Equivalence Between Average-Case Hardness of Learning and Cryptography for Mixed Quantum States

Alexandru Cojocaru*† Laura Lewis*‡

Abstract

The relationship between cryptography and learning theory has long been a central theme in the foundations of theoretical computer science: cryptographic primitives can imply hardness of learning, while hardness of learning can in turn be used to construct cryptographic schemes. Recent works have begun exploring analogous connections in the quantum setting, relating the average-case hardness of learning quantum states (AHL) to cryptographic primitives such as one-way state generators (OWSG). Despite recent progress exploring this for pure states, the relationship for mixed states has remained an open question.

In this work, we prove that the existence of AHL for mixed quantum states is equivalent to the existence of inefficiently verifiable one-way state generators (IV-OWSGs). As a consequence, this relates mixed-state AHL to EFI pairs. Moreover, as a corollary of existing results, we obtain a separation between IV-OWSGs and OWSGs relative to the SWAP oracle.

1 Introduction

Classically, there is a fundamental connection between cryptography and learning theory, where they are often viewed as each other's antithesis [Riv91]. Namely, cryptographic constructions can imply the impossibility of efficient learning [KV94, KS09, SZB21, DLSS14, Val84, HN23, NR06, Ang88, NY15], and the hardness of learning can be used for novel constructions of cryptographic primitives [IL90, OS16, BFKL93, Reg09, GPV08, LPR10, MR09]. For instance, early results show the intractability of PAC learning algorithms for classes of Boolean functions based on constructions of public-key cryptosystems [KV94, KS09]. Meanwhile, a prominent example in the opposite direction is the Learning with Errors problem, which is widely believed to be computationally intractable to solve, even for quantum computers, leading to many (post-quantum) cryptographic schemes based on it [Reg09, GPV08, LPR10] (see, e.g., [MR09] for a more detailed review). Several works in the classical world have also shown the equivalence between the non-existence of one-way functions and the existence of efficient (average-case) learning algorithms for various problems [IL90, NR06, NY15, HN23].

In contrast to this bountiful literature in the classical world, results investigating the intersection of quantum learning and cryptography have been significantly more scarce. Some works have established computational hardness for learning classes of quantum states/unitaries from the existence of pseudorandom quantum states/unitaries [ZLK+24, YE23, FPVY25]. Recently, researchers have begun to explore the converse: designing quantum cryptography from the *average-case hardness of learning (AHL)* classes of quantum states [HH24, FGSY25,

*Authors are listed alphabetically.
†University of Edinburgh. Email: a.cojocaru@ed.ac.uk
‡University of Edinburgh. UC Berkeley. Email: lllewis@berkeley.edu

HHM25, NLO+26][1]. Notably, [HH24, FGSY25] prove that AHL for classes of pure states is equivalent to the existence of pure *one-way state generators (OWSGs)*, i.e., efficient quantum algorithms that generate pure states that are easy to produce but computationally hard to invert. However, this equivalence is unknown for the case of *mixed states*, leading us to the central question of this work:

Is the average-case hardness of learning mixed states equivalent to mixed one-way state generators?

This can be viewed as a more complete analogue of the equivalence between the average-case hardness of learning and the existence of one-way functions [IL90, NR06, NY15, HN23] in the classical setting. Prior work [HH24] provided initial progress on this question by showing that the existence of EFI[2] implies AHL for classes of mixed states by using the equivalence of EFI with a variant of OWSGs, namely secretly-verifiable statistically-invertible OWSGs [MY22a]. However, it is not clear if proving the converse is possible with this approach, as this variant of OWSGs requires an orthogonality property (statistical invertibility) which an arbitrary AHL instance does not satisfy in general.

In this work, we identify the appropriate cryptographic counterpart of mixed-state AHL: *inefficiently verifiable one-way state generators (IV-OWSGs)*. Our main result shows that a class of mixed quantum states is hard to learn on average if and only if IV-OWSGs exist.

Our equivalence also transfers assumptions and applications between quantum learning and cryptography. Combining our main result with known relationships between IV-OWSGs and EFI pairs [MMWY24] shows that EFI implies mixed-state AHL, while a sufficiently strong, exponentially-hard AHL assumption implies EFI. Consequently, AHL in this parameter regime implies statistically binding quantum commitments and semi-honest oblivious transfer [BCQ22]. In the other direction, any assumption known to imply IV-OWSGs, in particular, assumptions yielding EFI pairs, also gives rise to a class of mixed states that is hard to learn on average. Thus, AHL provides a learning-theoretic lens on cryptographic assumptions that may lie below classical one-way functions.

Finally, by combining existing results, it can be shown that IV-OWSGs exist relative to the SWAP oracle [GZ25], whereas OWSGs do not. The construction proceeds from statistically secure single-copy pseudorandom states in the common Haar random state model [CCS25] and applies the lifting result of [GZ25] to obtain the corresponding primitive in the SWAP model. This oracle separation shows that our characterization cannot, in general, be upgraded from IV-OWSGs to efficiently-verifiable OWSGs by a relativizing argument.

We remark that this does not contradict the equivalence proved in the pure state case [HN23, FGSY25] because verification can always be made efficient for pure states (see Lemma B.1 in [MY22a]). Preliminary versions of our results appeared in [Lew25].

1.1 Problem Definition and Results

The average-case hardness of learning (AHL) problem intuitively states that a class of quantum states is hard to learn on average if, given polynomially many copies of a state randomly sampled from the class, any efficient adversary can only learn a state close to it with at most some probability δ. More formally, following [HH24], AHL is defined as follows[3]:

Definition 1 (Average-case hardness of learning (AHL); Definition 3.1 in [HH24])**.** *Let $\mathcal{C}$ be a class of n-qubit quantum states ρ_x, indexed by $x \in \{0,1\}^n$, where each ρ_x is generatable by a QPT*

[1]We remark that [HHM25] studies a different hardness-of-learning task, namely quantum distribution learning, where one wishes to learn classical bitstrings sampled from a distribution generated by a QPT algorithm.

[2]Here, an EFI pair is a pair of efficiently generatable quantum states that are statistically far but computationally indistinguishable [BCQ22].

[3]We have made some minor modifications to the definition from [HH24] to align better with notation used in the learning theory literature, but this definition nonetheless captures the same premise.

algorithm. The class $\mathcal{C}$ is (ϵ, δ)-hard-to-learn *if there exists an efficiently sampleable distribution $\mathcal{D}_\lambda$ over $\{0,1\}^n$ such that for all $t(\lambda) = \mathrm{poly}(\lambda)$ and uniform QPT learning algorithms $\mathcal{A}$,*

$$\Pr_{\substack{\rho_x \leftarrow \mathcal{D}_\lambda(\mathcal{C}) \\ y \leftarrow \mathcal{A}(\rho_x^{\otimes t(\lambda)})}} [\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)] \leq \delta(\lambda), \tag{1.1}$$

where we use $\mathcal{D}_\lambda(\mathcal{C})$ to denote the distribution over $\mathcal{C}$ induced by $\mathcal{D}_\lambda$, which is over the labels of the states. When $\epsilon(\lambda), \delta(\lambda) = 1/\mathrm{poly}(\lambda)$, we simply say that $\mathcal{C}$ is hard on average.

In the above definition, we think of λ as the security parameter, while $n = \mathrm{poly}(\lambda)$.

When $\epsilon(\lambda), \delta(\lambda) = 1/\mathrm{poly}(\lambda)$, this assumption states that efficient algorithms can only learn the target state up to inverse-polynomial accuracy with inverse-polynomial probability. One can also strengthen the assumption by choosing $\epsilon(\lambda) = 1/\mathrm{poly}(\lambda), \delta(\lambda) = 1/2^{cn}$ for some constant $c > 0$, which is still reasonable from our current understanding of tomography. Note that one must have $c < 1$, as otherwise, the AHL assumption becomes trivial [FGSY25]. We refer the reader to [FGSY25] for further discussion of noteworthy parameter regimes.

We also remark that this definition only considers *proper* learning, i.e., learning algorithms are only allowed to output a hypothesis from the class $\mathcal{C}$. This is consistent with previous work [HH24, FGSY25]. There is also a notion of improper learning, where the algorithm can approximate ρ_x using any (efficient) description of a quantum state, instead of being restricted to only states in $\mathcal{C}$, but we do not consider improper learning in this work.

The other main object we consider is the *one-way state generator* [MY22a]. Informally, a one-way state generator consists of a set of algorithms $(\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ defined as follows:

- KeyGen is a QPT algorithm that outputs a classical key k.
- $\mathsf{StateGen}(k)$ is a QPT algorithm that, on input key k, outputs a (potentially mixed) quantum state ρ_k.
- $\mathsf{Ver}(k', \rho_k)$ is a QPT algorithm that, on input ρ_k and a bitstring k', outputs $\top$ or $\bot$.

The security of a OWSG intuitively says that no QPT adversary can invert the output of $\mathsf{StateGen}$; in other words, given polynomially-many copies of ρ_k, it should be hard to recover k. An *inefficiently-verifiable* OWSG (IV-OWSG) is the same as a OWSG except the verification is not restricted to be QPT [MMWY24, BJ24]. We refer the reader to Definitions 3 and 4 for formal definitions. Our main result is as follows.

Theorem 1 (Equivalence between mixed AHL and IV-OWSGs)**.** *There exists a class of states that is average-case hard to learn if and only if there exist IV-OWSGs.*

As an immediate corollary, because IV-OWSG is equivalent to EFI with an exponential loss in the reduction (see Theorems 4.3 and 5.2 in [MMWY24]), we obtain that AHL is equivalent to EFI, also up to an exponential loss in the reduction.

Corollary 1 (Relationship between mixed AHL and EFI)**.** *If EFI pairs exist, then there exists a class of states that is average-case hard to learn. Moreover, if there exists a class of states that is average-case (ϵ, δ)-hard-to-learn, for $\epsilon(\lambda) = 1/\mathrm{poly}(\lambda)$ and $\delta(\lambda) = 2^{-0.75\lambda}$, then EFI pairs exist.*

This corollary imports known applications of EFI into the learning setting. In particular, AHL with the stated exponential hardness implies statistically binding quantum commitments and semi-honest oblivious transfer [BCQ22]. Conversely, EFI and any other assumption implying IV-OWSGs yield classes of mixed quantum states that are hard to learn on average.

More broadly, our results position AHL as a candidate foundational assumption for quantum cryptography in *Microcrypt*, where useful quantum cryptographic primitives may exist

even in the absence of quantum-secure one-way functions. Identifying natural hardness assumptions that suffice for cryptography but are potentially weaker than one-way functions is a central goal in quantum cryptography. For example, recent proposals have drawn on the hardness of learning quantum states and the hardness of implementing non-collapsing measurements [FGSY25, MSY25]. Our equivalence places mixed-state AHL within this broader program and motivates studying whether AHL can hold without quantum-secure one-way functions.

In addition, we collect several existing results in the literature to show a separation between OWSG and IV-OWSG relative to the SWAP oracle [GZ25] (see Definition 9). This shows that a strengthening of Theorem 1 to an equivalence between OWSGs and AHL is not possible.

Theorem 2. *Relative to the SWAP oracle, IV-OWSG exist while OWSG do not exist.*

We briefly discuss the proof of Theorem 1.

AHL implies IV-OWSG. Consider an AHL instance (as in Definition 1) specified by a class $\mathcal{C} = \{\rho_x\}$ of quantum states, an efficiently sampleable distribution $\mathcal{D}_\lambda$, and parameters $\epsilon(\lambda), \delta(\lambda)$. As a natural approach to construct an IV-OWSG, one may consider taking KeyGen to sample from $\mathcal{D}_\lambda$ and $\mathsf{StateGen}(x)$ to output the corresponding state ρ_x from $\mathcal{C}$. However, the correct choice of verification algorithm is not immediately clear: the canonical choice of verification algorithm for pure states [MY22a] does not apply for mixed states; quantum state certification algorithms [BOW19] require too many copies of the unknown state; the SWAP test for mixed states no longer reflects the trace distance between two states. Our key observation is that one can use quantum hypothesis selection [BO21] to instantiate the verification algorithm. Algorithms for quantum hypothesis selection are computationally inefficient, resulting in an IV-OWSG.

OWSG implies AHL. Consider an IV-OWSG defined by three algorithms $(\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$. For this direction, one can consider the simple construction of AHL: take the distribution $\mathcal{D}_\lambda$ to sample from the same distribution as KeyGen and the class of states $\mathcal{C}$ to be the same set of states output by $\mathsf{StateGen}$. Intuitively, the result of the theorem is reasonable: if one cannot invert a quantum state (i.e., find its key), then it should also be hard to learn a classical description of it, which is precisely specified by the key. However, in AHL, the problem definition (Definition 1) allows the learning algorithm to recover keys of nearby states as well. Meanwhile, a priori, the verification algorithm of the OWSG only accepts exactly the correct key. To reconcile these definitional differences, previous work [HH24] instead considers OWSGs with an additional property of statistical invertibility: any states with different keys are far apart in trace distance (see Definition 6). This circumvents the issue by ensuring that the learning algorithm in AHL can only learn exactly the correct key, as any other key will not be within the allowed margin of error. In contrast, our crucial observation is that the verification algorithm in fact also accepts keys corresponding to nearby states.

Organization

In Section 2.1, we introduce basic cryptographic primitives that we refer to throughout the work. In Section 2.2, we state a guarantee from quantum learning theory which we find useful. In Sections 3 and 4, we prove Theorem 1. Specifically, in Section 3, we prove that AHL implies IV-OWSG, and in Section 4, we prove that IV-OWSG implies AHL. In Section 5, we compile results from the literature to show an oracle separation between IV-OWSG and OWSG.

Acknowledgments

The authors thank Damiano Abram, Mohammed Barhoush, Taiga Hiroka, Dominik Leichtle, Thomas Vidick, and Chirag Wadhwa, for helpful discussions at various stages of this project. A.C. acknowledges support from the National Science Foundation grant CCF-1813814, from the AFOSR under Award Number FA9550-20-1-0108 and from the Quantum Advantage Pathfinder project. L.L. was supported by a Marshall Scholarship and a U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0026073.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

AI Use Disclosure

All ideas, results, and writing were generated by the authors without AI use. ChatGPT 5.6 Pro was used to check the correctness of the proofs after they were already completed by the authors.

2 Preliminaries

2.1 Cryptographic Tools and Definitions

In this section, we review some ideas from cryptography, which we will use throughout the paper. First, we define a one-way function, which is the most fundamental primitive in classical cryptography. Throughout, we let λ denote the security parameter and $n = \mathrm{poly}(\lambda)$. Hereafter, by a negligible function we mean a function that decays faster than any polynomial.

Definition 2 (One-Way Function (OWF))**.** *Let $\mathcal{F} = \{f_\lambda\}_{\lambda\in\mathbb{N}}$ be a family of efficiently-computable functions $f_\lambda : \{0,1\}^\lambda \to \{0,1\}^{\ell(\lambda)}$. $\mathcal{F}$ is a* one-way function *if for every polynomial-time probabilistic algorithm $\mathcal{A}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every security parameter $\lambda \in \mathbb{N}$*

$$\Pr_{\substack{x\leftarrow\{0,1\}^\lambda \\ y\leftarrow f_\lambda(x)}}[\mathcal{A}(1^\lambda, y) \in f_\lambda^{-1}(y)] \leq \mathsf{negl}(\lambda). \tag{2.1}$$

In other words, a OWF is a function which is easy to compute but hard to invert for any polynomial-time algorithm. In this work, we mainly focus on the *quantum version* of one-wayness. In particular, one can consider quantum objects that behave similarly to a OWF, e.g., a keyed quantum state for which it is hard to efficiently recover the key. This is exactly the notion of a one-way state generator. In the following definition, we parameterize the security.

Definition 3 (η-Secure One-Way State Generators (η-Secure OWSGs); Definition 3.1 in [MY22a])**.** *An η-secure* one-way state generator (OWSG) *is a set of algorithms (*$\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver}$*) such that*

- $\mathsf{KeyGen}(1^\lambda)$ *is a QPT algorithm tha t, on input the security parameter λ, outputs a classical key $k \in \{0,1\}^n$.*

- $\mathsf{StateGen}(k)$ *is a QPT algorithm that, on input key* $k \in \{0,1\}^n$*, outputs a (potentially mixed) quantum state* ρ_k*.*
- $\mathsf{Ver}(k', \rho_k)$ *is a QPT algorithm that, on input* ρ_k *and a bitstring* k'*, outputs* $\top$ *or* $\bot$*.*

Moreover, these algorithms satisfy the following properties

- *(Correctness) For any key* $k \in \{0,1\}^n$*, then*

$$\Pr_{\rho_k \leftarrow \mathsf{StateGen}(k)}[\top \leftarrow \mathsf{Ver}(k, \rho_k)] \geq 1 - \mathsf{negl}(\lambda). \tag{2.2}$$

- *(Security) For any uniform QPT adversary* $\mathcal{A}$ *and polynomial* t*,*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(k)}}[\top \leftarrow \mathsf{Ver}(\mathcal{A}(\rho_k^{\otimes t(\lambda)}), \rho_k)] \leq \eta. \tag{2.3}$$

If StateGen *only outputs pure states, we call these* pure OWSGs*. Moreover, when* $\eta = \mathsf{negl}(\lambda)$*, we simply call these OWSGs.*

It is instructive to note that for pure OWSGs, without loss of generality, one may consider Ver as the following algorithm [MY22a, MY22b]: on input k' and $|\phi_k\rangle$, measure ϕ_k with the projective measurement $\{|\phi_{k'}\rangle\langle\phi_{k'}|, I - |\phi_{k'}\rangle\langle\phi_{k'}|\}$. If the result is $|\phi_{k'}\rangle\langle\phi_{k'}|$, output $\top$; otherwise, output $\bot$.

We also note that in some definitions of OWSGs, the correctness is only required to be probabilistic over the choice of key. We modify this slightly to hold for all keys instead, following [FGSY25]. There are also several other variations of OWSGs. First, we consider inefficiently-verifiable OWSGs (IV-OWSGs).

Definition 4 (η-Secure Inefficiently-verifiable OWSGs (IV-OWSGs); Definition 15 in [BJ24] or Definition 3.2 in [MMWY24])**.** *An* inefficiently-verifiable one-way state generator (IV-OWSG) *is defined in the same way as a OWSG, except that* Ver *is allowed to be inefficient.*

We also introduce secretly-verifiable OWSGs (SV-OWSGs) [MY22a]. Here, one can think of the verification algorithm Ver as being given two classical keys k, k' instead of a classical key k' and a copy of an unknown quantum state ρ_k as in Definition 3. Then, without loss of generality, one can replace Ver with simply checking if $k = k'$.

Definition 5 (Secretly-verifiable OWSGs (SV-OWSGs); Definition 7.1 in [MY22a])**.** *A* secretly-verifiable one-way state generator (SV-OWSG) *is a set of algorithms* $(\mathsf{KeyGen}, \mathsf{StateGen})$*, where* $\mathsf{KeyGen}, \mathsf{StateGen}$ *are defined in the same way as for OWSGs. The only difference is that security is defined as*

- *(Security) For any uniform QPT adversary* $\mathcal{A}$ *and polynomial* t*,*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(k)}}[k \leftarrow \mathcal{A}(\rho_k^{\otimes t(\lambda)})] \leq \mathsf{negl}(\lambda). \tag{2.4}$$

Moreover, there is another variant which is similar to SV-OWSGs but additionally requires that all of the states generated are far apart in trace distance.

Definition 6 (Secretly-verifiable statistically-invertible OWSGs (SV-SI-OWSGs); Definition 7.3 in [MY22a])**.** *A* secretly-verifiable and statistically-invertible one-way state generator (SV-SI-OWSG) *is a set of algorithms* $(\mathsf{KeyGen}, \mathsf{StateGen})$ *such that*

- $\mathsf{KeyGen}(1^\lambda)$ *is a QPT algorithm that, on input the security parameter* λ*, outputs a classical key* $k \in \{0,1\}^n$*.*
- $\mathsf{StateGen}(k)$ *is a QPT algorithm that, on input key* $k \in \{0,1\}^n$*, outputs a (potentially mixed) quantum state* ρ_k*.*

Moreover, these algorithms satisfy the following properties

- *(Statistical invertibility) For any* k, k' *with* $k \neq k'$*, then*

$$\mathrm{d}_{\mathrm{tr}}(\rho_k, \rho_{k'}) \geq 1 - \mathsf{negl}(\lambda). \tag{2.5}$$

- *(Computational non-invertibility) For any uniform QPT adversary* $\mathcal{A}$ *and polynomial* t*,*

$$\Pr_{\substack{k \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_k \leftarrow \mathsf{StateGen}(k)}} [k \leftarrow \mathcal{A}(\rho_k^{\otimes t(\lambda)})] \leq \mathsf{negl}(\lambda). \tag{2.6}$$

While this may seem like a strange requirement at first, SV-SI-OWSGs are interesting because they are equivalent to EFI, a primitive believed to be a minimal assumption for quantum cryptography.

Definition 7 (EFI; Definition 3.1 in [BCQ22])**.** *An* EFI pair *is a uniform QPT algorithm* EFI *such that*

- *(Efficient generation)* EFI *is a uniform QPT algorithm that on input* $(1^\lambda, b)$ *for a security parameter* $\lambda \in \mathbb{N}$ *and a bit* $b \in \{0,1\}$*, outputs a (potentially mixed) quantum state* ρ_b*. That is,* $\mathsf{EFI}(1^\lambda, b) = \rho_b$*.*
- *(Statistically far)*

$$\mathrm{d}_{\mathrm{tr}}(\rho_0, \rho_1) \geq 1 - \mathsf{negl}(\lambda). \tag{2.7}$$

- *(Computational indistinguishability) For any uniform QPT adversary* $\mathcal{A}$*,*

$$|\Pr[1 \leftarrow \mathcal{A}(\rho_0)] - \Pr[1 \leftarrow \mathcal{A}(\rho_1)]| \leq \mathsf{negl}(\lambda). \tag{2.8}$$

Theorem 3 (Theorem 7.7 in [MY22a])**.** *SV-SI-OWSGs exist if and only if EFI pairs exist.*

2.2 Quantum Learning Theory

First, recall the definition of trace distance between two quantum states.

Definition 8 (Trace distance)**.** *Let* ρ, σ *be two quantum states. The* trace distance *between* ρ *and* σ *is*

$$\mathrm{d}_{\mathrm{tr}}(\rho, \sigma) \triangleq \frac{1}{2}\|\rho - \sigma\|_1, \tag{2.9}$$

where $\|M\|_1 = \mathrm{tr}\left(\sqrt{M^\dagger M}\right)$ *denotes the trace norm.*

We also find the following guarantee from quantum learning theory useful. It provides a sample complexity upper bound for the task of *hypothesis selection*, i.e., finding the best approximation to an unknown state from a class of hypothesis states.

Theorem 4 (Theorem 1.5 in [BO21]). *There is a quantum algorithm that, given m fixed hypothesis states $\mathcal{C} = \{\sigma_1, \ldots, \sigma_m\}$, where $\sigma_i \in \mathbb{C}^{d\times d}$, parameters $0 < \epsilon, \delta < 1/2$, and access to unentangled copies of a state $\rho \in \mathbb{C}^{d\times d}$, where $\rho \in \mathcal{C}$, uses*

$$N(\epsilon,\delta) = \min\left\{\frac{(\log^2 m + L_1)(\log d)}{\epsilon^4}\cdot\mathcal{O}(L_1), \frac{\log^3 m + \log(L_2/\delta)\cdot\log m}{\epsilon^2}\cdot\mathcal{O}(L_2\log(L_2/\delta))\right\} \tag{2.10}$$

copies of ρ, where $L_1 = \log\left(\frac{\log d}{\delta\epsilon}\right)$ and $L_2 = \log(1/\epsilon)$ and has the following guarantee: with probability at least $1-\delta$, it outputs k such that

$$\mathrm{d}_{\mathrm{tr}}(\rho,\sigma_k) \leq \epsilon. \tag{2.11}$$

3 AHL implies IV-OWSG

In this section, we prove that AHL implies IV-OWSG. For pure states, it is easy to construct IV-OWSGs from AHL because there is a canonical choice of verification algorithm (simply measure the projector). However, the main difficulty for mixed states is the absence of such a verification algorithm, where, e.g., quantum state certification algorithms or the SWAP test require too many copies of the unknown state. Our key observation is that one can use quantum hypothesis selection (Theorem 4) to instantiate the verification algorithm for mixed states. However, hypothesis selection is computationally inefficient, resulting in an IV-OWSG.

Proposition 1 (AHL $\Rightarrow$ IV-OWSG). *If there exists a class of states that is average-case hard to learn (as in Definition 1), then there exists an IV-OWSG.*

Proof. Suppose there exists an AHL instance specified by a class $\mathcal{C}$ of n-qubit quantum states indexed by bitstrings in $\mathcal{X} \subseteq \{0,1\}^n$, an efficiently sampleable distribution $\mathcal{D}_\lambda$ over $\mathcal{X}$, and $\epsilon(\lambda), \delta(\lambda) = 1/\mathrm{poly}(\lambda)$. By definition of AHL, for all polynomials $t(\lambda) = \mathrm{poly}(\lambda)$ and all uniform QPT learning algorithms $\mathcal{A}$, then

$$\Pr_{\substack{\rho_x\leftarrow\mathcal{D}_\lambda(\mathcal{C})\\ y\leftarrow\mathcal{A}(\rho_x^{\otimes t(\lambda)})}}[\mathrm{d}_{\mathrm{tr}}(\rho_x,\rho_y)\leq\epsilon(\lambda)] \leq \delta(\lambda). \tag{3.1}$$

Here, we use $\mathcal{D}_\lambda(\mathcal{C})$ to denote the distribution over $\mathcal{C}$ induced by $\mathcal{D}_\lambda$, which is over the labels of the states instead. We construct an IV-OWSG as follows:

- $\mathsf{KeyGen}(1^\lambda)$ outputs $x \in \{0,1\}^n$ sampled from $\mathcal{D}_\lambda$.
- $\mathsf{StateGen}(x)$ outputs $\rho_x^{\otimes N}$, where $\rho_x \in \mathcal{C}$ and $N = N(\epsilon(\lambda)/2, 1/2^n)$, where N is defined in Theorem 4.
- $\mathsf{Ver}(y, \rho_x^{\otimes N})$ first runs the hypothesis selection algorithm from Theorem 4 on $\rho_x^{\otimes N}$ and the class $\mathcal{C}$ to obtain an output x'. Then, it outputs $\top$ iff $\mathrm{d}_{\mathrm{tr}}(\rho_y,\rho_{x'}) \leq \epsilon(\lambda)/2$.

Note that $N = N(\epsilon(\lambda)/2, 1/2^n) = \mathrm{poly}(n)$ because $N(\epsilon,\delta) = \mathrm{poly}(1/\epsilon, \log(1/\delta), \log(m))$, and the size of the hypothesis class is $m = |\mathcal{C}| \leq 2^n$. Thus, $\mathsf{StateGen}$ is a QPT algorithm because it only needs to output a $\mathrm{poly}(n)$ copies of an efficiently generatable state (since all states in $\mathcal{C}$ can be generated by a QPT algorithm).

We also remark that this is only an IV-OWSG because the hypothesis selection algorithm from [BO21] is computationally inefficient. Also, because the verification is allowed to be inefficient and y, x' are known along with a classical description of the class $\mathcal{C}$, then the trace distance calculation can be performed.

First, we show correctness, i.e., for any $x \leftarrow \mathsf{KeyGen}(1^\lambda)$, then

$$\Pr_{\rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}[\top \leftarrow \mathsf{Ver}(x, \rho_x^{\otimes N})] \geq 1 - \mathsf{negl}(\lambda). \tag{3.2}$$

By the guarantees of hypothesis selection with our choice of parameters (Theorem 4), then running hypothesis selection on $\rho_x^{\otimes N}$ and the class $\mathcal{C}$ outputs some x' such that $\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_{x'}) \leq \epsilon(\lambda)/2$ with probability at least $1 - 1/2^n$. Then, the verification algorithm simply checks if $\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_{x'}) \leq \epsilon(\lambda)/2$, which clearly holds as long as hypothesis selection is successful. Thus, it is clear that we have

$$\Pr_{\rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}[\top \leftarrow \mathsf{Ver}(x, \rho_x^{\otimes N})] \geq 1 - \frac{1}{2^n}. \tag{3.3}$$

It remains to show security. We want to show that for any uniform QPT adversary $\mathcal{A}$ and any polynomial t,

$$\Pr_{\substack{x \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}}[\top \leftarrow \mathsf{Ver}(\mathcal{A}(\rho_x^{\otimes (N \cdot t(\lambda))}), \rho_x^{\otimes N})] \leq \mathsf{negl}(\lambda). \tag{3.4}$$

Suppose for the sake of contradiction that there exists a uniform QPT algorithm $\mathcal{B}$ and a polynomial t such that

$$\Pr_{\substack{x \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}}[\top \leftarrow \mathsf{Ver}(\mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))}), \rho_x^{\otimes N})] > \delta(\lambda) + \frac{1}{2^n}. \tag{3.5}$$

To arrive at a contradiction, we show that $\mathcal{B}$ can also learn $\mathcal{C}$, contradicting AHL. Define the set

$$S_{\mathrm{good}} \triangleq \left\{(x, y) \in \mathcal{X}^2 : \mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda), \rho_x \leftarrow \mathcal{C}(x), \rho_y \leftarrow \mathcal{C}(y)\right\}. \tag{3.6}$$

Recall here that $\mathcal{X}$ is the set of bitstring labels for states in the class $\mathcal{C}$. Moreover, we use $\rho_x \leftarrow \mathcal{C}(x)$ to denote the state $\rho_x \in \mathcal{C}$ corresponding to the bitstring label x. Then, we have

$$\delta(\lambda) + \frac{1}{2^n} \tag{3.7}$$

$$< \Pr_{\substack{x \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}}[\top \leftarrow \mathsf{Ver}(\mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))}), \rho_x^{\otimes N})] \tag{3.8}$$

$$= \sum_{x,y} \Pr\left[x \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \Pr_{\rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] \Pr_{\rho_x^{\otimes N} \leftarrow \mathsf{StateGen}(x)}[\top \leftarrow \mathsf{Ver}(y, \rho_x^{\otimes N})] \tag{3.9}$$

$$= \sum_{x,y \in S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[\top \leftarrow \mathsf{Ver}(y, \rho_x^{\otimes N})] \tag{3.10}$$

$$+ \sum_{x,y \notin S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[\top \leftarrow \mathsf{Ver}(y, \rho_x^{\otimes N})] \tag{3.11}$$

$$\leq \sum_{x,y \in S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] + \frac{1}{2^n} \sum_{x,y \notin S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] \tag{3.12}$$

$$= \sum_{x,y \in S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] + \frac{1}{2^n} \Pr_{\substack{\rho_x \leftarrow \mathcal{D}_\lambda(\mathcal{C}) \\ y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})}}[\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) > \epsilon(\lambda)] \tag{3.13}$$

$$\leq \sum_{x,y \in S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)}[y \leftarrow \mathcal{B}(\rho_x^{\otimes (N \cdot t(\lambda))})] + \frac{1}{2^n}, \tag{3.14}$$

where in the third line, we use the definition of our construction of IV-OWSGs. In the fifth line, we use the definition of our verification procedure, which we claim outputs $\top$ for $(x,y) \notin S_{\mathrm{good}}$ with probability at most $1/2^n$. This is because by Theorem 4 and our choice of parameters, with probability at least $1-1/2^n$, hypothesis selection outputs some x' such that $\mathrm{d_{tr}}(\rho_x,\rho_{x'}) \leq \epsilon(\lambda)/2$. Moreover, for $(x,y) \notin S_{\mathrm{good}}$, then $\mathrm{d_{tr}}(\rho_x,\rho_y) > \epsilon(\lambda)$. Then, by reverse triangle inequality,

$$\mathrm{d_{tr}}(\rho_y,\rho_{x'}) \geq |\,\mathrm{d_{tr}}(\rho_y,\rho_x) - \mathrm{d_{tr}}(\rho_x,\rho_{x'})| > \epsilon(\lambda)/2. \tag{3.15}$$

Hence, because $\mathsf{Ver}(y,\rho_x^{\otimes N})$ outputs $\top$ iff $\mathrm{d_{tr}}(\rho_y,\rho_{x'}) \leq \epsilon(\lambda)/2$, then in this case, the verification outputs $\bot$ unless hypothesis selection fails, which occurs with probability at most $1/2^n$.

Finally, writing the AHL security probability, we have, for a polynomial $t'(\lambda) = N\cdot t(\lambda)$, then

$$\Pr_{\substack{\rho_x\leftarrow\mathcal{D}_\lambda(\mathcal{C})\\ y\leftarrow\mathcal{B}(\rho_x^{\otimes t'(\lambda)})}}[\mathrm{d_{tr}}(\rho_x,\rho_y)\leq\epsilon(\lambda)] = \sum_{x,y}\Pr[x\leftarrow\mathcal{D}_\lambda]\Pr_{\rho_x\leftarrow\mathcal{C}(x)}[y\leftarrow\mathcal{B}(\rho_x^{\otimes t'(\lambda)})]\Pr_{\substack{\rho_x\leftarrow\mathcal{C}(x)\\ \rho_y\leftarrow\mathcal{C}(y)}}[\mathrm{d_{tr}}(\rho_x,\rho_y)\leq\epsilon(\lambda)] \tag{3.16}$$

$$= \sum_{x,y\in S_{\mathrm{good}}}\Pr[x\leftarrow\mathcal{D}_\lambda]\Pr_{\rho_x\leftarrow\mathcal{C}(x)}[y\leftarrow\mathcal{B}(\rho_x^{\otimes t'(\lambda)})] \tag{3.17}$$

$$> \delta(\lambda), \tag{3.18}$$

which is a contradiction. In the last line, we use Equation (3.14). Thus, we have proven that for all uniform QPT algorithms $\mathcal{A}$ and polynomials t, then

$$\Pr_{\substack{x\leftarrow\mathsf{KeyGen}(1^\lambda)\\ \rho_x^{\otimes N}\leftarrow\mathsf{StateGen}(x)}}[\top\leftarrow\mathsf{Ver}(\mathcal{A}(\rho_x^{\otimes(N\cdot t(\lambda))}),\rho_x^{\otimes N})] \leq \delta(\lambda)+\frac{1}{2^n} \tag{3.19}$$

for an inverse-polynomial δ. This can be amplified to negligible security by Theorem 3.7 of [MY22a]. □

4 IV-OWSG implies AHL

In this section we show that IV-OWSG implies AHL. Intuitively, this claim makes sense: if one cannot invert the state (i.e., find its key), then it should also be hard to learn a classical description of it, which is precisely specified by the key. Despite this intuition, we were not able to find a proof of this in the literature, and [HH24] only achieves a relationship between OWSGs and AHL in the mixed case for the different notion of SV-SI-OWSGs (see Definition 6). Their proof also uses the statistically-invertible property of these OWSGs crucially. Instead, we provide a simple proof showing that IV-OWSG implies AHL, even in the mixed state case.

Before presenting the proof of our result, we note that an initial, weaker result can be obtained easily by stitching together results from the literature.

Proposition 2. *Let $D>0$ be a constant and let $\eta(\lambda)\geq(0.5+D)\lambda$. If $2^{-\eta(\lambda)}$-secure IV-OWSGs exist, then AHL exists.*

Proof. This follows by combining Theorem 5.2 from [MMWY24] and Theorem 5.3 in [HH24]. Namely, Theorem 5.2 in [MMWY24] proves that exponentially-secure IV-OWSGs imply EFI. Then, Theorem 5.3 in [HH24] shows that EFI implies AHL for mixed states. □

This proves that IV-OWSG implies AHL, although with an exponential loss in the reduction. In contrast, we prove that standard IV-OWSGs are sufficient to construct AHL, with no loss in the reduction. The key intuition is that the Ver algorithm of the OWSG also accepts keys corresponding to states that are close in trace distance to the target state, just like AHL.

Proposition 3 (IV-OWSG $\Rightarrow$ AHL)**.** *If IV-OWSGs exist, then there exists a class of states that is average-case hard to learn (as in Definition 1).*

Proof. Suppose there exists an IV-OWSG specified by algorithms $(\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$. We construct a class of states that is hard to learn on average from this. This requires constructing a distribution $\mathcal{D}_\lambda$ and a class of states $\mathcal{C}$. We simply define $\mathcal{D}_\lambda$ as sampling from the same distribution as $\mathsf{KeyGen}(1^\lambda)$ and $\mathcal{C}$ as the same set of states output by $\mathsf{StateGen}$. We want to show that this class is indeed an AHL instance, i.e., there exist $\epsilon(\lambda), \delta(\lambda) = 1/\mathrm{poly}(\lambda)$ such that for all polynomials t and uniform QPT algorithms $\mathcal{A}$,

$$\Pr_{\substack{\rho_x \leftarrow \mathcal{D}_\lambda(\mathcal{C}) \\ y \leftarrow \mathcal{A}(\rho_x^{\otimes t(\lambda)})}} [\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)] \leq \delta(\lambda). \tag{4.1}$$

Suppose for the sake of contradiction that for all inverse polynomial ϵ, δ, there exists a polynomial t and QPT learner $\mathcal{B}$ such that

$$\Pr_{\substack{\rho_x \leftarrow \mathcal{D}_\lambda(\mathcal{C}) \\ y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})}} [\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)] > \delta(\lambda). \tag{4.2}$$

To arrive at a contradiction, we will show that $\mathcal{B}$ can also break the security of the OWSG. Define the set

$$S_{\mathrm{good}} \triangleq \{(x, y) \in \mathcal{X}^2 : \mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda),\ \rho_x \leftarrow \mathcal{C}(x), \rho_y \leftarrow \mathcal{C}(y)\}. \tag{4.3}$$

Notice that by our assumption on the adversary $\mathcal{B}$ we have

$$\Pr_{\substack{\rho_x \leftarrow \mathcal{D}_\lambda(\mathcal{C}) \\ y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})}} [\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)] \tag{4.4}$$

$$= \sum_{(x,y) \in \mathcal{X}^2} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)} [y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})] \Pr_{\substack{\rho_x \leftarrow \mathcal{C}(x) \\ \rho_y \leftarrow \mathcal{C}(y)}} [\mathrm{d}_{\mathrm{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)] \tag{4.5}$$

$$= \sum_{(x,y) \in S_{\mathrm{good}}} \Pr[x \leftarrow \mathcal{D}_\lambda] \Pr_{\rho_x \leftarrow \mathcal{C}(x)} [y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})] \tag{4.6}$$

$$> \delta(\lambda). \tag{4.7}$$

Now, we can lower bound the probability of verification passing for the OWSG using $\mathcal{B}$'s output:

$$\Pr_{\substack{x \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \rho_x \leftarrow \mathsf{StateGen}(x)}} [\top \leftarrow \mathsf{Ver}(\mathcal{B}(\rho_x^{\otimes t(\lambda)}), \rho_x)] \tag{4.8}$$

$$= \sum_{(x,y) \in \mathcal{X}^2} \Pr\left[x \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \Pr_{\rho_x \leftarrow \mathsf{StateGen}(x)} [y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})] \Pr_{\rho_x \leftarrow \mathsf{StateGen}(x)} [\top \leftarrow \mathsf{Ver}(y, \rho_x)] \tag{4.9}$$

$$\geq \sum_{(x,y) \in S_{\mathrm{good}}} \Pr\left[x \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \Pr_{\rho_x \leftarrow \mathsf{StateGen}(x)} [y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})] \Pr_{\rho_x \leftarrow \mathsf{StateGen}(x)} [\top \leftarrow \mathsf{Ver}(y, \rho_x)] \tag{4.10}$$

$$\geq \sum_{(x,y) \in S_{\mathrm{good}}} \Pr\left[x \leftarrow \mathsf{KeyGen}(1^\lambda)\right] \Pr_{\rho_x \leftarrow \mathsf{StateGen}(x)} [y \leftarrow \mathcal{B}(\rho_x^{\otimes t(\lambda)})] \cdot \left(\Pr_{\rho_y \leftarrow \mathsf{StateGen}(y)} [\top \leftarrow \mathsf{Ver}(y, \rho_y)] - \epsilon(\lambda) \right) \tag{4.11}$$

$$> \delta(\lambda)\left(1 - \mathsf{negl}(\lambda) - \epsilon(\lambda)\right). \tag{4.12}$$

In the fourth line, we use that, by definition of S_{good}, then for all $(x, y) \in S_{\mathrm{good}}$, $\mathrm{d_{tr}}(\rho_x, \rho_y) \leq \epsilon(\lambda)$. Thus, $|\operatorname{tr}(\Pi\rho_x) - \operatorname{tr}(\Pi\rho_y)| \leq \epsilon(\lambda)$ for any POVM Π. In general, the verification procedure $\mathsf{Ver}(y, \rho_x)$ performs a (computationally inefficient) POVM on ρ_x, so the probability of it outputting $\top$ is given by $\operatorname{tr}(\Pi\rho_x)$. In the last inequality, we used that $\Pr[\top \leftarrow \mathsf{Ver}(x, \rho_x)] \geq 1 - \mathsf{negl}(\lambda)$ for all x (see our definition of OWSGs in Definition 3) and Equation (4.7) with our definition of the AHL instance. Because δ, ϵ are inverse polynomial, this is a contradiction to the security of the IV-OWSG. □

5 Separation Between IV-OWSG and OWSG

In this section, we stitch together several results in the literature to show an oracle separation between IV-OWSG and OWSG. For our purposes, it is important that these results hold against uniform adversaries. This separation holds relative to the SWAP oracle [GZ25], which we define as follows.

Definition 9 (Oracle models; Definition 3.1 of [GZ25])**.** *Let $\mathcal{D} = \{\mathcal{D}_n\}_{n\in\mathbb{N}}$ be a family of distributions over pure states on n-qubits, which are elements of $\mathbb{C}^{2^n}$. Let $|\phi_n\rangle$ be a state sampled from $\mathcal{D}$ during initialization, before any party receives oracle access. Consider the space $\mathbb{C}^{2^n+1}$ which will be the space spanned by $\mathbb{C}^{2^n}$ of $\mathcal{D}_n$ and an orthogonal basis vector $|0\rangle$.*

- *Define the* common Haar random state model *over $\mathcal{D}$ to be the oracle $\mathrm{CHRS}_n : \mathbb{C} \to \mathbb{C}^{2^n}$ which acts as*
$$|0\rangle \mapsto |\phi_n\rangle. \tag{5.1}$$
- *Define the* swap structured state model (SWAP) *over $\mathcal{D}$ to be the oracle $\mathrm{SWAP}_n : \mathbb{C}^{2^n+1} \to \mathbb{C}^{2^n+1}$ which acts as*
$$\mathrm{SWAP}_n = I - |0\rangle\langle 0| - |\phi_n\rangle\langle\phi_n| + |0\rangle\langle\phi_n| + |1\rangle\langle 0|. \tag{5.2}$$

First, we state a result from [GZ25].

Theorem 5 (Corollary C.7 of [GZ25])**.** *There does not exist a OWSG relative to SWAP.*

Now, we show the following theorem, which establishes the oracle separation.

Theorem 6. *IV-OWSG exist relative to SWAP.*

To show this, we state two results, which, when combined, give us the above statement. Notably, all of these results hold for uniform adversaries.

Theorem 7 (Theorem 4.1 of [CCS25])**.** *Statistically-secure 1PRS exist relative to SWAP.*

Proof. By Theorem 4.1 of [CCS25], statistically-secure 1PRS exist in the common Haar random state model. By applying Corollary 8.4 of [GZ25], this implies a construction of statistically-secure 1PRS in the SWAP model. □

Theorem 8. *If 1PRS exist, then IV-OWSG exists.*

Proof. This follows by combining several existing results in the literature. First, Theorem 1.1 of [MY22b] shows that if 1PRS exist, then non-interactive quantum commitments with computational hiding and statistical binding exist. Note that non-interactive quantum commitments trivially imply interactive quantum commitments. Next, Theorem 4.6 of [BCQ22] shows that if quantum commitments with computational hiding and statistical binding exist, then EFI exists. Finally, Theorem 4.3 of [MMWY24] shows that if EFI exists, then IV-OWSGs exist. □

References

[Ang88] Dana Angluin. Queries and concept learning. *Machine learning*, 2:319–342, 1988.

[BCQ22] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. *arXiv preprint arXiv:2209.04101*, 2022.

[BFKL93] Avrim Blum, Merrick Furst, Michael Kearns, and Richard J Lipton. Cryptographic primitives based on hard learning problems. In *Annual international cryptology conference*, pages 278–291. Springer, 1993.

[BJ24] Rishabh Batra and Rahul Jain. Commitments are equivalent to statistically-verifiable one-way state generators. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1178–1192. IEEE, 2024.

[BO21] Costin Bădescu and Ryan O'Donnell. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1398–1411, 2021.

[BOW19] Costin Bădescu, Ryan O'Donnell, and John Wright. Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 503–514, 2019.

[CCS25] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: constructing and separating quantum pseudorandomness. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 108–137. Springer, 2025.

[DLSS14] Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 441–448, 2014.

[FGSY25] Bill Fefferman, Soumik Ghosh, Makrand Sinha, and Henry Yuen. The hardness of learning quantum circuits and its cryptographic applications. *arXiv preprint arXiv:2504.15343*, 2025.

[FPVY25] Ben Foxman, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. Random unitaries in constant (quantum) time. *arXiv preprint arXiv:2508.11487*, 2025.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 197–206, 2008.

[GZ25] Eli Goldin and Mark Zhandry. Translating between the common haar random state model and the unitary model. In *Annual International Cryptology Conference*, pages 269–300. Springer, 2025.

[HH24] Taiga Hiroka and Min-Hsiu Hsieh. Computational complexity of learning efficiently generatable pure states. *arXiv preprint arXiv:2410.04373*, 2024.

[HHM25] Taiga Hiroka, Min-Hsiu Hsieh, and Tomoyuki Morimae. Hardness of quantum distribution learning and quantum cryptography. *arXiv preprint arXiv:2507.01292*, 2025.

[HN23] Shuichi Hirahara and Mikito Nanashima. Learning in pessiland via inductive inference. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 447–457. IEEE, 2023.

[IL90] Russell Impagliazzo and Levin LA. No better ways to generate hard np instances than picking uniformly at random. In *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science*, pages 812–821. IEEE, 1990.

[KS09] Adam R Klivans and Alexander A Sherstov. Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer and System Sciences*, 75(1):2–12, 2009.

[KV94] Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994.

[Lew25] Laura Lewis. Computational complexity in quantum learning tasks. *Edinburgh Research Archive, Informatics Thesis and Dissertation Collection*, 2025.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology–EUROCRYPT 2010: 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30–June 3, 2010. Proceedings 29*, pages 1–23. Springer, 2010.

[MMWY24] Giulio Malavolta, Tomoyuki Morimae, Michael Walter, and Takashi Yamakawa. Exponential quantum one-wayness and efi pairs. In *International Conference on Security and Cryptography for Networks*, pages 121–138. Springer, 2024.

[MR09] Daniele Micciancio and Oded Regev. Lattice-based cryptography. In *Post-quantum cryptography*, pages 147–191. Springer, 2009.

[MSY25] Tomoyuki Morimae, Yuki Shirakawa, and Takashi Yamakawa. Quantum cryptography and hardness of non-collapsing measurements, 2025.

[MY22a] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint arXiv:2210.03394*, 2022.

[MY22b] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Annual International Cryptology Conference*, pages 269–295. Springer, 2022.

[NLO+26] Pradeep Niroula, Minzhao Liu, Sivaprasad Omanakuttan, David Amaro, Shouvanik Chakrabarti, Soumik Ghosh, Zichang He, Yuwei Jin, Fatih Kaleoglu, Steven Kordonowy, et al. Digital signatures with classical shadows on near-term quantum computers. *arXiv preprint arXiv:2602.04859*, 2026.

[NR06] Moni Naor and Guy N Rothblum. Learning to impersonate. In *Proceedings of the 23rd international conference on Machine learning*, pages 649–656, 2006.

[NY15] Moni Naor and Eylon Yogev. Bloom filters in adversarial environments. In *Annual Cryptology Conference*, pages 565–584. Springer, 2015.

[OS16] Igor C Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds and pseudorandomness. *arXiv preprint arXiv:1611.01190*, 2016.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.

[Riv91] Ronald L Rivest. Cryptography and machine learning. In *International Conference on the Theory and Application of Cryptology*, pages 427–439. Springer, 1991.

[SZB21] Min Jae Song, Ilias Zadik, and Joan Bruna. On the cryptographic hardness of learning single periodic neurons. *Advances in neural information processing systems*, 34:29602–29615, 2021.

[Val84] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[YE23] Lisa Yang and Netta Engelhardt. The complexity of learning (pseudo) random dynamics of black holes and other chaotic systems. *arXiv preprint arXiv:2302.11013*, 2023.

[ZLK+24] Haimeng Zhao, Laura Lewis, Ishaan Kannan, Yihui Quek, Hsin-Yuan Huang, and Matthias C Caro. Learning quantum states and unitaries of bounded gate complexity. *PRX Quantum*, 5(4):040306, 2024.