)}\rangle \langle 0^{\ell(n)}|\}$ on \tilde{W} . If the measurement rejects, return uniform $b' \leftarrow \{0, 1\}$, and end the game.
 Otherwise, apply a CNOT operation from the \tilde{M} (also denoted M_1) register onto the M_2 sub-register.
3. Sample $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, do nothing.
 - (b) If $b = 1$, measure M_e in the standard basis.
4. Apply $\text{Com}_{1,s}$ to (M_1, W_1) to obtain (C_1, D_1) , and $\text{Com}_{2,s}$ to (M_2, W_2) obtaining C_2, D_2 .
5. Return $\tilde{D} = (D_1, D_2)$ to \mathcal{A}_s .
6. Set b' to the output of \mathcal{A}_s .

As discussed above, Expmt_1 performs identical operations as Expmt_0 , thus we have

$$\Pr_{\text{Expmt}_1}[b' = b] = \Pr_{\text{Expmt}_0}[b' = b] \quad (8)$$

Finally, let Expmt_2 be identical to Expmt_1 except that M_1 is measured in Step 3 of Expmt_2 , instead of M_e . That is,

Expmt_2 :

1. \mathcal{B}_s obtains registers (\tilde{C}, \tilde{D}) from \mathcal{A}_s .
2. \mathcal{B}_s 's **Validity Check**. Apply $(\widetilde{\text{Com}}_s)^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then apply measurement $\{|0^{\ell(n)}\rangle \langle 0^{\ell(n)}|, \mathbb{I} - |0^{\ell(n)}\rangle \langle 0^{\ell(n)}|\}$ on \tilde{W} . If the measurement rejects, return uniform $b' \leftarrow \{0, 1\}$, and end the game. Otherwise, apply a CNOT operation from the \tilde{M} (also denoted M_1) register onto the M_2 sub-register.
3. Sample $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, do nothing.
 - (b) If $b = 1$, measure M_1 in the standard basis.
4. Apply $\text{Com}_{1,s}$ to (M_1, W_1) to obtain (C_1, D_1) and $\text{Com}_{2,s}$ to M_2, W_2 obtaining C_2, D_2 .
5. Return $\tilde{D} = (D_1, D_2)$ to \mathcal{A}_s .
6. Set b' to the output of \mathcal{A}_s .

Note that the CNOT operation (Step 2) on any state

$$\sum_{i \in [0, 2^{p(n)} - 1]} (\alpha_i |i\rangle)_{M_1} \otimes |0^p\rangle_{M_2}$$

results in the state

$$\sum_{i \in [0, 2^{p(n)} - 1]} \alpha_i |i\rangle_{M_1} |i\rangle_{M_2}.$$

This implies that measuring either one of the registers M_1 or M_2 results in an identical (mixed) state on the system. This implies that Expmt_1 and Expmt_2 are identical, and thus

$$\Pr_{\text{Expmt}_2}[b' = b] = \Pr_{\text{Expmt}_1}[b' = b] \quad (9)$$

Next, note that the only difference between Expmt_2 and $\text{Expmt-Binding}_{\widetilde{\text{Com}}_{s, \mathcal{A}, n}}$ occurs only in the case of $b = 1$. In this case in Expmt_2 , a CNOT is *first* applied from M_1 onto M_2 and then M_1 is measured in the standard basis. Whereas in $\text{Expmt-Binding}_{\widetilde{\text{Com}}_{s, \mathcal{A}, n}}$, M_1 is first measured in the standard basis and then a CNOT is applied from M_1 onto M_2 . Since these operations commute, the experiments are identical. Thus,

$$\Pr_{\text{Expmt}_2}[b' = b] = \text{Adv}_{\widetilde{\text{Com}}_{s, \mathcal{A}, n}}^{\text{Expmt-Binding}} \quad (10)$$

Combining equations (8), (9) and (10) with the fact that $\text{Adv}_{\widetilde{\text{Com}}_{s, \mathcal{A}, n}}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}$ gives us that for \mathcal{B}_s defined above,

$$\text{Adv}_{\text{Com}_{e,s, \mathcal{B}, n}}^{\text{Expmt-Binding}} = \text{Adv}_{\widetilde{\text{Com}}_{s, \mathcal{A}, n}}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}.$$

By definition of \mathcal{B}_s , this implies that at least one of the following is true.

- Either are infinitely many $n \in \mathbb{N}$ for which there exists $s \geq s^*(n)$ such that

$$\text{Adv}_{\text{Com}_{1,s,\mathcal{B},n}}^{\text{Expmt-Binding}} = \text{Adv}_{\text{Com}_{s,\mathcal{A},n}}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}.$$

- Or there are infinitely many $n \in \mathbb{N}$ for which there exists $s \leq s^*(n)$ such that

$$\text{Adv}_{\text{Com}_{2,s,\mathcal{B},n}}^{\text{Expmt-Binding}} = \text{Adv}_{\text{Com}_{s,\mathcal{A},n}}^{\text{Expmt-Binding}} \geq \frac{1}{q(n)}.$$

which completes the proof of the claim. \square

This concludes the proof of Lemma 8.3. \square

The construction given therefore satisfies Definition 8.5, completing the proof of Theorem 8.1. \square

Lemma 8.4. *For non-uniform parameter $s^*(n)$, $\widetilde{\text{Com}}$ satisfies computational hiding according to Definition 8.5.*

Proof. The proof follows by a straightforward hybrid argument. Consider an intermediate unitary $\widehat{\text{Com}}_s$ which, in place of applying a CNOT from \widetilde{M} onto M_2 (and thereby committing to the same message twice), instead initializes M_2 to 0, thereby potentially committing to two different messages via the unitaries $\text{Com}_{1,s}$ and $\text{Com}_{2,s}$. By computational hiding of $\text{Com}_{1,s}$ for $s = s^*(n)$, there is a negligible function $\mu(\cdot)$ such that the state of \mathcal{A}_s in $\text{Expmt-Hiding}_{\mathcal{A},n}^{\widehat{\text{Com}}_s}$ when the challenger picks $b = 1$ is computationally $\mu(n)$ -close to its state when $\widehat{\text{Com}}_s$ is applied instead of $\widetilde{\text{Com}}_s$. Furthermore, by the statistical hiding of Com_2 for $s = s^*(n)$, the state of \mathcal{A}_s in $\text{Expmt-Hiding}_{\mathcal{A},n}^{\widehat{\text{Com}}_s}$ when the challenger picks $b = 0$ is statistically $\mu'(n)$ -close to its state when $\widehat{\text{Com}}_s$ is applied instead of $\widetilde{\text{Com}}_s$. This implies that for every polynomial-sized quantum circuit \mathcal{A}_s and every $n \in \mathbb{N}$,

$$\left| \Pr_{y \leftarrow \widehat{\text{Com}}_{s^*(n)}(1^n,0)} [\mathcal{A}(y) = 1] - \Pr_{y \leftarrow \widetilde{\text{Com}}_{s^*(n)}(1^n,1)} [\mathcal{A}(y) = 1] \right| \leq \mu(n) + \mu'(n) = \text{negl}(n)$$

which completes the proof of the lemma. \square

8.4 Obtaining a (Uniform) Commitment

In the previous section, we obtained a commitment that is binding for every choice of (non-uniform advice) s , but only hiding when the advice string matches the value $s^*(n)$. We show that this implies a uniform construction of commitments, as long as $|s^*(n)|$ is not too large. Formally, we prove the following theorem.

Theorem 8.2. *Assume that there exists a function $s^*(\cdot)$ such that (1) there exists a polynomial $p(\cdot)$ such that for all $n \in \mathbb{N}$, $|s^*(n)| \leq (\log p(n))$, and (2) there exists an s^* -non-uniform hiding commitment according to Definition 8.5. Then there exists a (standard) uniform commitment satisfying computational collapse binding according to Definition 8.1 and computational hiding according to Definition 8.2.*

Proof. Let $t = t(n) = 2^{|s^*(n)|} \leq p(n)$, and let $\text{Com}_1, \text{Com}_2, \dots, \text{Com}_t$ denote the commitment unitaries for m quantum string commitments, one corresponding to each possible value of $s^*(n)$.

Construction. Define unitary $\widetilde{\text{Com}}$ as follows:

1. The input is a single bit b , and there is a $t(n) \cdot (\ell(n) + 1)$ qubit auxiliary register \widetilde{W} , divided into sub-registers $M_1, M_2, M_3, \dots, M_t, W_1, W_2, \dots, W_t$ where the M_i registers are initialized to $|0\rangle$ and the W_i registers are initialized to $|0^{\ell(n)}\rangle$.

2. $\widetilde{\text{Com}}$ applies the H gate to each of the registers M_1, \dots, M_{t-1} . Next, it applies unitary U that maps

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (m, x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [t-1]} x_i \oplus y \oplus m)$$

to registers $\widetilde{M}, M_1, \dots, M_t$.

3. Next for $i \in [t]$, it applies the Com_i unitary to registers M_i, W_i to obtain (C_i, D_i) .

4. The output registers are $C = (C_1, \dots, C_t), D = (D_1, \dots, D_t)$.

Lemma 8.5. $\widetilde{\text{Com}}$ satisfies computational hiding according to Definition 8.2.

Proof. Let $s^* = s^*(n) \in t[n]$ be such that Com_{s^*} satisfies statistical (resp., computational) hiding according to Definition 8.2.

Further, assume towards a contradiction that the lemma is not true. Then there exists an unbounded (resp., quantum polynomial-sized) adversary \mathcal{A} and a polynomial $q(\cdot)$ such that $\text{Adv}_{\widetilde{\text{Com}}, \mathcal{A}, n}^{\text{Expmt-Hiding}}$ (in Definition 8.2) is at least $\frac{1}{q(n)}$. Consider the following intermediate experiment.

Expmt-intermediate $_{\mathcal{A}, n}$:

- \mathcal{A} outputs a quantum state on register M .
- The challenger samples $b \leftarrow \{0, 1\}$, then swaps the contents of M with $|0\rangle$ if $b = 0$. It initializes \widetilde{W} , divided into sub-registers $M_1, M_2, M_3, \dots, M_t, W_1, W_2, \dots, W_t$ where for all $j \in [t]$, the M_j registers are initialized to $|0\rangle$ and the W_j registers are initialized to $|0^{\ell(n)}\rangle$.

It then applies the H gate to each of the registers M_1, \dots, M_{t-1} . Next, it applies unitary U mapping

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (m, x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [n-1]} x_i \oplus y \oplus m)$$

to registers M, M_1, \dots, M_t . Next, it swaps out the value on register $M_{s^*(n)}$ to $|0\rangle$, and finally for $j \in [t]$, it applies the Com_j unitary to registers M_j, W_j to obtain (C_j, D_j) .

It sends $\widetilde{C} = (C_1, \dots, C_t)$ to \mathcal{A}_s .

- Denote the output of \mathcal{A} by b' .

Note that the contents of registers C_1, \dots, C_t in this intermediate experiment are independent of the bit b (due to swapping out M_j with $|0\rangle$), therefore for any \mathcal{A} , $\Pr[b' = b] = \frac{1}{2}$ in the intermediate experiment.

This implies that there exists a fixing of $b \in \{0, 1\}$ such that \mathcal{A} 's output in the intermediate experiment $\text{Expmt-Hiding}_{\mathcal{A}, n}^{\widetilde{\text{Com}}}$ with this fixing of b , is at least $\frac{1}{2q(n)}$ -far from its output in $\text{Expmt-intermediate}_{\mathcal{A}, n}$. Suppose this holds for $b = 1$ (the case of $b = 0$ follows similarly). We build a (non-uniform) reduction \mathcal{B} that interacts with an external challenger to break the purported hiding of Com_{s^*} , as follows.

1. Obtain a quantum state on register M from \mathcal{A} .
2. Prepare a $t(n) \cdot (\ell(n) + 1)$ qubit register \widetilde{W} , divided into sub-registers $M_1, M_2, M_3, \dots, M_t$, and W_1, W_2, \dots, W_t where for all $j \in [t]$, the M_j registers are initialized to $|0\rangle$ and the W_j registers are initialized to $|0^{\ell(n)}\rangle$.

3. Apply the H gate to each of the registers M_1, \dots, M_{t-1} . Next, apply unitary U that maps

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [t-1]} x_i \oplus y \oplus m)$$

to registers M, M_1, \dots, M_t .

4. Send the M_{s^*} register to the external challenger, and obtain register C_{s^*} .
5. For $j \in [t] \setminus s^*$, apply the Com_j unitary to registers M_j, W_j to obtain (C_j, D_j) .
6. Send C_1, \dots, C_t to \mathcal{A} and return the output b' of \mathcal{A} .

By construction, the size of \mathcal{B} is only polynomially larger than that of \mathcal{A} , and

$$\left| \Pr[\mathcal{B} \text{ outputs } 1 | b = 0] - \Pr[\mathcal{B} \text{ outputs } 1 | b = 1] \right| \geq \frac{1}{2q(n)}$$

This implies that

$$\left| \Pr[\mathcal{B}(\text{Com}_{s^*}(1^n, 0)) = 1] - \Pr[\mathcal{B}(\text{Com}_{s^*}(1^n, 1)) = 1] \right| \geq \frac{1}{2q(n)}$$

which is a contradiction to the hiding of Com_{s^*} , as desired, and thus the lemma must be true. \square

Lemma 8.6. $\widetilde{\text{Com}}$ satisfies statistical (resp., computational) collapse binding according to Definition 8.1 as long as all of $\text{Com}_1, \dots, \text{Com}_t$ satisfy statistical (resp., computational) collapse binding according to Definition 8.1.

Proof. Suppose, towards a contradiction, that the statement of the theorem is false. Then, there exists a (malicious) committer \mathcal{A} and a polynomial $q(\cdot)$ such that in $\text{Expmt}\text{-Binding}_{\mathcal{A}, n}^{\widetilde{\text{Com}}}$,

1. \mathcal{A} (arbitrarily) prepares and outputs registers $(\widetilde{C}, \widetilde{D})$ to a challenger.
2. **Validity Check.** The challenger applies $(\widetilde{\text{Com}})^\dagger$ to $(\widetilde{C}, \widetilde{D})$ to obtain $(\widetilde{M}, \widetilde{W})$, then projects \widetilde{W} onto $|0\rangle^{t(n) \cdot (1+\ell(n))}$. If the projection rejects, the challenger returns uniform b' , and ends the experiment (this corresponds to the case where decommitment fails to verify).
3. Otherwise, the challenger samples $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, the challenger does nothing.
 - (b) If $b = 1$, the challenger measures \widetilde{M} in the standard basis.
4. Finally, the challenger applies $\widetilde{\text{Com}}$ to $(\widetilde{M}, \widetilde{W})$ to obtain $(\widetilde{C}, \widetilde{D})$.

This entails applying the H gate to each of the registers M_1, \dots, M_{t-1} . Next, applying unitary U that maps

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (m, x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [t-1]} x_i \oplus y \oplus m)$$

to registers $\widetilde{M}, M_1, \dots, M_t$. Finally for $i \in [t]$, applying the Com_i unitary to registers M_i, W_i to obtain (C_i, D_i) .

5. The challenger returns \widetilde{D} to \mathcal{A}_s .

6. Denote the output of \mathcal{A}_s by b' , then it holds that (for infinitely many n),

$$\Pr[b' = b] \geq \frac{1}{2} + \frac{1}{q(n)}$$

We will now consider a sequence of $t(n)$ hybrids, where Hyb_0 is identical to the one outlined above, and for each $j \in [t(n)]$, Hyb_j is identical to Hyb_{j-1} , except that the M_j register is measured in the standard basis before applying the Com_i unitary at the end of Step 4. We write the description of Hyb_j for completeness below, with the difference from the binding experiment underlined.

Hyb_j :

1. \mathcal{A} (arbitrarily) prepares and outputs registers (\tilde{C}, \tilde{D}) to a challenger.
2. **Validity Check.** The challenger applies $(\widetilde{\text{Com}})^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then projects \tilde{W} onto $|0\rangle^{t(n) \cdot (1+\ell(n))}$. If the projection rejects, the challenger returns uniform b' , and ends the experiment (this corresponds to the case where decommitment fails to verify).
3. Otherwise, the challenger samples $b \leftarrow \{0, 1\}$ and:
 - (a) If $b = 0$, the challenger does nothing.
 - (b) If $b = 1$, the challenger measures \tilde{M} in the standard basis.

4. The challenger applies the H gate to each of the registers M_1, \dots, M_{t-1} . Next, it applies unitary U that maps

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (m, x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [t-1]} x_i \oplus y \oplus m)$$

to registers $\tilde{M}, M_1, \dots, M_t$. Next, it measures registers M_1, \dots, M_j in the standard basis.

Finally for $i \in [t]$, it applies the Com_i unitary to registers M_i, W_i to obtain (C_i, D_i) .

5. The challenger returns \tilde{D} to \mathcal{A} .
6. Denote the output of \mathcal{A} by b' .

Note that in Hyb_t , $\Pr[b' = b] = \frac{1}{2}$ for any \mathcal{A} , because measuring all registers M_1, \dots, M_t in the standard basis implies a measurement of M in the standard basis. Moreover by our assumption above, in Hyb_0 , $\Pr[b' = b] \geq \frac{1}{2} + \frac{1}{q(n)}$ for adversary \mathcal{A} and some polynomial $q(\cdot)$. This implies that for every n , there exists $j^* = j^*(n) \in [t(n)]$ such that

$$\left| \Pr[b' = b | \text{Hyb}_{j^*}] - \Pr[b' = b | \text{Hyb}_{j^*-1}] \right| \geq \frac{1}{q(n)t(n)} \quad (11)$$

To complete the contradiction, we build a reduction \mathcal{B} that breaks statistical (resp., computational) binding of Com_{j^*} , by doing the following.

1. Obtain registers (\tilde{C}, \tilde{D}) from \mathcal{A}_s .
2. **Validity Check.** \mathcal{B} applies $(\widetilde{\text{Com}})^\dagger$ to (\tilde{C}, \tilde{D}) to obtain (\tilde{M}, \tilde{W}) , then projects \tilde{W} onto $|0\rangle^{t(n) \cdot (1+\ell(n))}$. If the projection rejects, it returns uniform b' , and ends the experiment (this corresponds to the case where decommitment fails to verify).

3. Otherwise, \mathcal{B} samples $b \leftarrow \{0, 1\}$ and:

- (a) If $b = 0$, it does nothing.
- (b) If $b = 1$, it measures \tilde{M} in the standard basis.

4. \mathcal{B} applies the H gate to each of the registers M_1, \dots, M_{t-1} . Next, it applies unitary U that maps

$$(m, x_1, x_2, \dots, x_{t-1}, y) \mapsto (m, x_1, x_2, \dots, x_{t-1}, \bigoplus_{i \in [t-1]} x_i \oplus y \oplus m)$$

to registers $\tilde{M}, M_1, \dots, M_t$. Next, it measures registers M_1, \dots, M_{j^*-1} in the standard basis.

Finally for $i \in [t]$, it applies the Com_i unitary to registers M_i, W_i to obtain (C_i, D_i) .

5. \mathcal{B}_s then sends C_{j^*}, D_{j^*} to the binding challenger for Com_{j^*} , and then obtains D_{j^*} from the challenger.

6. \mathcal{B}_s returns $\tilde{D} = (D_1, \dots, D_t)$ to \mathcal{A}_s .

7. Let b' denote the output of \mathcal{A} . Output 1 if $b' = b$, otherwise output 0.

Note that when the external challenger for Com_{j^*} samples its challenge $c = 0$, the interaction of \mathcal{B} with \mathcal{A} corresponds to Hyb_{j^*-1} , and otherwise to Hyb_{j^*} . Thus by equation (11), we have that:

$$\left| \Pr[\mathcal{B} \text{ outputs } 1 | c = 1] - \Pr[\mathcal{B} \text{ outputs } 1 | c = 0] \right| \geq \frac{1}{q(n)t(n)}.$$

Since the size of \mathcal{B} is only polynomially larger than that of \mathcal{A} , this contradicts⁹ the statistical (resp., computational) binding of Com_{j^*} , as desired. \square

This concludes the proof of Theorem 8.2. \square

Theorem 8.3. *There exists a constant $c > 0$ such that cn copy secure one-way state generators with pure state outputs (Definition 3.1) imply quantum bit commitments satisfying computational collapse binding according to Definition 8.1 and computational hiding according to Definition 8.2.*

Proof. The theorem follows from Corollary 7.1, Theorem 8.1 and Theorem 8.2. \square

Computationally hiding and computationally collapse binding quantum bit commitments are known [BCQ23] to imply EFI pairs, which in turn are known to imply secure computation for all classical and quantum functionalities [BCKM21, GLSV21, AQY22, BCQ23]. We therefore also have the following corollary.

Corollary 8.2. *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.*

9 Acknowledgments

We thank James Bartusek, Yanyi Liu, Amit Sahai, and Taiga Hiroka for illuminating discussions and useful comments. Both 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 number FA9550-23-1-0543.

⁹Note that any \mathcal{B} satisfying the equation above can be converted (with polynomial overhead) to an adversary that has advantage at least $\frac{1}{2q(n)t(n)}$ in the binding game.

References

- [AA14] Scott Aaronson and Alex Arkhipov. Bosonsampling is far from uniform. *Quantum Inf. Comput.*, 14(15-16):1383–1423, 2014.
- [Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016.
- [Aar20] Scott Aaronson. Shadow tomography of quantum states. *SIAM Journal on Computing*, 49(5):STOC18–368–STOC18–394, 2020.
- [AC02] Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In Helmut Alt and Afonso Ferreira, editors, *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science, Antibes - Juan les Pins, France, March 14-16, 2002, Proceedings*, volume 2285 of *Lecture Notes in Computer Science*, pages 323–334. Springer, 2002.
- [ACC⁺22] Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Advances in Cryptology – CRYPTO 2022, Part II*, Lecture Notes in Computer Science, pages 165–194, Santa Barbara, CA, USA, August 2022. Springer, Heidelberg, Germany.
- [AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 237–265. Springer, 2022.
- [AK07] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory Comput.*, 3(1):129–157, 2007.
- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Advances in Cryptology – CRYPTO 2022, Part I*, Lecture Notes in Computer Science, pages 208–236, Santa Barbara, CA, USA, August 2022. Springer, Heidelberg, Germany.
- [BB84] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- [BCKM21] 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, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 467–496, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.
- [BCM⁺18] Zvika Brakerski, Paul F. Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 320–331. IEEE Computer Society, 2018.

- [BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10–13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPICs*, pages 24:1–24:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [Bel02] Mihir Bellare. A note on negligible functions. *J. Cryptol.*, 15(4):271–284, 2002.
- [CX22] Shujiao Cao and Rui Xue. On constructing one-way quantum state generators, and more. *IACR Cryptol. ePrint Arch.*, page 1323, 2022.
- [DS23] Marcel Dall’Agnol and Nicholas Spooner. On the necessity of collapsing for post-quantum and quantum commitments. In Omar Fawzi and Michael Walter, editors, *18th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2023, July 24–28, 2023, Aveiro, Portugal*, volume 266 of *LIPICs*, pages 2:1–2:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [GJMZ23] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20–23, 2023*, pages 1579–1588. ACM, 2023.
- [GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In David S. Johnson, editor, *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14–17, 1989, Seattle, Washington, USA*, pages 25–32. ACM, 1989.
- [GLSV21] 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, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 531–561, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, Oct 2020.
- [HMY23] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. From the hardness of detecting superpositions to cryptography: Quantum public key encryption and commitments. In Carmit Hazay and Martijn Stam, editors, *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*, volume 14004 of *Lecture Notes in Computer Science*, pages 639–667. Springer, 2023.
- [HRV13] Iftach Haitner, Omer Reingold, and Salil P. Vadhan. Efficiency improvements in constructing pseudorandom generators from one-way functions. *SIAM J. Comput.*, 42(3):1405–1430, 2013.
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, Research Triangle Park, NC, USA, October 30 – November 1, 1989. IEEE Computer Society Press.

- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st Annual ACM Symposium on Theory of Computing*, pages 12–24, Seattle, WA, USA, May 15–17, 1989. ACM Press.
- [IR90] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages 8–26, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Heidelberg, Germany.
- [JLS18] 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.
- [KNY23] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Publicly verifiable deletion from minimal assumptions. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 228–245, Cham, 2023. Springer Nature Switzerland.
- [KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1589–1602. ACM, 2023.
- [Kre21] 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.
- [LC97] Hoi-Kwong Lo and Hoi Fung Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78(17):3410, 1997.
- [LMW23] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. *Cryptology ePrint Archive*, Paper 2023/1602, 2023. <https://eprint.iacr.org/2023/1602>.
- [LR86] Michael Luby and Charles Rackoff. Pseudo-random permutation generators and cryptographic composition. In *18th Annual ACM Symposium on Theory of Computing*, pages 356–363, Berkeley, CA, USA, May 28–30, 1986. ACM Press.
- [May97] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical review letters*, 78(17):3414, 1997.
- [MY22a] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *Cryptology ePrint Archive*, Paper 2022/1336, 2022. <https://eprint.iacr.org/2022/1336>.
- [MY22b] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Advances in Cryptology – CRYPTO 2022, Part I*, *Lecture Notes in Computer Science*, pages 269–295, Santa Barbara, CA, USA, August 2022. Springer, Heidelberg, Germany.

- [MY23a] Tomoyuki Morimae and Takashi Yamakawa. Proofs of quantumness from trapdoor permutations. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPICs*, pages 87:1–87:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [MY23b] Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. *CoRR*, abs/2302.04749, 2023.
- [Nao89] Moni Naor. Bit commitment using pseudo-randomness. In Gilles Brassard, editor, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 128–136. Springer, 1989.
- [Ren08] Renato Renner. Security of quantum key distribution. *International Journal of Quantum Information*, 6(01):1–127, 2008.
- [Unr14] Dominique Unruh. Quantum position verification in the random oracle model. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15:78–88, 1983.
- [Yan21] Jun Yan. Quantum computationally predicate-binding commitments with application in quantum zero-knowledge arguments for NP. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part I*, volume 13090 of *Lecture Notes in Computer Science*, pages 575–605, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany.
- [Yan22] Jun Yan. General properties of quantum bit commitments (extended abstract). In *Advances in Cryptology – ASIACRYPT 2022, Part IV*, *Lecture Notes in Computer Science*, pages 628–657. Springer, Heidelberg, Germany, December 7–11, 2022.
- [YK] Henry Yuen and Asa Kosto. Scribe notes on Shadow Tomography. <https://www.henryyuen.net/spring2022/lec5-shadow.pdf>.

A One-Way Puzzles are Necessary for QCCC Cryptography

In this section, we show that a variety of protocols, including encryption, commitments and digital signatures in the hybrid or local or quantum computation classical communication (QCCC) model imply one-way puzzles.

A.1 Public Key Encryption

First, we define public-key encryption in the QCCC model. Here, public keys and ciphertexts are classical while secret keys may be quantum. We provide a formal definition and then prove that this type of PKE implies one-way puzzles. Later we will also show that PKE with classical public and secret keys, but quantum ciphertexts, implies one-way puzzles.

Definition A.1 (Public Key Encryption). *A public key encryption scheme consists of a set of QPT algorithms (KeyGen, Enc, Dec) where*

- $\text{KeyGen}(1^n)$ on input the security parameter outputs a classical public key pk and a quantum secret key ρ_{sk} .
- $\text{Enc}(pk, m)$ takes a public key and a classical message as input and outputs a classical ciphertext.
- $\text{Dec}(\rho, c)$ takes a secret key ρ and a classical ciphertext c as input and outputs a message.

It has the following properties.

- **Correctness.** For all $m \in \{0, 1\}^n$,

$$\Pr_{(pk, \rho_{sk}) \leftarrow \text{KeyGen}(1^n)} [\text{Dec}(\rho_{sk}, \text{Enc}(pk, m)) = m] \geq 1 - \text{negl}(n)$$

- **Hiding.** For all QPT adversaries \mathcal{A} ¹⁰,

$$\Pr_{\substack{(pk, \rho_{sk}) \leftarrow \text{KeyGen}(1^n) \\ m \leftarrow \{0, 1\}^n}} [\mathcal{A}(pk, (\text{Enc}(pk, m))) = m] \leq \text{negl}(n)$$

Theorem A.1. *Public key encryption in the QCCC model implies one-way puzzles.*

Proof. Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a PKE scheme. We define algorithms $(\text{Samp}, \text{Ver})$ as follows:

- $\text{Samp}(1^n)$: Sample $(pk, \cdot) \leftarrow \text{KeyGen}(1^n)$, $m \leftarrow \{0, 1\}^n$ and return $(k := m, s := (pk, \text{Enc}(pk, m)))$.
- $\text{Ver}(k, s)$: Parse $k := m$ and $s := (pk, ct)$. Let ρ_{sk} denote the mixed secret key corresponding to public key s . Accept if and only if $m = \text{Dec}(\rho_{sk}, ct)$.

By the correctness of the PKE scheme, with overwhelming probability over the sampling of keys, for any message m , $\text{Dec}(\rho_{sk}, \text{Enc}(pk, m)) = m$. Correctness of the one-way puzzle follows. Security of the one-way puzzle also follows immediately by hiding of the PKE scheme. \square

We also show that public-key encryption with classical public and secret keys implies one-way puzzles, even when the ciphertexts are quantum.

Theorem A.2. *Public key encryption with classical keys implies one-way puzzles.*

Proof. Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a PKE scheme. We define algorithms $(\text{Samp}, \text{Ver})$ as follows:

- $\text{Samp}(1^n)$: Sample $(pk, sk) \leftarrow \text{KeyGen}(1^n)$ and return (sk, pk) .
- $\text{Ver}(k, s)$: Sample $m \leftarrow \{0, 1\}^n$ and accept if and only if $m = \text{Dec}_k(\text{Enc}_s(m))$

Claim A.1. $(\text{Samp}, \text{Ver})$ satisfies Definition 3.2.

Proof. By the correctness of the PKE scheme, with overwhelming probability over the sampling of keys (pk, sk) , for any message m , $\text{Dec}_{sk}(\text{Enc}_{pk}(m)) = m$. Correctness of the one-way puzzle follows.

Suppose there exists an adversary that breaks the security of the one-way puzzle. That is there exists a non-negligible function $p(\cdot)$ such that for all $n \in \mathbb{N}$

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

¹⁰We note that this is a weaker property than CPA-security, we use this because it will give us a stronger result.

Rewriting using the definitions of Samp and Ver we get

$$\Pr_{\substack{(pk,sk) \leftarrow \text{KeyGen}(1^n) \\ m \leftarrow \{0,1\}^n}} [sk^* \leftarrow \mathcal{A}(pk) \wedge m = \text{Dec}_{sk^*}(\text{Enc}_{pk}(m))]$$

Let \mathcal{B} be an algorithm that takes (pk, z) as input, runs $\mathcal{A}(pk)$ to get sk^* and outputs $\text{Dec}_{sk^*}(z)$. Then

$$\Pr_{\substack{(pk,sk) \leftarrow \text{KeyGen}(1^n) \\ m \leftarrow \{0,1\}^n}} [\mathcal{B}(pk, (\text{Enc}_{pk}(m)) = m) \geq p(n)]$$

which contradicts the hiding of the PKE scheme. □

The theorem follows trivially from the claim. □

A.2 Digital Signatures

We define signatures in the QCCC model. Here, verification keys and signatures are classical, while signing keys may be quantum. We provide a formal definition and then prove that such signatures imply one-way puzzles. Later, we will also show that signature schemes with classical (signing and verification) keys, but quantum signatures, also imply one-way puzzles.

Definition A.2 (Signature Scheme). *A digital signature scheme consists of a set of QPT algorithms (KeyGen, S, V) where*

- $\text{KeyGen}(1^n)$ takes the security parameter as input and outputs a signing key ρ_{sk} and a verification key vk .
- $S(\rho_{sk}, m)$ takes a signing key and a message as input and outputs a signature σ .
- $V(vk, \sigma, m)$ takes a verification key, a signature, and a message as input and outputs a bit b .

A signature scheme has the following properties.

- **Correctness.** For all $m \in \{0, 1\}^n$,

$$\Pr_{(\rho_{sk}, vk) \leftarrow \text{KeyGen}(1^n)} [V(vk, S(\rho_{sk}, m), m) = 1] \geq 1 - \text{negl}(n)$$

- **Unforgeability.** For all QPT adversaries \mathcal{A} ,

$$\Pr_{(\rho_{sk}, vk) \leftarrow \text{KeyGen}(1^n)} [(t, m) \leftarrow \mathcal{A}(vk) \wedge V(vk, t, m) = 1] \leq \text{negl}(n)$$

Theorem A.3. *QCCC signature schemes imply one-way puzzles.*

Proof. Let (KeyGen, S, V) be a PKE scheme. We define the algorithms $(\text{Samp}, \text{Ver})$ as follows:

- $\text{Samp}(1^n)$: Sample $(\rho_{sk}, vk) \leftarrow \text{KeyGen}(1^n)$. Use ρ_{sk} to sign $m = 0^n$, obtaining $t = S(\rho_{sk}, 0^n)$. Return $(k := t, s := vk)$.
- $\text{Ver}(k, s)$: Accept if and only if $V(s, k, 0^n) = 1$.

By the correctness of the signature scheme, with overwhelming probability over the sampling of keys (sk, vk) , for any message m (and in particular for $m = 0^n$), $V(vk, (S(sk, m), m)) = 1$. Correctness of the one-way puzzle follows. Furthermore, any adversary that breaks security of the one-way puzzle outputs a valid signature on 0^n with respect to a randomly sampled verification key vk , breaking unforgeability as defined above. \square

Theorem A.4. *Signature schemes with classical verification and signing keys imply one-way puzzles.*

Proof. Let (KeyGen, S, V) be a PKE scheme. We define the algorithms Samp, Ver as follows:

- $\text{Samp}(1^n)$: Sample $(sk, vk) \leftarrow \text{KeyGen}(1^n)$ and return (vk, sk) .
- $\text{Ver}(k, s)$: Sample $m \leftarrow \{0, 1\}^n$ and accept if and only if $V(vk, (S(sk, m), m)) = 1$

Claim A.2. $(\text{Samp}, \text{Ver})$ satisfies Definition 3.2.

Proof. By the correctness of the signature scheme, with overwhelming probability over the sampling of keys (sk, vk) , for any message m , $V(vk, (S(sk, m), m)) = 1$. Correctness of the one-way puzzle follows.

Suppose there exists an adversary that breaks the security of the one-way puzzle. That is there exists a non-negligible function $p(\cdot)$ such that for all $n \in \mathbb{N}$

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

Rewriting using the definitions of Samp and Ver we get

$$\Pr_{\substack{(sk,vk) \leftarrow \text{KeyGen}(1^n) \\ m \leftarrow \{0,1\}^n}} [sk^* \leftarrow \mathcal{A}(vk) \wedge V(vk, (S(sk^*, m), m)) = 1]$$

Let \mathcal{B} be an algorithm that takes vk as input, runs $\mathcal{A}(vk)$ to get sk^* , samples $m \leftarrow \{0, 1\}^n$ and outputs $S(sk^*, m), m$. Then

$$\Pr_{(sk,vk) \leftarrow \text{KeyGen}(1^n)} [(\sigma, m) \leftarrow \mathcal{B}(vk) \wedge V(vk, (\sigma, m)) = 1] \geq p(n)$$

which contradicts the unforgeability of the signature scheme. \square

The theorem follows trivially from the claim. \square

A.3 Bit Commitments

We define commitments in the QCCC model below. The definition here differs from the one in Section 8.1 since unlike the commitment scheme we build from pure OWSGs, general QCCC commitments may be interactive.

Definition A.3 (Bit Commitment Scheme, Syntax). *A (QCCC) bit commitment scheme is an efficient two-party protocol $\text{Com} = \langle C, R \rangle$ between a committer C and a receiver R consisting of a commit stage and an opening stage.*

- **Commit Stage.** Both parties receive the security parameter 1^n and the committer C receives a private input $b \in \{0, 1\}$. It interacts with the receiver R to produce a classical transcript z . At the end of the stage each party outputs a (private) quantum state, denoted by q_C and q_R respectively.

- **Opening Stage.** Both parties receive the transcript z produced in the first stage and their respective output states. They then interact with classical communication and at the end of this stage the receiver outputs a bit or the reject symbol \perp .

A bit-commitment scheme satisfies correctness if there exists a negligible function $\mu(\cdot)$ such that for all $n \in \mathbb{N}$ and all $b \in \{0, 1\}$, when C and R are honest, R outputs b at the end of the opening stage with probability at least $1 - \mu(n)$.

Definition A.4 (Computational Hiding). A bit-commitment scheme is computationally hiding if there exists a negligible function $\mu(\cdot)$ such that for an honest committer C that receives a bit b as input, no QPT adversarial receiver \mathcal{R} can distinguish interactions where C receives $b = 0$ and interactions where C receives $b = 1$ with advantage greater than $\mu(n)$.

Below, we define a weak notion of binding for commitments, that we label computational weak honest binding. This is implied by other standard notions such as honest binding [Yan22].

Definition A.5 (Computational Weak Honest Binding). A quantum bit commitment scheme is weakly honest binding if there exists a negligible function $\mu(\cdot)$ such that for $b \in \{0, 1\}$, no QPT adversary \mathcal{A} wins the following game with probability greater than $\mu(n)$.

- Run the commit stage of the commitment with an honest committer C that receives input b and an honest receiver R. Let z be the transcript and let q_R be the receiver state at the end of the commit stage.
- Run the opening stage between \mathcal{A} and R, where \mathcal{A} receives z (but not the state output by C) and R receives q_R . The adversary wins if R outputs $1 - b$.

We say that a bit-commitment scheme has non-interactive opening if the opening stage consists of a single message sent from the committer to the receiver.

Theorem A.5. QCCC Bit commitments that are computationally hiding, computationally weak honest binding, and have non-interactive openings imply one-way puzzles.

Proof. Let $\text{Com} = (C, R)$ be a computationally hiding and computationally weak honest binding bit-commitment scheme. We define the algorithms Samp , Ver as follows:

- $\text{Samp}(1^n)$:
 1. Run the commit stage with an honest committer C that receives input 0 and an honest receiver R. Let z be the transcript and let q_C and q_R be the committer and receiver outputs respectively.
 2. Run the opening stage with an honest committer C that receives input (z, q_C) and an honest receiver R that receives input (z, q_R) . Let the opening message be d .
 3. Output (d, z)
- $\text{Ver}(k, s)$:
 1. Consider an honest execution of the commit stage between C with input 0 and R. Sample a random output state q_R^* for R conditioned on the transcript being s . This may be achieved in exponential time by rejection sampling. Output \perp if no such q_R^* exists.
 2. Return \top if the output of R on input q_R^* and message k is 0. Return \perp otherwise.

Claim A.3. $(\text{Samp}, \text{Ver})$ satisfies Definition 3.2.

Proof. By construction, for a transcript z generated by an honest execution of commitment to 0, the state q_R^* generated by $\text{Ver}(d, z)$ for arbitrary d is identically distributed to the state q_R output by the receiver during the honest execution. $\text{Ver}(d, z)$ accepts whenever an honest receiver accepts message d as an opening to 0 given transcript z and state q_R^* . Therefore, for any message d , the probability that $\text{Ver}(d, z)$ accepts is the probability that an honest receiver accepts message d as an opening to 0 given transcript z and state q_R .

Correctness. Samp generates the message d as an honest opening for z , an honest commitment to 0. Therefore by the correctness of the commitment scheme the probability that $\text{Ver}(d, z)$ accepts is at least $1 - \text{negl}(n)$. Therefore

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

which proves correctness of the oneway puzzle.

Security. We prove by contradiction. Suppose there exists a QPT adversary \mathcal{A} that breaks one-way puzzle security. That is there exists a non-negligible function $\epsilon(\cdot)$ such that for all $n \in \mathbb{N}$:

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

Since s is sampled as the transcript of an honest commitment to zero, this means that an honest receiver accepts the output of $\mathcal{A}(s)$ as an opening to zero given transcript s and the receiver output from the commitment execution. This may be written as

$$\Pr [0 \leftarrow R(z, q_R, \mathcal{A}(z)) \mid z, q_C, q_R \leftarrow \langle C(0, 1^n), R(1^n) \rangle] \geq \epsilon(n)$$

By the computational weak honest binding property of the commitment, the probability that $\mathcal{A}(z)$ is accepted as an opening to zero when the value committed to is one must be negligible. That is,

$$\Pr [0 \leftarrow R(z, q_R, \mathcal{A}(z)) \mid z, q_C, q_R \leftarrow \langle C(1, 1^n), R(1^n) \rangle] = \text{negl}(n)$$

We now build an algorithm \mathcal{B} that breaks the hiding of the commitment scheme. \mathcal{B} interacts with an honest committer C as follows:

- During the commit stage, simulate an honest receiver in interactions with C to obtain a transcript z and a receiver state q_R .
- Run \mathcal{A} on input z to obtain opening transcript d^* .
- Return 1 if $R(z, q_R, d^*)$ outputs 0. Else output 0.

By the two previous inequalities, the probability that \mathcal{B} outputs 1 when the commitment is to 0 is at least $\epsilon(n)$. Likewise, the probability that \mathcal{B} outputs 1 when the commitment is to 1 is at most $\text{negl}(n)$. \mathcal{B} therefore distinguishes both interactions with advantage at least $\epsilon(n) - \text{negl}(n)$, which contradicts the hiding of the commitment scheme. □

The theorem follows trivially from the previous claim. □

A.4 Symmetric Encryption with Classical Keys and Ciphertexts

We prove that symmetric encryption schemes for classical messages, with classical keys and ciphertexts imply one-way puzzles.

Definition A.6 (Symmetric Encryption). *A symmetric encryption scheme consists of a set of QPT algorithms (KeyGen, Enc, Dec) where*

- $\text{KeyGen}(1^n)$ takes the security parameter as input and outputs a secret key k .
- $\text{Enc}_{sk}(m)$ takes a key and a message as input and outputs a ciphertext c .
- $\text{Dec}_k(c)$ takes a key and a ciphertext as input and outputs a message.

A symmetric encryption scheme has the following properties.

- **Correctness.** For all $m \in \{0, 1\}^n$,

$$\Pr_{k \leftarrow \text{KeyGen}(1^n)} [\text{Dec}_k(\text{Enc}_k(m)) = m] \geq 1 - \text{negl}(n)$$

- **Hiding.** All QPT adversaries \mathcal{A} win the following game with probability less than $1/2 + \text{negl}(n)$.

- $m_0, m_1 \leftarrow \mathcal{A}(1^n)$
- $b \leftarrow \{0, 1\}$
- $k \leftarrow \text{KeyGen}(1^n)$
- $b' \leftarrow \mathcal{A}(\text{Enc}_k(m_b))$

\mathcal{A} wins the game if $b' = b$.

We additionally assume that the size of the keyspace is a negligible fraction of the message space, i.e. $\frac{|\text{Supp}(\text{KeyGen}(1^n))|}{2^n} = \text{negl}(n)$. This is without loss of generality for symmetric encryption schemes that support multi-message encryption.

Theorem A.6. Symmetric encryption in the QCCC model implies one-way puzzles.

Proof. Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric encryption scheme. We define the algorithms Samp, Ver as follows:

- $\text{Samp}(1^n)$: Sample $k \leftarrow \text{KeyGen}(1^n)$ and $m \leftarrow \{0, 1\}^n$ and return $k, (m, \text{Enc}_k(m))$.
- $\text{Ver}(k, s)$: Parse s as (m, c) and accept if $m = \text{Dec}_k(c)$. Reject otherwise.

Claim A.4. $(\text{Samp}, \text{Ver})$ satisfies Definition 3.2.

Proof. By the correctness of the symmetric encryption scheme, with overwhelming probability over the sampling of key k , for any message m , $\text{Dec}_k(\text{Enc}_k(m)) = m$. Correctness of the one-way puzzle follows.

Suppose there exists an adversary that breaks the security of the one-way puzzle. That is there exists a non-negligible function $\epsilon(\cdot)$ such that for all $n \in \mathbb{N}$

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

Rewriting using the definitions of Samp and Ver we get

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^n) \\ m \leftarrow \{0,1\}^n}} [k^* \leftarrow \mathcal{A}(m, \text{Enc}_k(m)) \wedge m = \text{Dec}_{k^*}(\text{Enc}_k(m))] \geq p(n)$$

We now build an algorithm \mathcal{B} that breaks the hiding of the encryption scheme. \mathcal{B} interacts with an external CPA challenger as follows.

1. Sample $m_0, m_1 \leftarrow \{0, 1\}^n$.

2. Send m_0, m_1 to the external challenger.
3. Receive $c = \text{Enc}_k(m_b)$ for some freshly sampled k and uniform bit b .
4. Compute $k^* \leftarrow \mathcal{A}(m_0, c)$
5. If $m_0 = \text{Dec}_{k^*}(c)$ then return 0, else return a uniform bit.

We calculate the probability that \mathcal{B} guess the uniform bit b .

$$\begin{aligned}
\Pr[\mathcal{B} \text{ returns } b] &= \Pr[b = 0 \wedge m_0 = \text{Dec}_{k^*}(c)] + \frac{1}{2} \cdot \Pr[b = 0 \wedge m_0 \neq \text{Dec}_{k^*}(c)] \\
&\quad + \frac{1}{2} \cdot \Pr[b = 1 \wedge m_0 \neq \text{Dec}_{k^*}(c)] \\
&= \frac{1}{2} \cdot \Pr[m_0 = \text{Dec}_{k^*}(\text{Enc}_k(m_0))] + \frac{1}{4} \cdot \Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_0))] \\
&\quad + \frac{1}{4} \cdot \Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_1))] \\
&= \frac{1}{4} + \frac{1}{4} \cdot \Pr[m_0 = \text{Dec}_{k^*}(\text{Enc}_k(m_0))] + \frac{1}{4} \cdot \Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_1))] \\
&\geq \frac{1}{4} + \frac{\epsilon(n)}{4} + \frac{1}{4} \cdot \Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_1))]
\end{aligned}$$

For arbitrary z , define \mathbb{M}_z as follows:

$$\mathbb{M}_z := \{m \in \{0, 1\}^n : \exists k^* \in \text{Supp}(\text{KeyGen}(1^n)) \text{ s.t. } \text{Dec}_{k^*}(z) = m\}$$

Since Dec is deterministic, for all z

$$|\mathbb{M}_z| \leq |\text{Supp}(\text{KeyGen}(1^n))|$$

Then

$$\Pr_{\substack{m_0, m_1 \leftarrow \{0, 1\}^n \\ k \leftarrow \text{KeyGen}(1^n)}}[m_0 \in \mathbb{M}_{\text{Enc}_k(m_1)}] \leq \frac{|\mathbb{M}_{\text{Enc}_k(m_1)}|}{2^n} \leq \frac{|\text{Supp}(\text{KeyGen}(1^n))|}{2^n} = \text{negl}(n)$$

which by definition of \mathbb{M} means that

$$\Pr_{\substack{m_0, m_1 \leftarrow \{0, 1\}^n \\ k \leftarrow \text{KeyGen}(1^n)}}[\exists k^* \text{ s.t. } m_0 = \text{Dec}_{k^*}(\text{Enc}_k(m_1))] = \text{negl}(n)$$

which means that in the hiding game

$$\Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_1))] = 1 - \text{negl}(n)$$

Substituting this back in the expression for \mathcal{B} guessing b

$$\begin{aligned}
\Pr[\mathcal{B} \text{ returns } b] &\geq \frac{1}{4} + \frac{\epsilon(n)}{4} + \frac{1}{4} \cdot \Pr[m_0 \neq \text{Dec}_{k^*}(\text{Enc}_k(m_1))] \\
&\geq \frac{1}{4} + \frac{\epsilon(n)}{4} + \frac{1}{4} \cdot (1 - \text{negl}(n)) \\
&\geq \frac{1}{2} + \frac{\epsilon(n)}{4} - \text{negl}(n)
\end{aligned}$$

which contradicts the hiding of the symmetric encryption scheme. □

The theorem follows immediately from the claim. □

B Proving Lemma 5.2

In this section, we provide a complete proof of Lemma 5.2. First, we list some preliminary theorems/claims that will be useful in proving the lemma.

B.1 Preliminary Claims

The following well-known theorem states that the boolean inner product function is a good extractor. This can be derived as a simple consequence of the leftover hash lemma.

Theorem B.1 (Inner Product is a Good Extractor). *Let X be a random variable distributed over $\{0, 1\}^n$ where $H_{\min}(X) \geq k$, for some $k > 0$. Then*

$$\text{SD}\left((U_n, \langle X, U_n \rangle), (U_n, U_1)\right) \leq 2^{(1-k)/2}$$

Proof. The hash function that takes a uniform seed $h \leftarrow \{0, 1\}^{n+1}$ and $x \in \{0, 1\}^n$ and outputs $\langle x, y \rangle \oplus z$ is known to be universal, where y indicates the first n bits of h and z is the last bit of h . Then by Theorem 3.1

$$\text{SD}\left((U_n, U_1, \langle X, U_n \rangle \oplus U_1), (U_n, U_2)\right) \leq \epsilon$$

for $\epsilon \leq 2^{(1-k)/2}$. We now consider the distribution obtained by replacing the last bit with the XOR of the last two bits in both distributions. Since performing the same operation on two distributions can only reduce the statistical distance,

$$\text{SD}\left((U_n, U_1, \langle X, U_n \rangle), (U_n, U_2)\right) \leq \epsilon$$

which implies that $\text{SD}\left((U_n, \langle X, U_n \rangle), (U_n, U_1)\right) \leq \epsilon$, as desired. \square

We will also use the following claim that relates the bias in an arbitrary single-bit distribution to its Shannon entropy.

Claim B.1. *Let X be an arbitrary d -biased distribution on a single bit, i.e.,*

$$\Pr[X = 1] = \frac{1+d}{2} \text{ and } \Pr[X = 0] = \frac{1-d}{2}.$$

Then,

$$H(X) \leq 1 - \frac{d^2}{2}$$

Additionally, whenever $d \leq 1/2$, we have

$$H(X) \geq 1 - d^2$$

Proof. Note that

$$\begin{aligned} H(X) &= \left(\frac{1+d}{2}\right) \cdot \log\left(\frac{2}{1+d}\right) + \left(\frac{1-d}{2}\right) \cdot \log\left(\frac{2}{1-d}\right) \\ &= 1 - (1/2) \cdot [(1+d) \log(1+d) + (1-d) \log(1-d)] \end{aligned}$$

Rearranging

$$1 - H(X) = (1/2) \cdot [\log(1-d^2) + d \cdot (\log(1+d) - \log(1-d))]$$

Since $\log y = \ln y / \ln 2$, we can rewrite as

$$1 - H(X) = (1/2 \ln 2) \cdot [\ln(1 - d^2) + d \cdot (\ln(1 + d) - \ln(1 - d))]$$

By the Taylor series expansion, for $y > -1$

$$\ln(1 + y) = y - y^2/2 + y^3/3 - y^4/4 + \dots$$

Noting that $0 < d < 1$ implies that $d, -d$ and $-d^2$ are all greater than -1 , we can apply this expansion to the \ln terms in the previous equation. This gives

$$\begin{aligned} 1 - H(X) &= (1/2 \ln 2) \cdot [(-d^2 - d^4/2 - d^6/3 - d^8/4 \dots) \\ &\quad + d \cdot ((d - d^2/2 + d^3/3 - d^4/4 \dots) - (-d - d^2/2 - d^3/3 - d^4/4 \dots))] \\ &= (1/2 \ln 2) \cdot [(-d^2 - d^4/2 - d^6/3 - d^8/4 \dots) + d \cdot (2d + 2d^3/3 + 2d^5/5 \dots)] \\ &= (1/2 \ln 2) \cdot [(d^2 + d^4/6 + d^6/15 + \dots + d^{2n}/(n \cdot (2n - 1)) + \dots)] \end{aligned}$$

We may pull out the common factor d^2 to obtain

$$1 - H(X) = (d^2/2 \ln 2) \cdot [(1 + d^2/6 + d^4/15 + \dots + d^{2n-2}/(n \cdot (2n - 1)) + \dots)] \quad (12)$$

which implies

$$1 - H(X) \geq (d^2/2 \ln 2)$$

or equivalently,

$$\begin{aligned} H(X) &\leq 1 - \frac{d^2}{2 \ln 2} \\ &\leq 1 - \frac{d^2}{2} \end{aligned}$$

Additionally, equation (12) implies

$$\begin{aligned} 1 - H(X) &\leq (d^2/2 \ln 2) \cdot [(1 + d^2 + d^4 + \dots + d^{2n} + \dots)] \\ &= \frac{d^2}{2 \ln 2 \cdot (1 - d^2)} \end{aligned}$$

If $d \leq 1/2$, then

$$\begin{aligned} 1 - H(X) &\leq \frac{d^2}{2 \ln 2 \cdot (1 - \frac{1}{4})} \\ &= \frac{2d^2}{3 \ln 2} \end{aligned}$$

which implies

$$H(X) \geq 1 - \frac{2d^2}{3 \ln 2} \geq 1 - d^2$$

which concludes the proof. □

B.2 Proof of the Lemma

In this section, we proceed to prove the main lemma.

Lemma 5.2 Restated. Let $X = \{X_n\}_{n \in \mathbb{N}}$ be an ensemble of distributions over $\{0, 1\}^n$ s.t. for sufficiently large $n \in \mathbb{N}$, there exists $x_n^* \in \{0, 1\}^n$ satisfying:

1. $\Pr_{x \leftarrow X_n}[x = x_n^*] \geq 1/6n$
2. $\forall x' \neq x_n^*, \Pr_{x \leftarrow X_n}[x = x'] \leq 2/n^{600}$

Let R_n be uniformly distributed on $\{0, 1\}^n$. Define $\alpha_{0,n}$ and $\alpha_{1,n}$ as follows.

$$\begin{aligned} \alpha_{0,n}(x, r) &:= \langle x, r \rangle \\ \alpha_{1,n}(x, r) &:= \begin{cases} U_1 & \text{if } x = x_n^* \\ \langle x, r \rangle & \text{otherwise} \end{cases} \end{aligned}$$

For $b \in \{0, 1\}$, define distribution $A_{b,n} := R_n, \alpha_{b,n}(X, R)$. Then for sufficiently large $n \in \mathbb{N}$,

$$H(A_{1,n}) - H(A_{0,n}) \geq 1/100n^2$$

To prove this lemma, we will first define random variables that indicate biases in the two distributions. Fix any sufficiently large n , and any $r \in \{0, 1\}^n, b \in \{0, 1\}$. Define

$$\begin{aligned} B_{r,n}^b &:= \left| \Pr_{x \leftarrow X_n} [\alpha_{b,n}(x, r) = 1] - \Pr_{x \leftarrow X_n} [\alpha_{b,n}(x, r) = 0] \right| \\ B_{r,n}^* &:= \left| \Pr_{x \leftarrow X_n} [\langle x, r \rangle = 1 | x \neq x_n^*] - \Pr_{x \leftarrow X_n} [\langle x, r \rangle = 0 | x \neq x_n^*] \right| \end{aligned}$$

Intuitively, the variable $B_{r,n}^b$ represents the bias of $\alpha_{b,n}(X_n, r)$ away from uniform. $B_{r,n}^*$ represents the bias conditioned on $X_n \neq x_n^*$ (which is identical for both distributions). Then,

$$\begin{aligned} B_{r,n}^0 &= \left| \Pr_{X_n} [x_n^*] + \Pr_{x \leftarrow X_n} [x \neq x_n^*] \left(\Pr_{x \leftarrow X_n} [\langle x, r \rangle = \langle x_n^*, r \rangle | x \neq x_n^*] - \Pr_{x \leftarrow X_n} [\langle x, r \rangle \neq \langle x_n^*, r \rangle | x \neq x_n^*] \right) \right| \\ &\geq \Pr_{X_n} [x_n^*] - \Pr_{x \leftarrow X_n} [x \neq x_n^*] \left| \left(\Pr_{x \leftarrow X_n} [\langle x, r \rangle = \langle x_n^*, r \rangle | x \neq x_n^*] - \Pr_{x \leftarrow X_n} [\langle x, r \rangle \neq \langle x_n^*, r \rangle | x \neq x_n^*] \right) \right| \\ &\geq \Pr_{X_n} [x_n^*] - \Pr_{x \leftarrow X_n} [x \neq x_n^*] \cdot B_{r,n}^* \end{aligned} \tag{13}$$

where the second inequality follows due to a triangle inequality. Furthermore,

$$\begin{aligned} B_{r,n}^1 &= \left| \Pr_{x \leftarrow X_n} [x \neq x_n^*] \left(\Pr_{x \leftarrow X_n} [\langle x, r \rangle = \langle x_n^*, r \rangle | x \neq x_n^*] - \Pr_{x \leftarrow X_n} [\langle x, r \rangle \neq \langle x_n^*, r \rangle | x \neq x_n^*] \right) \right| \\ &= \Pr_{x \leftarrow X_n} [x \neq x_n^*] \cdot B_{r,n}^* \end{aligned} \tag{14}$$

The proof of the lemma follows immediately from the next two claims, which analyse the distribution X in different ways depending on whether or not it is heavily biased towards x_n^* .

Claim B.2. Consider any $n \in \mathbb{N}$ for which X_n is heavily biased towards x_n^* , that is,

$$\Pr_{x \leftarrow X_n}[x = x_n^*] > 1 - 1/n.$$

Then, if n is large enough, it holds that for all $r \in \{0, 1\}^n$,

$$H(\alpha_1(X, r)) - H(\alpha_0(X, r)) \geq 1/100n^2.$$

Proof. Since $B_{r,n}^* \leq 1$ for every $n, r \in \mathbb{N}$, equations (13) and (14) imply

$$B_{r,n}^0 \geq \Pr_{X_n}[x_n^*] - \Pr_{x \leftarrow X_n}[x \neq x_n^*] \cdot B_{r,n}^* \geq 1 - 2/n$$

and

$$B_{r,n}^1 = \Pr_{x \leftarrow X_n}[x \neq x_n^*] \cdot B_{r,n}^* \leq 1/n.$$

Therefore, by Claim B.1 applied to $\alpha_0(X, r)$ and $\alpha_1(X, r)$, for large enough n

$$\begin{aligned} H(\alpha_0(X, r)) &\leq 1 - \frac{(1 - 2/n)^2}{2} \leq 2/3 \\ H(\alpha_1(X, r)) &\geq 1 - 1/n^2 \end{aligned}$$

Thus for all $r \in \{0, 1\}^n$,

$$H(\alpha_1(X, r)) - H(\alpha_0(X, r)) \geq 1/3 - 1/n^2 \geq 1/100n^2$$

as desired. □

Claim B.3. Consider any $n \in \mathbb{N}$ for which

$$\Pr_{x \leftarrow X_n}[x = x_n^*] \leq 1 - 1/n.$$

Then, if n is large enough, it holds that for all $r \in \{0, 1\}^n$,

$$H(A_{1,n}) - H(A_{0,n}) \geq 1/100n^2.$$

Proof. The premise of the claim implies

$$\Pr_{x \leftarrow X_n}[x \neq x_n^*] \geq 1/n.$$

Let X'_n denote the distribution $(X_n | X_n \neq x_n^*)$. For all $x' \in \text{Supp}(X'_n)$,

$$\begin{aligned} \Pr_{x \leftarrow X'_n}[x = x'] &= \frac{\Pr_{x \leftarrow X_n}[x = x']}{\Pr_{x \leftarrow X_n}[x \neq x_n^*]} \\ &\leq n \cdot \Pr_{x \leftarrow X_n}[x = x'] \\ &\leq 2n^{-599} \end{aligned}$$

which implies that for all $x' \in \text{Supp}(X'_n)$,

$$H_{X'_n}(x') \geq 599 \log n - \log 2 \geq 598 \log n \quad (\text{for } n \geq 2)$$

which implies that $H_{\min}(X'_n) \geq 598 \log n$. Then, Theorem B.1 applied to X'_n implies that:

$$\text{SD}\left((R_n, \langle X'_n, R_n \rangle), (R_n, U_1)\right) \leq \sqrt{2n^{-598}} < n^{-200}$$

The statistical distance may be rewritten as

$$\text{SD}\left((R_n, \langle X'_n, R_n \rangle), (R_n, U_1)\right) = \sum_r \frac{1}{2} \cdot \Pr_{R_n}[r] \cdot \left| \Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 1] - 1/2 \right| \leq n^{-200}$$

Since $\Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 1] = 1 - \Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 0]$, we can rewrite the above equation as

$$\sum_r \frac{1}{4} \cdot \Pr_{R_n}[r] \cdot \left| \Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 1] - \Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 0] \right| \leq n^{-200}$$

Substituting $B_{r,n}^* = |\Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 1] - \Pr_{x \leftarrow X'_n}[\langle x, r \rangle = 0]|$, we obtain

$$\sum_r \frac{1}{4} \cdot \Pr_{R_n}[r] \cdot B_{r,n}^* \leq n^{-200}$$

By a Markov argument on r ,

$$\Pr_{r \leftarrow R_n} [B_{r,n}^* \geq n^{-100}] \leq n^{-100}$$

Let $\mathbb{B}_n := \{r | B_{r,n}^* \leq n^{-100}\}$. Therefore, $\Pr_{r \leftarrow R_n} [r \notin \mathbb{B}_n] \leq n^{-100}$. Moreover, equations (13) and (14) imply that for all $r \in \mathbb{B}$,

$$\begin{aligned} B_{r,n}^0 &\geq \left(\Pr_{X_n}[x_n^*] - \Pr_{x \leftarrow X_n} [x \neq x_n^*] \cdot B_{r,n}^* \right) \geq 1/6n - n^{-100} \geq 1/7n, \text{ and} \\ B_{r,n}^1 &= \Pr_{x \leftarrow X_n} [x \neq x_n^*] \cdot B_{r,n}^* \leq n^{-100} \end{aligned}$$

Finally, by applying the chain rule,

$$\begin{aligned} H(A_{1,n}) - H(A_{0,n}) &= H(\alpha_{1,n}(X_n, R_n), R_n) - H(\alpha_{0,n}(X_n, R_n), R_n) \\ &= H(\alpha_{1,n}(X_n, R_n) | R_n) + H(R_n) - H(\alpha_{0,n}(X_n, R_n) | R_n) - H(R_n) \\ &= H(\alpha_{1,n}(X_n, R_n) | R_n) - H(\alpha_{0,n}(X_n, R_n) | R_n) \\ &= \sum_r \Pr_{R_n}[r] \cdot (H(\alpha_{1,n}(X_n, r)) - H(\alpha_{0,n}(X_n, r))) \\ &= \sum_{r \in \mathbb{B}_n} \Pr_{R_n}[r] \cdot (H(\alpha_{1,n}(X_n, r)) - H(\alpha_{0,n}(X_n, r))) \\ &\quad + \sum_{r \notin \mathbb{B}_n} \Pr_{R_n}[r] \cdot (H(\alpha_{1,n}(X_n, r)) - H(\alpha_{0,n}(X_n, r))) \\ &\geq \sum_{r \in \mathbb{B}_n} \Pr_{R_n}[r] \cdot (H(\alpha_{1,n}(X_n, r)) - H(\alpha_{0,n}(X_n, r))) - \sum_{r \notin \mathbb{B}_n} \Pr_{R_n}[r] \\ &\geq \sum_{r \in \mathbb{B}_n} \Pr_{R_n}[r] \cdot (H(\alpha_{1,n}(X_n, r)) - H(\alpha_{0,n}(X_n, r))) - \Pr_{r \leftarrow R_n} [r \notin \mathbb{B}_n] \end{aligned}$$

Applying Claim B.1 to the distributions $\alpha_{0,n}(X_n, r)$ and $\alpha_{1,n}(X_n, r)$ with biases $B_{r,n}^0$ and $B_{r,n}^1$ respectively, and noting that $B_{r,n}^1 < \frac{1}{2}$ for $n \geq 2$,

$$\begin{aligned} H(A_{1,n}) - H(A_{0,n}) &\geq \sum_{r \in \mathbb{B}_n} \Pr_R[r] \cdot \left(\frac{(B_{r,n}^0)^2}{2} - (B_{r,n}^1)^2 \right) - \Pr_{r \leftarrow R} [r \notin \mathbb{B}_n] \\ &\geq \sum_{r \in \mathbb{B}_n} \Pr_R[r] \cdot \left(\frac{1}{98n^2} - \frac{1}{n^{100}} \right) - \Pr_{r \leftarrow R} [r \notin \mathbb{B}_n] \\ &\geq \Pr_{r \leftarrow R} [r \in \mathbb{B}_n] \cdot \left(\frac{1}{99n^2} \right) - \Pr_{r \leftarrow R} [r \notin \mathbb{B}_n] \\ &\geq (1 - 1/n^{100}) \cdot \left(\frac{1}{99n^2} \right) - (1/n^{100}) \\ &\geq 1/100n^2 \quad (\text{for large enough } n) \end{aligned}$$

This completes the proof of the claim. \square