

Instantiating Microcrypt: Obstacles and opportunities via tailored state certification

Jose Carrasco^{*1}, Jens Eisert^{1,2,3}, Soumik Ghosh^{4, 12}, Dominik Hangleiter^{11,12}, Nicky Kai Hong Li^{5,6,7}, and Ryan Sweke^{†8,9,10}

¹Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany

²Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany

³Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany

⁴Center for Theoretical Physics, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA

⁵Technische Universität Wien, Atominstitut, Stadionallee 2, 1020 Vienna, Austria

⁶Vienna Center for Quantum Science and Technology, TU Wien, 1020 Vienna, Austria

⁷Institute for Quantum Optics and Quantum Information (IQOQI), Austrian Academy of Sciences, Boltzmanngasse 3, 1090 Vienna, Austria

⁸African Institute for Mathematical Sciences (AIMS), South Africa

⁹Department of Mathematical Sciences, Stellenbosch University, Stellenbosch 7600, South Africa

¹⁰National Institute for Theoretical and Computational Sciences (NITheCS), South Africa

¹¹Institute for Theoretical Physics, ETH Zürich, Switzerland

¹²Simons Institute for the Theory of Computing, University of California at Berkeley, USA

September 23, 2026

Abstract

Recent work has introduced the Hamiltonian phase state (HPS) assumptions, which postulate that Hamiltonian phase states can be used to instantiate pseudorandom and one-way state generators [BHH+25]. Additionally, it has been conjectured that these assumptions can be true, even if one-way functions do not exist. This is exciting, because if true, then the HPS assumptions provide a route to the instantiation of cryptography which is genuinely in Microcrypt – the world of (quantum) cryptographic protocols and primitives which can exist even if one-way functions do not. In this work we falsify this conjecture, by proving that if the HPS assumptions are true, then one-way functions exist. While this removes the possibility of instantiating genuine Microcrypt protocols and primitives with Hamiltonian phase states, it shows that the HPS assumptions provide novel inherently quantum assumptions for the construction of classical cryptography. The technical contribution that allows us to do this is a method for the construction of one-way puzzles from one-way state generators via tailored "measure first, ask later" state certification protocols. This generalizes prior constructions of one-way puzzles from one-way state generators via classical shadows and allows us to relate properties of the resulting one-way puzzle to properties of the state certification protocol used in the construction. Specifically, if the state certification protocol admits efficient classical post-processing then one obtains an efficiently verifiable one-way puzzle, and if the state certification protocol can be efficiently classically simulated in a certain sense, then one obtains a classical one-way puzzle, which implies one-way functions. Indeed, the latter observation allows us to prove that the HPS assumptions imply one-way functions, by exploiting properties of recent state certification protocols for phase states. The former observation provides a new toolbox for the construction of efficiently verifiable one-way puzzles by exploiting tailored state certification protocols for pseudorandom and one-way state generators.

^{*}jose.carrasco@fu-berlin.de

[†]rsweke@aims.ac.za

Contents

1	Introduction	3
1.1	Our contributions	4
1.1.1	Obstacles: Hamiltonian phase state assumptions imply one-way functions	4
1.1.2	Opportunities: Efficiently verifiable one-way puzzles via tailored state certification protocols	6
1.1.3	Towards simpler one-way functions from the Search HPS assumption	8
1.2	Related work	8
1.3	Discussion and open questions	11
1.4	A short story on the highs and lows of developing this work	12
1.5	Structure of this work	13
2	Preliminaries	13
2.1	Classical cryptographic primitives	13
2.2	Microcrypt primitives	15
2.3	Hamiltonian phase states and Microcrypt	18
3	State certification	20
4	Certifiable Microcrypt primitives	22
5	One-way puzzles from certifiable one-way state generators	24
6	One-way functions from simulable certifiable one-way state generators	26
6.1	OWFs from simulable certifiable OWSGs via distributional OWFs	26
6.2	Direct OWF construction via efficient verifiability	28
7	Hamiltonian phase state assumptions imply one-way functions	32
7.1	Computationally efficient certification of phase states	33
7.2	Classical simulability of phase state certification	35
7.3	Proof of Theorem 7.1	39
8	Towards concrete instantiations of efficiently verifiable one-way puzzles	40
A	Proof of Corollary 2.8	42
B	Physical implementation of Hamiltonian phase states	43
C	Proof of Theorem 7.3	44
D	Proof of Lemma 6.3	46
E	The one-way function obtained from Search HPS	46

1 Introduction

The existence of *one-way functions* (OWFs) is a minimal assumption for classical cryptography. More specifically, all meaningful classical cryptographic primitives imply OWFs, and thus OWFs are necessary (but not sufficient) for classical cryptography. Remarkably, however, over the last few years, evidence has emerged that meaningful cryptography may be possible in a *quantum* world even if $P = NP$ (and hence OWFs do not exist). More specifically:

1. A series of breakthrough works [Kre21; KQS+23; KQT25] have culminated in the construction of oracle worlds in which $P = NP$, but quantum cryptographic primitives such as quantum computable OWFs and pseudorandom state generators (PRSG) do exist.
2. Khurana and Tomer have shown that one can construct one-way puzzles (OWPs) from “quantum advantage assumptions” for quantum random sampling schemes, which can be true even if $P = NP$ [KT25], under extremely mild complexity-theoretic assumptions.

In addition, a variety of works have also shown how to construct non-trivial quantum cryptographic protocols from such inherently quantum primitives, including variants of bit-commitments, secure multiparty computation and digital signatures amongst others [AQY22; AGQ+22; MY22; GY25; KT24; ALY23; MY24; CGG24; KQT25].

In line with the “world building” tradition of theoretical cryptography, the name *Microcrypt* has been given to the world in which OWFs do not exist, but genuinely quantum cryptographic primitives such as quantum computable OWFs, OWPs and PRSGs (amongst others) do exist. Motivated by how unexpected and exciting this world is, recent years have witnessed significant effort to map this world [Sat24], i.e.:

1. To propose and define Microcrypt primitives, understand which primitives imply which, and which complexity-theoretic conditions are necessary for a primitive to exist.
2. To understand which cryptographic protocols can be built from which Microcrypt primitives.
3. To propose *concrete instantiations* of Microcrypt primitives under suitable *quantum* assumptions, which are plausibly independent of OWFs.

Given this work, we now understand that Microcrypt has a hierarchical structure – i.e., it consists of multiple distinct “subworlds”. Each such subworld is defined by a necessary upper bound on the computational power of quantum computers, and has an associated candidate minimal primitive – i.e., a primitive which can be constructed from every other primitive in the subworld, and is thus necessary for its existence. More specifically, as discussed and illustrated in Ref. [GMM+25], we have the following Microcrypt hierarchy:

1. **Quantumania:** The cryptographic primitives which are broken if $BQP = QCMA$. Nearly every primitive in Quantumania implies the existence of *efficiently-verifiable* OWPs (EV-OWPs) [KT24; KT25], which can therefore be considered as a minimal primitive for this world. This world is particularly interesting, given that it contains a variety of cryptographic protocols which can be executed with local quantum computation but only classical communication – i.e., QCCC cryptography [CGG24]. Additionally, Ref. [KQT25] has recently provided an oracle relative to which $P = NP$, yet quantum computable trapdoor OWFs (which imply EV-OWPs) exist, thereby giving evidence that Quantumania may indeed be non-empty even if OWFs do not exist.
2. **Countcrypt:** The cryptographic primitives which are not necessarily broken if $BQP = QCMA$, but are broken if $BQP = PP$. Amongst others, this world contains PRSGs [JLS18], one-way state generators (OWSGs) and standard OWPs [KT24; KT25] (which appear to be a minimal primitive for this world).
3. **Nanocrypt:** The cryptographic primitives which may exist even if $BQP = PP$, and for which efficiently generated, statistically far-apart, *computationally indistinguishable* (EFI) pairs [BCQ23] appear to be a minimal primitive.

In order to realize any of the cryptography possible in any subworld of Microcrypt, one requires proposals for *concrete instantiations* of the relevant minimal primitive. Of course, any such concrete instantiation will come with its own assumptions, and importantly, for the instantiation to be genuinely in Microcrypt, its assumptions should

not imply the existence of OWFs. With the goal of providing such concrete instantiations under inherently quantum assumptions, Ref. [BHH+25] recently proposed the *Hamiltonian phase state (HPS) assumptions*. Specifically, these assumptions postulate that ensembles of Hamiltonian phase states provide concrete PRSGs and OWSGs, which can then be used to construct OWPs [KT24]. Importantly, Ref. [BHH+25] also provided some preliminary evidence that the HPS assumptions may be true, even if OWFs do not exist. This is exciting, as if this is the case, then the HPS assumptions indeed provide a clear foundation for the instantiation of the Countcrypt subworld of Microcrypt. However, given that the HPS assumptions are new and untested, the natural question that we address in this work is the following:

Can the Hamiltonian phase state assumptions be true if one-way functions do not exist?

We resolve this question in the *negative*, showing that if the Hamiltonian phase state assumptions are true, then OWFs exist. This unfortunately provides a clear obstacle for instantiating Microcrypt, by showing that the HPS assumptions are not sufficient for the instantiation of genuine Microcrypt cryptography. It does however mean that the HPS assumptions provide a novel *inherently quantum* assumption for the construction of classical cryptography, complementing recent work aiming precisely to provide such assumptions [PQS25; LPQ+26].

The primary technical contribution that allows us to resolve the question above is a method for the construction of OWP variants from OWSGs via "measure first, ask later" state certification protocols for the set of states defining the OWSG. This generalizes and abstracts the existing construction of OWPs from OWSGs via classical shadows [KT24]. Importantly, it also allows us to relate properties of the resulting OWP to properties of the state certification protocol that was used. In particular, if the state certification measurement protocol can be classically simulated in a specific sense, then the resulting OWP is a *classical* OWP, which then implies OWFs [KT24; KT25]. Indeed, this fact is what allows us to prove that the HPS assumptions imply OWFs, by exploiting properties of existing state certification protocols for (Hamiltonian) phase states [HPS25].

Importantly, however, the above method for constructing OWPs from OWSGs and state certification protocols also allows us to show that if the state certification protocol admits efficient classical post-processing, then the resulting OWP is *efficiently verifiable*. This provides a new toolbox for constructing *efficiently verifiable* OWPs, and may open up new opportunities for identifying concrete ensembles of states with which to instantiate Quantumania.

1.1 Our contributions

Motivated by the above question, we make a variety of contributions, which can be loosely categorized into those providing *obstacles* and those providing *opportunities* (hopefully) for the concrete instantiation of Microcrypt.

1.1.1 Obstacles: Hamiltonian phase state assumptions imply one-way functions

We postpone formal definitions to Section 2.3, but informally, the *Hamiltonian phase state (HPS) assumptions*, recently introduced in Ref. [BHH+25], are the following:

Assumption 1.1: (Hamiltonian phase state assumptions [BHH+25])

1. **Decision HPS assumption** (informal version of Assumption 2.18): There exists a set of Hamiltonian phase states, and an efficiently sampleable distribution over this set, which can be used to construct a pseudorandom state generator.
2. **Search HPS assumption** (informal version of Assumption 2.20): There exists a set of Hamiltonian phase states, and an efficiently sampleable distribution over this set, which can be used to construct a one-way state generator.

We note that the Search HPS assumption is a weaker assumption, in that the Decision HPS assumption implies the Search HPS assumption. Given that there are a wide variety of cryptographic primitives and protocols that can be built from PRSGs and OWSGs [Sat24; GMM+25], the HPS assumptions provide a route for the *concrete instantiation* of these primitives and protocols via a family of quantum states that can be easily prepared experimentally via




Figure 1: Summary of the main definitions and constructions developed in this work (where n is taken as the security parameter of the OWSG). Theorem 1.2 below is proven by showing that there exists an η -simulable and computationally efficient "measure first, ask later" state certification protocol for the OWSG obtained via the Search HPS assumption, and therefore one can construct a OWF by the implications of either Theorem 1.6 or Theorem 1.7. As we discuss in Section 1.1.3, Theorem 1.7 yields a simpler OWF, by virtue of not passing through a distributonal OWF.

instantaneous quantum polynomial-time (IQP) circuits. Additionally, Ref. [BHH+25] has provided some preliminary evidence that the HPS assumptions are *independent of the existence of OWFs* – i.e., the HPS assumptions could be true, even if OWFs do not exist. If correct, then the concrete protocols and primitives constructed via the HPS assumptions would indeed be in Microcrypt.

In this work, we establish that, contrary to prior evidence, the HPS assumptions are *not* independent of OWFs. Specifically, in Section 7, we prove the following:

Theorem 1.2: (One-way function from Hamiltonian phase state assumptions) *If the Search HPS assumption is true, then one can explicitly construct a quantum-secure one-way function.*

Given that the Decision HPS assumption implies the Search HPS assumption, an immediate corollary of the above result is that one can also construct one-way functions from the Decision HPS assumption. Unfortunately, Theorem 1.2 shows that one cannot use the HPS assumptions to obtain cryptography in Microcrypt. However, given that we are able to explicitly construct a OWF from either HPS assumption, it also shows that the HPS assumptions provide a novel, inherently quantum assumption for the construction of classical cryptography. As such, Theorem 1.2 simultaneously provides an *obstacle* for the concrete instantiation of Microcrypt, and an *opportunity* for the instantiation of Minicrypt from novel inherently quantum assumptions. This complements recent work aimed at providing inherently quantum constructions for cryptography [PQS25; LPQ+26], and opens up a wide variety of open questions and directions for future research, which we discuss in more detail in Section 1.3.

At a high level, our proof of Theorem 1.2 is enabled by a new toolbox for constructing *variants* of OWPs from OWSGs and copy efficient "measure first, ask later" state certification protocols for the set of states defining the OWSG, which is illustrated in Figure 1. In particular, this method abstracts and generalizes an existing construction of OWPs from OWSGs and classical shadows [KT24], and as illustrated in Figure 1, allows us to relate properties of the resulting OWP to properties of the state certification protocol used in the construction. Importantly, when the state certification protocol admits classically efficient post-processing then the resulting OWP is *efficiently verifiable*. If the state certification protocol is classically simulable in a certain sense, then the OWP can be used to construct a OWF. With this in hand, we prove Theorem 1.2 by showing that existing state certification protocols for phase states [HPS25] satisfy all the necessary efficiency and simulability properties for (a) the construction of a OWP from

the OWSG obtained from the HPS assumptions, and (b) the construction of an explicit OWF from this OWP. We stress that while the prior construction of OWPs from OWSGs via classical shadows [KT24] allows one to obtain a OWP from the OWSG obtained from the Search HPS assumption, it is not clear if one can use this construction to then obtain a OWF from this OWP.

Given that our method for the construction of OWP variants from OWSGs and tailored state certification protocols provides a potential route for the instantiation of EV-OWPs (and therefore Quantumania), we discuss these techniques in more detail in Section 1.1.2 below.

1.1.2 Opportunities: Efficiently verifiable one-way puzzles via tailored state certification protocols

As mentioned above, apart from allowing us to prove Theorem 1.2, our technique for constructing OWP variants from OWSGs and state certification protocols, illustrated in Figure 1, also provides a new toolbox for constructing *efficiently-verifiable* OWPs, and therefore a potential route towards instantiating Quantumania. Central to this toolbox is the notion of a "measure first, ask later" state certification protocol for a set of states $\{|\phi_k\rangle\}$.

Definition 1.3: ("Measure first, ask later" state certification protocol for $\{|\phi_k\rangle\}$) Let $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$ be a set of n -qubit states. Consider a QPT algorithm $\mathcal{M}(|\psi\rangle^{\otimes t}) \rightarrow s \in \{0, 1\}^{g(t,n)}$ and let $\mathcal{F} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\text{Accept}, \text{Reject}\}$ be such that $\mathcal{F}(\tilde{k}, \cdot) = \text{reject}$ for all $\tilde{k} \notin \mathbb{K}$. We say that $(\mathcal{M}, \mathcal{F})$ is a "measure first, ask later" state certification protocol for $\{|\phi_k\rangle\}$ from $t = t(n, \epsilon, \delta, |\mathbb{K}|)$ copies if, for all states $|\psi\rangle$ and $(\epsilon, \delta) \in (0, 1)$:

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} \left[\forall k \in \mathbb{K} \begin{cases} \mathcal{F}(k, s) = \text{Accept} & \text{if } |\psi\rangle = |\phi_k\rangle \\ \mathcal{F}(k, s) = \text{Reject} & \text{if } |\langle \psi | \phi_k \rangle|^2 < 1 - \epsilon \end{cases} \right] > 1 - \delta. \quad (1.1)$$

We say that $(\mathcal{M}, \mathcal{F})$ is:

1. *Copy efficient* if $t = \text{poly}(n, \epsilon^{-1}, \log \delta^{-1}, \log |\mathbb{K}|)$ is sufficient.
2. *Computationally efficient* if it is copy efficient and the function $\mathcal{F}$ is computationally efficient.
3. η -*simulable* if there exists a classical PPT algorithm $\mathcal{M}_C(k, 1^n, 1^t) \rightarrow s \in \{0, 1\}^{g(t,n)}$ such that

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq \eta(|k|, n, t) \quad (1.2)$$

for all $k \in \mathbb{K}$ and $n, t \in \mathbb{N}$, where again $\mathcal{M}_C(k, 1^n, 1^t)$ and $\mathcal{M}(|\phi_k\rangle^{\otimes t})$ are understood as distributions over $\{0, 1\}^{g(t,n)}$.

Informally, $(\mathcal{M}, \mathcal{F})$ is a "measure first, ask later" state certification protocol for the set of states $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$ if, when $s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})$, the quantity $\mathcal{F}(k, s)$ allows one to decide whether $|\psi\rangle = |\phi_k\rangle$, for all $k \in \mathbb{K}$ and all states $|\psi\rangle$. We note that by exploiting existing classical shadow protocols for the set of observables $\{|\phi_k\rangle\langle\phi_k|\}$, one can immediately construct a *copy efficient* "measure first, ask later" state certification protocol for *any* set of states $\{|\phi_k\rangle\}$ [HKP20]. However, we stress that this standard classical shadow based state certification protocol will typically be neither computationally efficient nor η -simulable for meaningful values of η . Indeed, the idea behind the definition above is to provide an abstraction of such "classical-shadow-type" state certification protocols, and to make clear the properties beyond copy-efficiency which, as we show below, will be useful for cryptographic constructions.

With this in hand, we can then define variants of *certifiable* OWSGs in a natural way:

Definition 1.4: (Certifiable one-way state generators (informal version of Definition 4.1)) Given a one-way state generator $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ (defined formally in Definition 2.10) with output states $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$, and a "measure first, ask later" state certification protocol $(\mathcal{M}, \mathcal{F})$, we say that the tuple of one-way state generator and state certification protocol $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ is:

1. A *certifiable* one-way state generator if $(\mathcal{M}, \mathcal{F})$ is copy efficient.
2. An η -*simulable certifiable* one-way state generator if $(\mathcal{M}, \mathcal{F})$ is certifiable and η -simulable.

3. An *efficiently certifiable* one-way state generator if $(\mathcal{M}, \mathcal{F})$ is computationally efficient.
4. An η -*simulable, efficiently certifiable* one-way state generator if $(\mathcal{M}, \mathcal{F})$ is efficiently certifiable and η -simulable.

We note that one can easily give an analogous definition for certifiable PRSG (see Definition 4.1), and prove that standard output length certifiable PRSGs imply certifiable OWSGs (see Theorem 4.4). Additionally, we stress again that *any* OWSG is a certifiable OWSG when equipped with the “measure first, ask later” state certification protocol provided by standard classical shadows [HKP20]. However, this OWSG will typically fail to be either computationally efficient or η -simulable for meaningful values of η . With this established, we then show that certifiable OWSGs can be used to construct variants of OWPs. Specifically, we prove the following:

Theorem 1.5: (One way puzzles from certifiable one-way state generators)

1. Given a *certifiable* one-way state generator, one can construct a one-way puzzle (Theorem 5.1).
2. Given an *efficiently certifiable* one-way state generator, one can construct an *efficiently verifiable* one-way puzzle (Corollary 5.3).

The constructions used to prove Theorem 1.5 generalize and abstract Khurana and Tomer’s construction of OWPs via OWSGs and classical shadows [KT24]. Indeed, our contribution is to show that this prior construction works for *any* copy-efficient “measure first, ask later” state certification protocol—not just the one derived from classical shadows—and that, if this protocol is in fact *computationally efficient*, then the resulting OWP is efficiently verifiable. This is interesting because it allows us to obtain *variants* of OWPs by using state certification protocols *tailored* to the set of states defining the OWSG. Indeed, our hope is that this can be used as a tool for translating progress in state certification into progress on the proposal of explicit candidates for *efficiently verifiable* OWPs – i.e., for the instantiation of Quantumania. As discussed in Section 1.2 below, there has recently been significant progress in the development of state certification protocols [HPS25; GHO25; LZ26], and Theorem 1.5 shows that if such progress yields computationally efficient “measure first, ask later” state certification protocols for the output states of existing candidate OWSGs (or PRSGs), then one in fact obtains a candidate *efficiently verifiable* OWP.

There is however an important caveat! As we show in Theorem 1.6 below, if the certifiable OWSG has a classical key generation algorithm, and is also 1/3-simulable then the resulting OWP can be used to construct a (quantum-secure) OWF. As such, the OWP is by definition not in Microcrypt!

Theorem 1.6: (OWFs from 1/3-simulable certifiable OWSGs (informal version of Theorem 6.1)) *Given a $\frac{1}{3}$ -simulable certifiable one-way state generator, with a classical key generation algorithm, one can construct a quantum-secure one-way function.*

We prove Theorem 1.6 by showing that 1/3-simulability of the certifiable OWSG implies the existence of an efficient approximate *classical* algorithm for simulating the *quantum* key/puzzle sampling algorithm of the OWP resulting from the construction used to prove Theorem 1.5. Given this, it follows from existing results that

1. One can construct a quantum-secure *distributional* OWF from the OWP [KT25].
2. One can then construct a quantum-secure *weak* OWF from the distributional OWF [IL89; KK07], from which one can construct a OWF [Gol+01; RS19].

Taken together, we see that Theorem 1.5 and Theorem 1.6 provide a new opportunity for the construction of EV-OWPs which are genuinely in Microcrypt, via any OWSG which:

1. Is itself not constructed from OWFs (i.e., relies on a purely quantum assumption which is plausibly independent of OWFs).
2. Admits a computationally efficient “measure first, ask later” state certification protocol, which is *not* also 1/3-simulable.

In Section 8 we discuss a variety of candidate OWSGs and the extent to which they may or may not satisfy the criteria above. However, as already mentioned in Section 1.1.1, we also stress that when a OWSG fails to provide a OWP that is genuinely in Microcrypt by virtue of being simulable, this implies that the quantum assumptions used for the OWSG provide new assumptions for *classical* cryptography, by virtue of the fact that the OWSG can be used to construct OWFs. As mentioned before, the obstacle for instantiating Microcrypt becomes an opportunity for instantiating Minicrypt from novel inherently quantum assumptions.

1.1.3 Towards simpler one-way functions from the Search HPS assumption

Given Theorem 1.6, it is clear that to prove Theorem 1.2 it is sufficient to provide a "measure first, ask later" state certification protocol for Hamiltonian phase states, which is both copy efficient *and* 1/3-simulable. However, as we discussed in the previous section, this allows one to directly construct a distributional OWF, which can be compiled into a weak OWF, which can then be compiled into a OWF, with each step adding complexity to the construction of the OWF. With the goal of providing a simpler concrete OWF construction from the Search HPS assumption, we show that given an η -simulable *efficiently* certifiable OWSG, for an η which is negligible in a specific sense, then one can directly construct a simple quantum-secure OWF, without going via distributional OWFs. In other words, we leverage both the increased accuracy of the measurement simulation algorithm, and the computational efficiency of the post-processing functions, to obtain a simpler OWF.

Theorem 1.7: (Simple OWFs from simulable and efficiently certifiable OWSGs (informal version of Theorem 6.4))
Given an η -simulable efficiently certifiable one-way state generator with classical key generation algorithm, one can directly construct a quantum-secure one-way function, without going via a distributional one-way function, whenever η is such that $\eta(|k|, n, t)$ is negligible with respect to n , for all $|k|$ and t that are at most polynomial in n .

With the above in mind, to prove Theorem 1.2 we don't simply prove the sufficient statement that there exists a copy efficient *and* 1/3-simulable state certification protocol for Hamiltonian phase states. Instead, we prove the stronger statement that the existing state certification protocol from Ref. [HPS25] is in fact an η -simulable computationally efficient state certification protocol, for negligible η in the appropriate sense, which allows us to invoke Theorem 1.7 to construct a OWF from the Search HPS assumption.

1.2 Related work

Below we describe a variety of existing and active research directions, which intersect with the goals and contributions of this work.

Microcrypt: Our work contributes to the rapidly growing literature on understanding and characterizing Microcrypt. As this area has become too extensive to survey completely here, we refer the reader to the *Microcrypt Zoo* [Sat24] for a broad overview. Nevertheless, we highlight the following points to place our work in context:

1. Over the past years, a wide variety of inherently quantum cryptographic primitives – such as EFI pairs [BCQ23], variants of PRSGs [JLS18; BS20], pseudorandom unitaries [MH25], OWSGs [MY22], and (efficiently verifiable) OWPs [KT24; CGG24; KNY23] amongst others – have been proposed and studied. Our work introduces two new primitives, namely certifiable and efficiently certifiable OWSGs, and shows their utility for the construction of *variants* of OWPs via state certification protocols tailored to the underlying OWSG. Apart from allowing us to prove Theorem 1.2, this is particularly interesting as it provides a new route for the construction of *efficiently verifiable* OWPs, the minimal primitive of Quantumania [GMM+25]. Previous work has shown that EV-OWPs can be constructed from quantum pseudorandom generators [CGG24] (which can be constructed from *logarithmic output length* PRSGs [ALY23], which in turn can be constructed from quantum computable pseudorandom functions [KQT25]) as well as from a variety of (non-interactive) QCCC primitives [CGG24], and our work adds a new method to this toolbox.
2. In order to actually implement any cryptographic protocol, one requires a *concrete instantiation* of the primitive on which the protocol is built. To this end, there has been a wide variety of work aimed at providing candidate

concrete instantiations of Microcrypt primitives, from a wide variety of alternative assumptions [MPS+24; BS19; CDX+24; BHH+25; FGS+25]. However, we note that many of these constructions assume the existence of (quantum-secure) OWFs as a starting point! As Microcrypt is particularly interesting due to its potential existence even if OWFs do not exist, one ideally wants candidate constructions from assumptions which are potentially independent of OWFs. Prior to our work, the Hamiltonian phase state assumptions [BHH+25] were precisely such assumptions, and one of our main contributions is to show that these assumptions are in fact *not* independent of the existence of OWFs, as originally hoped.

Cryptography from (quantum) hardness of learning: There is a long line of work on both learning theoretic characterizations of cryptographic primitives, as well as concrete instantiations of cryptographic protocols and primitives from *hardness of learning* assumptions – such as learning parities with noise, and learning with errors [Reg09]. Inspired by this, and motivated by the growing number of inherently quantum cryptographic protocols and primitives, recent work has attempted to understand the extent to which inherently *quantum* learning theoretic assumptions can be used to both characterize and instantiate (quantum) cryptography. Indeed, the Hamiltonian phase state assumptions [BHH+25], are precisely an assumption on the hardness of learning Hamiltonian phase states (or equivalently, random IQP circuits) in a specific sense, and one of our primary contributions is to show that while these assumptions *cannot* be used to instantiate genuine Microcrypt primitives, they are sufficient for the construction of OWFs, and therefore provide a means for the concrete instantiation of Minicrypt (at least).

However, a wide variety of other quantum learning theoretic assumptions have also been recently proposed and explored in the context of quantum cryptography. Similar in spirit to the HPS assumptions are the *computational no-learning* and *computational no-cloning* assumptions introduced in Ref. [FGS+25], which posit the hardness of learning and cloning the output states of sufficiently deep *brickwork* random quantum circuits from copies of the output state of the circuit. Specifically, the authors of Ref. [FGS+25] have shown that this assumption is sufficient for the construction of OWSGs, quantum commitments and digital signatures. We note that these assumptions differ from the HPS assumptions central to this work in both the families of circuits considered, and the details of the learning task which is assumed to be hard. While not strictly learning theoretic, we also mention that Ref. [KT25] has also explored the cryptographic potential of random quantum circuits, by showing that the standard “quantum advantage assumptions” – namely, the $\#P$ -hardness of estimating the output probabilities of a random quantum circuit from its classical description [HE23] – is sufficient for the construction of OWPs.

Moving away from random quantum *circuits*, Refs. [HH24; CL26] have given characterizations of both EFI pairs and OWSGs in terms of the average-case complexity of learning (efficiently generatable) random quantum *states*. Additionally, Ref. [HHM25] has provided a characterization of OWP in terms of the average-case hardness of proper quantum *distribution* learning. Ref. [ADD+21] has also studied the possibility of building quantum cryptography from *hardware* assumptions via quantum Physically Unclonable Functions. Finally, we mention the recent series of works [PQS25; LPQ+26; KLP+25] that have proposed and studied the *learning stabilizers with noise* assumption – which at a high level posits the average-case hardness of decoding random quantum stabilizer codes and generalizes the well studied learning parities with noise assumption – and shown that this assumption is sufficient for instantiating Cryptomania – i.e., this inherently quantum assumption is useful for *classical* cryptography. Indeed, our proof that the HPS assumptions imply OWFs, place the HPS assumptions on a similar footing to that of the learning stabilizers with noise assumptions.

State certification: State certification protocols are central to the results and contributions of this work. Indeed, our hope is that by using the notion of certifiable OWSGs as a tool, the community can more easily translate progress on quantum state certification into concrete instantiations of Microcrypt primitives. To aid with this, it’s helpful to put (significant) recent progress on state certification [HPS25; GHO25; CLS+26; CLS26] into context. In particular, the following is a list of criteria with respect to which different state certification protocols are often compared and contrasted (adapted and extended from the comparative table in Figure 1 of Ref. [CLS26]), together with a brief discussion of the relevance of a given criterion for our applications here:

1. **Set of states:** The set of states for which the state certification protocol is guaranteed to work for. In our setting, we require the protocol to work for the output states of the candidate OWSG we would like to use as input to

our constructions.

2. Target dependence of measurements: Whether or not the protocol only requires *target-independent* measurements and target-state dependent classical post-processing, as opposed to target-dependent measurement protocols. In the latter case, one can also distinguish between adaptive and non-adaptive measurement strategies. Borrowing from Ref. [EFH+23] we have called the former protocols "measure first, ask later". Here, we require "measure first, ask later" state certification protocols.
3. Copy complexity: The number of measurements required in the worst case. We only require polynomial copy complexity.
4. Measurement complexity: The complexity of the largest measurement required. Again, we only require efficient measurements.
5. Oracle model: The type of oracle access to the target state which is assumed. For example, many protocols assume the ability to query specific amplitudes of the target state in the computational and/or Hadamard basis. For the construction of *certifiable* OWSGs (and therefore OWP), the oracle model is not relevant to us. However, for the construction of *efficiently certifiable* OWSGs (and therefore *efficiently verifiable* OWPs) we require that the oracle can be *efficiently simulated* classically, for all target states of interest.
6. Computational complexity: Combined time complexity of the entire protocol, including any classical post-processing of measurement outcomes, assuming oracle access to the target state. As above, for the construction of certifiable OWSGs this is not relevant for us. However, for the construction of efficiently certifiable OWSGs, we require protocols which are computationally efficient *including* the simulation of the oracle model.
7. Robustness: The extent to which the state certification protocol can be made *tolerant*. For our purposes even *non-tolerant* state certification protocols suffice.

With the above in mind, we note that the paradigmatic "measure first, ask later" state certification protocol is that of classical shadows and its variants [HKP20; HW23; GPS24; KG22; ZRM21; WHL+23; BHH+24; CEF26; BDL+24]. Indeed, global Clifford shadows provide a state certification protocol which works for *all* target states, and requires only polynomially many efficient measurements [HKP20] – which is precisely why Khurana and Tomer are able to use such classical shadows to construct OWPs from *any* OWSG [KT24]. Unfortunately however, global Clifford shadows are *not* computationally efficient for arbitrary target states, and therefore cannot be used out of the box to obtain *efficiently verifiable* OWPs from arbitrary OWSGs. While variants of classical shadows can be made computationally efficient for some sets of structured states (such as stabilizer states [HKP20], fermionic Gaussian states [ZRM21; WHL+23] and bosonic Gaussian states [BDL+24]) it is not clear whether such states can be used for the construction of OWSGs.

With the view to identifying state certification protocols which satisfy *all* the criteria for the construction of an *efficiently certifiable* OWSG (especially for the HPS derived OWSG we are particularly concerned with) we note that a significant amount of recent work [HPS25; GHO25; CLS+26; CLS26] has gone into developing novel state certification protocols, which try to simultaneously optimize as many as the above criteria as possible (we refer to Ref. [CLS26] for a detailed comparison). These protocols are also computational efficient for any set of target states for which the assumed oracle model can be efficiently simulated. As a result, at least for the OWSG we are primarily concerned with here, which has Hamiltonian phase states as output states, the only protocol which is both "measure first, ask later" *and* admits computationally efficient post-processing is that of Ref. [HPS25] (as a result of the fact that the oracle which is assumed can be simulated efficiently for phase states). While the measurement complexity, copy complexity and robustness of this protocol is not as good as the more recent state certification protocols developed in [CLS+26; CLS26] this is not a concern for us, given that we do not require any robustness, and polynomial copy and measurement complexity suffices. Finally, we note that a state certification protocol which satisfies all our desired criterion for phase states is also developed implicitly in Ref. [PCK+26], however for simplicity of presentation we focus on the protocol from Ref. [HPS25].

1.3 Discussion and open questions

We highlight the following open questions and directions arising from our contributions here:

Concrete instantiations of efficiently certifiable OWSGs: One of the primary contributions of this work is to provide a concrete set of sufficient conditions for a state ensemble, in terms of its learnability and certifiability, for the construction of efficiently certifiable OWSGs, and therefore for efficiently verifiable OWPs. Given this, perhaps the most natural open question is whether one can identify explicit ensembles of quantum states which are (a) pseudorandom under a plausible conjecture, and (b) admit a computationally efficient “measure first, ask later” state certification protocol which is *not* also simulable. Of course, we are particularly interested in state ensembles which are plausibly pseudorandom *even if* OWFs do not exist – i.e., state ensembles which are not constructed assuming the existence of a OWF, and whose pseudorandomness would not imply OWFs. Before this work, Hamiltonian phase states were precisely such a candidate ensemble, and our hope is that by understanding precisely the requirements in terms of certifiability, new candidate ensembles for the construction of EV-OWPs can be identified. We discuss a variety of existing candidates and their shortcomings in Section 8.

QCCC cryptography via efficiently certifiable OWSGs: Quantumia is particularly interesting because it contains a wide variety of QCCC cryptographic protocols and primitives [CGG24; KQT25]. With this in mind, are there QCCC cryptographic protocols which can be constructed directly from efficiently certifiable OWSGs? Could one construct interactive QCCC bit commitments [ALY23], or QCCC public-key encryption, directly from efficiently certifiable OWSGs? Said simply, can we show that the primitives that we introduce in this work are *useful*?

Separations and implications between Quantumia primitives: Refs. [BMM+25; BCN24] proved the existence of an oracle world in which OWPs exist, but OWSGs do not exist, showing that unlike classical Minicrypt, the world of Countcrypt *cannot* be collapsed to a single minimal primitive [CGG24]. Does Quantumia also have “sub-worlds”? More specifically, can one prove an analogous black box separation between *efficiently certifiable* OWSGs and *efficiently verifiable* OWPs? More generally, how are efficiently certifiable OWSGs related to other Quantumia primitives such as quantum computable OWFs [KQT25], logarithmic depth output PRSGs, and quantum pseudorandom generators [ALY23]?

Refinements of certifiable primitives: One can straightforwardly adapt our definition of a “measure first, ask later” state certification protocol (Definition 1.3) to define constrained notions such as a single-copy or non-adaptive “measure first, ask later” state certification protocols. This subsequently allows us to define single-copy and non-adaptive versions of certifiable PRSGs and OWSGs. Intuitively, since single-copy and non-adaptive state certification protocols for a given set of states are *harder* to construct than their multi-copy and adaptive counterparts, it may be the case that single-copy and non-adaptive versions of the efficiently certifiable primitives we have introduced here are *more powerful* than the unconstrained versions which allow multi-copy and adaptive measurement protocols for certification. Said another way, it is by now well established that in the context of learning and testing, multi-copy and adaptive measurements provide a powerful resource which enables exponential separations (see e.g., Ref. [CCH+22]). Do such separations manifest cryptographically? Can one construct cryptographic protocols using single-copy/non-adaptive certifiable primitives that one cannot using the unconstrained certifiable primitives we define here?

Evidence for efficiently-certifiable OWSGs without OWFs: The major motivation for studying many Microcrypt primitives comes from evidence that they can exist, even if $P = NP$ and OWFs do not exist [Kre21; KQS+23; KQT25]. Can we provide similar evidence that efficiently certifiable OWSGs can exist even if OWFs do not exist? From Ref. [KQT25] we know that there exists an oracle world in which $P = NP$ but quantum computable OWFs, and therefore also logarithmic output length PRSGs, quantum pseudorandom generators and EV-OWPs exist. As such, one may be able to answer this question by better understanding the implications between existing Quantumia primitives, as per the previous question.

Classical cryptography from quantum assumptions: Given the fact that the HPS assumptions are sufficient for the construction of an explicit OWF (Theorem 1.2), one can instantiate all cryptography in *Minicrypt* under the HPS assumption. An immediate natural question is whether one could in fact instantiate cryptographic primitives and protocols in *Cryptomania* under the HPS assumptions? To this end, an immediate direction would be to try to

construct suitable *trapdoor* OWFs from the HPS assumptions. Additionally, our results show that one can construct OWFs from any OWSG which admits a simulable and computationally efficient state certification protocol. Could one construct other OWSGs under inherently quantum assumptions, which also admit such state certification protocols, and can therefore be used to construct OWFs under inherently quantum assumptions?

Independence and relation of HPS assumptions from classical assumptions: As has been mentioned before, one of the primary reasons that Theorem 1.2 is interesting, is because it shows that the HPS assumptions provide a novel *inherently quantum* assumption for the construction of classical cryptography. However, unlike existing and well studied assumptions like "Learning with Errors" [Reg09] the HPS assumptions have not been tested in any way, and their relation to existing assumptions is largely unclear. With this in mind, if one is to take the HPS assumptions seriously as a foundation for classical cryptography, then significantly more effort is needed to investigate both the plausibility of these assumptions, and their potential relation to existing known assumptions.

Quantumania via state certification protocols with quantum post-processing: We note that our definition of a computationally efficient "measure first, ask later" state certification protocol implicitly requires the post-processing functions to be classically efficient to compute. This is done to ensure that the constructions using such protocols lead to EV-OWPs as normally understood – i.e. OWPs with classically efficient verification algorithms. However, if one was to allow for efficient deterministic *quantum* verification algorithms in the definition of an EV-OWP, then in principle any copy-efficient state certification protocol with efficient *quantum* postprocessing algorithms would suffice. We believe it is an interesting direction to understand the power of such "measure first, ask later" state certification protocols with *quantum* post-processing functions, and their relation to quantum computable OWFs and the instantiation of concrete QCCC cryptographic protocols like one-time digital signatures or interactive quantum bit commitments.

1.4 A short story on the highs and lows of developing this work

This work grew out of the observation that one could construct *efficiently verifiable* OWPs from OWSGs, if one had a computationally efficient "measure first, ask later" state certification protocol for the output states of the OWSG. With this in mind, we went looking for such a state certification protocol for Hamiltonian phase states, knowing that they were conjectured to provide a OWSG (plausibly independent of OWFs), with the hope of providing a concrete candidate instantiation of an EV-OWP, and therefore of a QCCC one-time digital signature, which was plausibly genuinely in Microcrypt. In particular, a recent experimental work implemented an "almost" QCCC digital signature [NLO+26], which was just lacking efficient verifiability, by virtue of using standard classical shadows. We know that if we could find a computationally efficient "measure first, ask later" state certification protocol tailored for Hamiltonian phase states, this would give us an explicit EV-OWP (under the HPS assumptions), from which we would be able to construct a genuine and experimentally feasible QCCC one-time digital signature, answering an open question from Ref. [FGS+25].

We then found the state certification protocol for (Hamiltonian) phase states from Ref. [PCK+26] and happily put together the pieces, writing a draft whose main claim was the construction of an explicit candidate for an EV-OWP, and therefore experimentally feasible QCCC one-time digital signature, which was plausibly independent of OWFs, under the HPS assumptions. In other words, a plausible concrete and explicit instantiation of Quantumania!

After uploading the manuscript to the arXiv, and sharing the completed draft among friends and colleagues, Matthias Caro asked us why we hadn't used the state certification protocol from Ref. [HPS25]? Quickly we realized that one *could* of course also use that state certification protocol, and that if we did this, the resulting OWP sampling algorithm would be *classically* simulable, and imply OWFs. In essence, we *did* have a construction of an efficiently verifiable OWP under the HPS assumptions, but one could construct a OWF from this OWP, and therefore *the HPS assumptions implied OWFs*. Nothing we had written in our original draft was wrong, except the foundational assumption! We pulled the paper from the arXiv before it was announced (luckily), and the result is the work you are now reading.

1.5 Structure of this work

We start in Section 2 by introducing all relevant notation, existing definitions and assumptions. Following this, we provide in Section 3 formal definitions for the "measure first, ask later" state certification protocols which are central to this work. Building on this, we then define a variety of *certifiable* Microcrypt primitives in Section 4. With this established we then prove in Section 5 that one can build variants of OWPs from certifiable OWSGs. We then show in Section 6 that if a OWSG is *simulable* in addition to being certifiable, then one can construct variants of OWFs from the OWPs constructed from the OWSG. Given this, we can then finally prove Theorem 1.2 in Section 7. We conclude in Section 8 with a discussion of candidate state ensembles for instantiating Quantumania via the tools developed in this work.

AI Disclosure

The first version of this manuscript, containing all the essential ideas, constructions, proofs and method of presentation, was obtained without the assistance of generative AI. Both Claude Opus 4.8 (Max) and ChatGPT 5.6 Sol (Ultra) were then used to generate a detailed and critical referee reports, with a focus on verifying mathematical correctness. These report identified a variety of subtle mathematical issues, which were fixed via interaction with both models. This materially affected details in the proofs of Corollary 2.8, Theorem 5.1, Theorem 6.1 (as it relies on Corollary 2.8), Theorem 6.4 and Lemma 7.6. The authors verified the correctness and originality of all content including references.

Acknowledgements

We are grateful for helpful and inspiring conversations with Carlos Cid, Janek Denzler, David Elkouss, Bill Fefferman, Elies Gil-Fuster, Tommaso Guaita, Manuel Goulão, Jonas Haferkamp, Zephрина Aniska Be, Mina Doosti and Martina Onetti. JC thanks the support of Berlin Quantum. NKHL acknowledges support from the Austrian Science Fund (FWF) (10.55776/P36478), the Austrian Federal Ministry of Education, Science and Research via the Austrian Research Promotion Agency (FFG) through the flagship project FO999921415 (Vanessa-QC) funded by the European Union—NextGenerationEU, the European Research Council (Consolidator grant ‘Cocoquest’ 101043705), and the Croucher Foundation. RS thanks the Alexander von Humboldt foundation for their support, under the German Research Chair program at the African Institutes for Mathematical Sciences. JE has been funded by the BMFTR (QuSol, Hybrid++), Berlin Quantum, the Munich Quantum Valley, the Quantum Flagship (Millenion, PasQuans2), the DFG (CRC 183, SPP 2514), the Clusters of Excellence (ML4Q, MATH+), and the European Research Council (DebuQC). Part of this work was carried out while JC and JE visited the African Institute for Mathematical Sciences (AIMS), Cape Town, for the “1st AIMS Workshop on the Theory of Quantum Learning Algorithms” (2025). DH was supported by a Simons postdoctoral fellowship through DOE QSA and NSF QLCI Grant No. 2016245, and by the Swiss National Science Foundation through Ambizione Grant No. 223764.

2 Preliminaries

Throughout this work we use the notation $x \leftarrow \mathcal{X}$ to denote x sampled uniformly from the set $\mathcal{X}$, and the notation $x \leftarrow D$ to denote x sampled from a distribution D . Given some function $f : \mathcal{X} \rightarrow \mathcal{Y}$ we will use the notation $(f(x))_{x \leftarrow \mathcal{X}}$ to denote the distribution over $\mathcal{Y}$ which one samples from by first sampling $x \leftarrow \mathcal{X}$ and then outputting $f(x)$. Additionally, we use the notation negl to represent a negligible function – i.e. any function $f : \mathbb{N} \rightarrow \mathbb{R}$ such that for every positive polynomial p there exists an N such that for all $n \geq N$ one has $f(n) < 1/p(n)$. We use the abbreviations QPT and PPT for “quantum polynomial time” and “probabilistic polynomial time” respectively.

2.1 Classical cryptographic primitives

We begin by defining the classical cryptographic primitives relevant to this work. The first such primitives are (quantum-secure) one-way functions.

Definition 2.1: (One-way function (OWF)) Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be some fixed polynomial. A family of efficiently computable functions $\{F_\lambda : \{0, 1\}^{m(\lambda)} \rightarrow \{0, 1\}^*\}_{\lambda \in \mathbb{N}}$ is:

1. A (quantum-secure) *one-way function* if for all PPT (QPT) algorithms $\mathcal{A}$ there exists a negligible function negl such that

$$\Pr_{\substack{x \leftarrow \{0,1\}^{m(\lambda)} \\ y \leftarrow F_\lambda(x)}} [\mathcal{A}(1^\lambda, y) \in F_\lambda^{-1}(y)] \leq \text{negl}(\lambda) \quad (2.1)$$

for all sufficiently large $\lambda \in \mathbb{N}$.

2. A (quantum-secure) *weak one-way function* if there exists a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that for all (QPT) PPT algorithms $\mathcal{A}$

$$\Pr_{\substack{x \leftarrow \{0,1\}^{m(\lambda)} \\ y \leftarrow F_\lambda(x)}} [\mathcal{A}(1^\lambda, y) \notin F_\lambda^{-1}(y)] > \frac{1}{p(\lambda)} \quad (2.2)$$

for all sufficiently large $\lambda \in \mathbb{N}$.

In order to break a candidate OWF $\{F_\lambda\}$, on input $F_\lambda(x)$ the adversary needs to output a *single* element of the preimage of $F_\lambda(x)$. Given a weak OWF, there exists a polynomial p such that for all adversaries the failure probability is at least $1/p$ with respect to x drawn uniformly randomly. Given a OWF, the success probability of any adversary is negligible. We note that weak OWFs imply OWFs [Gol+01] and that quantum-secure weak OWFs imply quantum-secure OWFs [KK07; RS19].

The second classical cryptographic primitive relevant to this work is a *distributional* one-way function, for which, on input $F_\lambda(x)$, the adversary has the harder task of sampling from the uniform distribution over the preimage of $F_\lambda(x)$. A function family is a distributional one-way function if there is an inverse polynomial lower bound on the accuracy achievable by any PPT adversary:

Definition 2.2: (Distributional one-way function (D-OWF) [IL89]) Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be some fixed polynomial. A family of efficiently computable functions $\{F_\lambda : \{0,1\}^{m(\lambda)} \rightarrow \{0,1\}^*\}_{\lambda \in \mathbb{N}}$ is a *distributional one-way function* if there exists a polynomial p such that for all PPT algorithms $\mathcal{A}$, one has

$$d_{\text{TV}}((x, F_\lambda(x))_{x \leftarrow \{0,1\}^{m(\lambda)}}, (\mathcal{A}(1^\lambda, F_\lambda(x)), F_\lambda(x))_{x \leftarrow \{0,1\}^{m(\lambda)}}) \geq \frac{1}{p(\lambda)} \quad (2.3)$$

all sufficiently large $\lambda \in \mathbb{N}$.

In the above definition, only classical PPT adversaries have been considered. To define a *quantum-secure* distributional one-way function, one insists that, on average with respect to function inputs, no QPT adversary can prepare a state close to a uniform superposition over preimages. To make this precise, given any two quantum states ρ, σ let's denote by $F_{\text{sq}}(\rho, \sigma) = \|\sqrt{\rho}\sqrt{\sigma}\|_1^2$ the squared fidelity between the two states. We then have:

Definition 2.3: (Quantum-secure distributional one-way function (Adapted from [KK07])) Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be some fixed polynomial. A family of efficiently computable functions $\{f_\lambda : \{0,1\}^{m(\lambda)} \rightarrow \{0,1\}^*\}_{\lambda \in \mathbb{N}}$ is a quantum-secure *distributional one-way function* if there exists a polynomial p such that for all QPT algorithms $\mathcal{A}$ and all sufficiently large λ one has that

$$\bar{F}_{\mathcal{A}}(\lambda) := \frac{1}{2^{m(\lambda)}} \sum_{r \in \{0,1\}^{m(\lambda)}} F_{\text{sq}}(\sigma_{f_\lambda(r)}, \rho_{f_\lambda(r)}^{\mathcal{A}}) \leq 1 - \frac{1}{p(\lambda)} \quad (2.4)$$

where $\sigma_y = |H_y\rangle\langle H_y|$ and $\rho_y^{\mathcal{A}} = \mathcal{A}(1^\lambda, |y\rangle\langle y|)$ with

$$|H_y\rangle = \frac{1}{\sqrt{|f_\lambda^{-1}(y)|}} \sum_{z \in f_\lambda^{-1}(y)} |z\rangle. \quad (2.5)$$

We note that the definition we have given above, where the adversary can output mixed quantum states, is a strengthening of the definition given in [KK07] and implies their definition. As sampling from the preimage of $F_\lambda(x)$ is a

harder task than outputting a single element of the preimage of $F_\lambda(x)$, any (quantum-secure) OWF is immediately also a (quantum-secure) distributional OWF. The other direction however does not hold. In particular, it is possible for a family of functions $\{F_\lambda\}$ to be a distributional OWF but *not* a OWF. Perhaps surprisingly, however, Impagliazzo and Luby showed that given a distributional OWF, one can construct a OWF [IL89]. Particularly, we have the following result:

Theorem 2.4: (D-OWFs imply OWFs (Lemma 1 [IL89]) and Theorem 4.2.2 [Imp92])) *Given a distributional OWF, one can construct a OWF.*

We note that the construction underneath Theorem 2.4 works by first constructing a *weak* OWF from the distributional OWF, then constructing a OWF from the weak OWF. An analogous result for *quantum-secure* distributional OWF was proven in Ref. [KK07]:

Theorem 2.5: (Quantum-secure D-OWFs imply quantum-secure OWFs (Theorem 5 [KK07])) *Given a quantum-secure distributional OWF, one can construct a quantum-secure OWF.*

We note that in Ref. [KK07] they actually only proved that given a quantum-secure distributional OWF one can construct a quantum-secure *weak* OWF, however as mentioned, quantum-secure weak OWFs imply quantum-secure OWFs [RS19].

2.2 Microcrypt primitives

In this section, we provide definitions for all existing Microcrypt primitives that are relevant for this work. We begin with the definition of a *one-way puzzle*, originally introduced by Khurana and Tomer in Ref. [KT24].

Definition 2.6: (One-way puzzle (OWP)) A *one-way puzzle* is a pair $(\text{Samp}, \text{Ver})$ consisting of a sampling algorithm Samp and a verification function Ver . Here, $\text{Samp}(1^\lambda) \rightarrow (k, s)$ is either a uniform quantum polynomial time (QPT) or probabilistic polynomial time (PPT) algorithm that on input of a security parameter λ outputs a pair of classical strings; we call k the *key* and s the *puzzle*. The function $\text{Ver}(k, s) = b$ takes a key and a puzzle as inputs and returns a bit $b \in \{0, 1\}$. When $\text{Ver}(k, s) = 1$, we say that the pair (k, s) is *accepted*. For all sufficiently large λ , these must satisfy the following properties:

1. *Correctness:* Outputs of the sampling algorithm are accepted with high probability,

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(k, s) = 1] \geq 1 - \text{negl}(\lambda). \quad (2.6)$$

2. *Security:* Given just a puzzle s , it is computationally hard to find a key k such that the pair (k, s) is accepted. That is, for all QPT algorithms $\mathcal{A}(1^\lambda, s) \rightarrow k'$ that take a puzzle as input and output a key one has

$$\Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^\lambda) \\ k' \leftarrow \mathcal{A}(1^\lambda, s)}} [\text{Ver}(k', s) = 1] \leq \text{negl}(\lambda). \quad (2.7)$$

We note that the definition above allows for classical PPT or quantum QPT sampling algorithms. However, when the sampling algorithm is classical, it is known that one can construct a distributional OWF from the OWP:

Theorem 2.7: (Classical sampling for one-way puzzles implies a distributional one-way function (Claim D.1 [KT25])) *Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be some fixed polynomial. Given a one-way puzzle $(\text{Samp}, \text{Ver})$, assume that there exists an efficient deterministic classical algorithm Samp_C such that*

$$d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq 1/3, \quad (2.8)$$

where $\text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}} = (k_\lambda(r), s_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}$ is understood as the distribution over outputs of $\text{Samp}_C(1^\lambda, r)$ with respect to input bitstrings $r \in \{0, 1\}^{m(\lambda)}$ drawn uniformly at random. Then, the family of functions $\{f_\lambda : \{0, 1\}^{m(\lambda)} \rightarrow \{0, 1\}^*\}_{\lambda \in \mathbb{N}}$ with $f_\lambda(r) = s_\lambda(r)$ is a distributional one-way function.

In light of Theorem 2.4, the above means that one can construct OWFs from OWPs with classical sampling algorithms. The result cited above only explicitly proved the ability to construct a distributional OWF, however, the proof can be generalized to obtain a quantum-secure distributional OWF. In particular, we have the following corollary:

Corollary 2.8: (Classical sampling for one-way puzzles implies quantum-secure distributional one-way function) *Under the same conditions as Theorem 2.7, the family of functions $\{f_\lambda\}_{\lambda \in \mathbb{N}}$ is in fact a quantum-secure distributional one-way function.*

For completeness, we provide a proof of Corollary 2.8 in Appendix A. In light of Theorem 2.5, the above means that one can construct quantum-secure OWFs from OWPs with classical sampling algorithms. Given these results, from the perspective of Microcrypt, we are most interested in OWPs genuinely *quantum* sampling algorithms. As discussed in the introduction, multiple recent works have given evidence that such *quantum* OWP (i.e., with QPT sampling algorithm Samp) can be constructed even if OWFs do not exist [KT25; KQT25]. Additionally, OWPs are the minimal assumption for *Countcrypt* [GMM+25].

While the standard definition of a OWP given in Definition 2.6 does not require the verification algorithm Ver to be efficient, in this work we will be particularly interested in OWP where the verification algorithm Ver is efficient. In this case, we obtain an *efficiently verifiable* OWP defined formally as follows.

Definition 2.9: (Efficiently verifiable one-way puzzle (EV-OWP)) An *efficiently verifiable* one-way puzzle is a one-way puzzle $(\text{Samp}, \text{Ver})$ as above, with the added requirement that the verification function Ver is efficiently computable.

Note that since $\text{Samp}(1^\lambda) \rightarrow (k, s)$ is efficient, one implicitly has that $(k, s) \in \{0, 1\}^{O(\text{poly}(\lambda))}$. Therefore, $\text{Ver}(k, s) = b$ is efficiently computable if the runtime is $O(\text{poly}(\lambda))$ – i.e., polynomial with respect to the security parameter. Additionally, as discussed at length in the introduction, the existence of EV-OWP is the minimal assumption for Quantumania [GMM+25].

With this, we proceed to define the notion of a *pure one-way state generator*, initially introduced by Morimae and Yamakawa [MY22].

Definition 2.10: (Pure one-way state generator (OWSG)) A *pure one-way state generator* is a triple $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ of QPT algorithms satisfying the following conditions. The sampling algorithm $\text{KeyGen}(1^\lambda) \rightarrow k$ takes as input a security parameter λ and outputs a classical key $k \in \mathbb{K}_\lambda := \text{supp}(\text{KeyGen}(\lambda))$. The state generator algorithm $\text{StateGen}(k) \rightarrow |\phi_k\rangle$ takes a key $k \in \mathbb{K}_\lambda$ as input and outputs an $n(\lambda)$ -qubit pure quantum state $|\phi_k\rangle$. The verification algorithm $\text{Ver}(k, |\phi\rangle) \rightarrow b$ takes as input a key and a pure quantum state $|\phi\rangle$ and returns a bit $b \in \{0, 1\}$ as the result of applying the projective measurement $\{\Pi_k^1, \Pi_k^0\} = \{|\phi_k\rangle\langle\phi_k|, \mathbb{I} - |\phi_k\rangle\langle\phi_k|\}$ on the state $|\phi\rangle$, i.e., $\Pr[\text{Ver}(k, |\phi\rangle) = b] = \langle\phi|\Pi_k^b|\phi\rangle$. For sufficiently large λ , these algorithms must satisfy the following properties:

1. *Correctness:* Outputs of the sampling algorithm $\text{KeyGen}(1^\lambda) \rightarrow k$ and the state generator algorithm $\text{StateGen}(k) \rightarrow |\phi_k\rangle$ on input k are such that the pair $(k, |\phi_k\rangle)$ is accepted with high probability,

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ |\phi_k\rangle \leftarrow \text{StateGen}(k)}} \left[\text{Ver}(k, |\phi_k\rangle) \rightarrow 1 \right] \geq 1 - \text{negl}(\lambda). \quad (2.9)$$

2. *Security:* Given just copies of the states $|\phi_k\rangle$, it is computationally hard to find a key k' such that the pair $(k', |\phi_k\rangle)$ passes verification. That is, for all QPT algorithms $\mathcal{A}(1^\lambda, |\phi\rangle^{\otimes t}) \rightarrow k'$ that, on input t copies of a state, outputs a key, one has

$$\Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ |\phi_k\rangle \leftarrow \text{StateGen}(k) \\ k' \leftarrow \mathcal{A}(1^\lambda, |\phi_k\rangle^{\otimes t})}} \left[\text{Ver}(k', |\phi_k\rangle) \rightarrow 1 \right] \leq \text{negl}(\lambda), \quad (2.10)$$

for all $t = O(\text{poly}(\lambda))$.

Before proceeding, a number of remarks concerning the above definition are in order.

1. In principle, as discussed in Ref. [MY24], one can generalize the definition above in two ways: (a) allowing the QPT algorithm $\text{StateGen}(k) \rightarrow \rho_k$ to output mixed states and (b) allowing arbitrary QPT algorithms $\text{Ver}(k, \rho) \rightarrow b$ where the bit b need not be the result of a rank-1 projective measurement.
2. However, following Ref. [MY24], we note that if all $\rho_k = |\phi_k\rangle\langle\phi_k|$ are pure, as we consider in the above definition, then one can always define $\text{Ver}(k, \rho) \rightarrow b$ as the result of applying the projective measurement $\{\Pi_k^1, \Pi_k^0\} = \{|\phi_k\rangle\langle\phi_k|, \mathbb{I} - |\phi_k\rangle\langle\phi_k|\}$ on the state ρ . Therefore, for OWSG with pure-state outputs, as we consider here, our definition does not lose generality.
3. As KeyGen is efficient we have $\mathbb{K}_\lambda \subseteq \{0, 1\}^{O(\text{poly}(\lambda))}$. Given this, it follows from the efficiency of StateGen that $n(\lambda) = O(\text{poly}(\lambda))$.

Given that we only consider OWSG with pure outputs in this work, we will from now on drop the explicit *pure* qualification. Additionally, we note that in their original work defining OWP [KT24], Khurana and Tomer showed that OWSG are sufficient for the construction of OWP, by utilizing classical shadows to obtain a classical puzzle from the output of a OWSG. One of the main contributions of this work, in Section 5, will be to generalize this construction to allow for arbitrary "measure first, ask later" state certification protocols in place of classical shadows.

Finally, the last Microcrypt primitive relevant to us is the *pseudorandom state generator*, originally defined by Ji, Liu and Song [JLS18].

Definition 2.11: (Pseudorandom state generator (PRSG)) A *pseudorandom state generator* is a pair of QPT algorithms $(\text{KeyGen}, \text{StateGen})$ consisting of a sampling algorithm KeyGen and a state generator algorithm StateGen . The sampling algorithm $\text{KeyGen}(1^\lambda) \rightarrow k$ takes as input a security parameter λ and outputs a classical key $k \in \mathbb{K}_\lambda := \text{supp}(\text{KeyGen}(\lambda))$. The state generator algorithm $\text{StateGen}(k) \rightarrow |\phi_k\rangle$ takes a key $k \in \mathbb{K}_\lambda$ as input and outputs an $n(\lambda)$ -qubit pure quantum state $|\phi_k\rangle$. For sufficiently large λ and for all QPT algorithms $\mathcal{A}(1^\lambda, |\phi\rangle^{\otimes t}) \rightarrow b$ that, on input t copies of a state, output a bit $b \in \{0, 1\}$, one has

$$\left| \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ |\phi_k\rangle \leftarrow \text{StateGen}(k)}} [\mathcal{A}(1^\lambda, |\phi_k\rangle^{\otimes t}) = 1] - \Pr_{|\psi\rangle \leftarrow \text{Haar}(n)} [\mathcal{A}(|\psi\rangle^{\otimes t}) = 1] \right| \leq \text{negl}(\lambda), \quad (2.11)$$

for all $t = O(\text{poly}(\lambda))$, where $\text{Haar}(n)$ is the Haar distribution over n -qubit states.

As for a OWSG, given that $(\text{KeyGen}, \text{StateGen})$ are both efficient in the above definition, one again implicitly has that $\mathbb{K}_\lambda \subseteq \{0, 1\}^{O(\text{poly}(\lambda))}$ and $n = O(\text{poly}(\lambda))$. As alluded to in the introduction, we will refer to the function $n(\lambda)$ as the *stretch* of the PRSG. We note that it is helpful to distinguish the following regimes:

1. $n(\lambda) = \omega(\log \lambda)$. This is the regime in which (a) Ji, Liu and Song were able to show that PRSGs can be built from (quantum-secure) OWF [JLS18], (b) PRSG immediately yield OWSG [CGG+25] which in turn yield OWP [KT25] and (c) in which Kretschmer gave a black box separation between PRSG and OWF [Kre21] (giving evidence that such PRSG can exist even if OWFs do not). We will call such PRSG *standard* PRSG.
2. $n(\lambda) = c \log \lambda$ with $(c \geq 1)$. These are called *short* or *logarithmic length* PRSGs. Given that one can perform tomography on a quantum state of $O(\log \lambda)$ qubits to any desired precision ϵ in time $\text{poly}(\lambda, \epsilon)$, short PRSG behave more like cryptographic objects with classical output. Using this insight, such PRSG can be used to build quantum pseudorandom generators (QPRG), which are *pseudodeterministic* generators of pseudorandom bit strings [ALY23], from which one can construct EV-OWP [CGG24]. Such short PRSG can also be constructed from (quantum-secure) OWFs [BS20], and as a result the recent black box separation between OWFs and *quantum computable* OWFs gives evidence that short PRSG can also exist in a world without OWFs [KQT25].
3. $n(\lambda) \leq c \log \lambda$ for $c \in (0, 1)$. Here, there exist c such that PRSGs can exist *unconditionally* [BS20; AGQ+22].

In this work, given that we are interested in novel constructions of EV-OWP from OWSGs, we focus on the regime of *standard* PRSG. Finally, we note that one can also define a weaker notion of a *single-copy secure* PRSG, in which

the condition of Definition 2.11 is only required to hold for $t = 1$. We will, however, be concerned here with standard multi-copy secure PRSG.

2.3 Hamiltonian phase states and Microcrypt

In this section, we introduce the notion of phase states, and the specific set of phase states known as *Hamiltonian phase states* (HPS). Doing this allows us to introduce the central assumptions in this work – namely, the search and decision HPS assumptions recently introduced in Ref. [BHH+25]. We start with the definition of phase states.

Definition 2.12: (Phase states) Let n be a positive integer and consider arbitrary functions $f : \{0, 1\}^n \rightarrow [0, 2\pi)$ that map classical strings to angles. Every such *phase function* f defines a *phase state* $|\phi_f\rangle = \sum_{\mathbf{x}} \alpha_f(\mathbf{x}) |\mathbf{x}\rangle$ on n qubits with $\alpha_f(\mathbf{x}) = 2^{-n/2} \exp(iff(\mathbf{x}))$. The class PS_n of n -qubit phase states is defined as

$$\text{PS}_n := \left\{ |\phi_f\rangle := \frac{1}{\sqrt{2^n}} \sum_{\mathbf{x} \in \mathbb{F}_2^n} \exp(iff(\mathbf{x})) |\mathbf{x}\rangle \right\}_{f: \{0,1\}^n \rightarrow [0,2\pi)}. \quad (2.12)$$

There are uncountably many such states. We will be interested in restricted families of states associated to sets $\{f_k\}_{k \in \mathbb{K}}$ of phase functions that one can at least label by the elements $k \in \mathbb{K}$ of some numerable set. Given a family of functions $\{f_k\}_{k \in \mathbb{K}}$ with $f_k : \{0, 1\}^n \rightarrow [0, 2\pi)$ for all $k \in \mathbb{K}$, we define its corresponding set of phase states $\text{PS}_n(\{f_k\}_{k \in \mathbb{K}})$ simply as

$$\text{PS}_n(\{f_k\}_{k \in \mathbb{K}}) := \left\{ |\phi_k\rangle := \frac{1}{\sqrt{2^n}} \sum_{\mathbf{x} \in \mathbb{F}_2^n} \exp(iff_k(\mathbf{x})) |\mathbf{x}\rangle : k \in \mathbb{K} \right\} \subset \text{PS}_n. \quad (2.13)$$

Given this, we will primarily be interested in sets of phase states defined by *efficiently computable* phase functions, defined as follows:

Definition 2.13: (Efficiently computable phase functions) Given a set of phase functions $\{f_k\}_{k \in \mathbb{K}}$ with $f_k : \{0, 1\}^n \rightarrow [0, 2\pi)$, we say that $\{f_k\}_{k \in \mathbb{K}}$ is *efficiently computable* if there exists a deterministic algorithm $\mathcal{A}$ which on input $k \in \mathbb{K}$, $x \in \{0, 1\}^n$ and 1^l with $l \in \mathbb{N}$, runs in time $\text{poly}(|k|, n, l)$ and outputs a dyadic rational $\mathcal{A}(k, x, l) = \hat{f}_k(x, l)$ satisfying

$$|\hat{f}_k(x, l) - f_k(x)| \leq \frac{1}{2^{l+1}}, \quad (2.14)$$

where $|k|$ denotes the bit-length of the encoding of index k .

With this established, we now define the specific set of phase states known as Hamiltonian phase states.

Definition 2.14: (Hamiltonian phase states) Let q, m, n be positive integers and $\Theta_q = \{0, 2\pi/q, 4\pi/q, \dots, 2(q-1)\pi/q\}$. The class $\text{HPS}_{q,m,n}$ of *Hamiltonian phase states* (HPS) contains n -qubit states of the form $|\phi(\theta, \mathbf{A})\rangle = U(\theta, \mathbf{A})|+\rangle^n$, where the unitary $U(\theta, \mathbf{A})$ is diagonal and given by

$$U(\theta, \mathbf{A}) = \exp \left(i \sum_{j=1}^m \theta_j Z^{A_{j,1}} \otimes Z^{A_{j,2}} \otimes \dots \otimes Z^{A_{j,n}} \right) = \sum_{\mathbf{x} \in \mathbb{F}_2^{mn}} \exp \left(i \sum_{j=1}^m \theta_j (-1)^{A_j \cdot \mathbf{x}} \right) |\mathbf{x}\rangle \langle \mathbf{x}| \quad (2.15)$$

for all $\theta = (\theta_1, \dots, \theta_m) \in \Theta_q^m$, and $\mathbf{A} = ((A_{1,1}, \dots, A_{1,n}), \dots, (A_{m,1}, \dots, A_{m,n})) \in \mathbb{F}_2^{mn}$. We define $\mathbb{K}_{q,m,n} := \Theta_q^m \times \mathbb{F}_2^{mn}$ for the set of labels and write $k = (\theta, \mathbf{A}) \in \mathbb{K}_{q,m,n}$ for its elements. Therefore,

$$\text{HPS}_{q,m,n} := \left\{ |\phi_k\rangle := \frac{1}{\sqrt{2^n}} \sum_{\mathbf{x} \in \mathbb{F}_2^n} \exp(iff_k(\mathbf{x})) |\mathbf{x}\rangle \in (\mathbb{C}^2)^{\otimes n} : k \in \mathbb{K}_{q,m,n} \right\} \quad (2.16)$$

where, for all $k = (\theta, \mathbf{A}) \in \mathbb{K}_{q,m,n}$, we have

$$f_k(\mathbf{x}) := \sum_{j=1}^m \theta_j (-1)^{A_j \cdot \mathbf{x}} = \sum_{j=1}^m \theta_j (-1)^{A_{j,1}x_1 \dots (-1)^{A_{j,n}x_n}. \quad (2.17)$$

Remark 2.15: (Preparing Hamiltonian phase states) As discussed in detail in Section 2.4 of Ref. [BHH+25], states in the class $\text{HPS}_{q,m,n}$ can be prepared via *instantaneous quantum polynomial-time* (IQP) circuits with $\text{poly}(n, m)$ gates acting on $|0^n\rangle$, independently of the value of q . Specifically, preparation of any state $|\phi_k\rangle \in \text{HPS}_{q,m,n}$ only requires a single layer of Hadamard gates, followed by $\lceil m/n \rceil$ alternating layers of single-qubit Z rotations and CNOT circuits. The commuting structure of these circuits makes Hamiltonian phase states natural candidates for implementation in programmable quantum platforms, including Rydberg-atom and superconducting architectures, as discussed in Appendix B.

Remark 2.16: (On parameter regimes for poly-size keys) For integer values of q , the class $\text{HPS}_{q,m,n}$ contains $|\mathbb{K}_{q,m,n}| = q^m 2^{nm}$ states that can be labeled by keys k with $|k| = m(n + \lceil \log q \rceil)$ bits. In general, we will take the number of qubits n as our security parameter for cryptographic applications and thus require that both m and $\log q$ are at most $\text{poly}(n)$, so that the class $\text{HPS}_{q,m,n}$ can be labeled by keys with $\text{poly}(n)$ bit-length encodings.

Remark 2.17: (Efficient computability of HPS phase functions) For any positive integers q, m, n , the set of phase functions $\{f_k\}_{k \in \mathbb{K}_{q,m,n}}$ defining the class $\text{HPS}_{q,m,n}$ in Eq. (2.17) is efficiently computable. More specifically, there exists an algorithm $\mathcal{A}$ which on input (k, x, l) runs in time $\text{poly}(n, m, \log q, l)$ and outputs $\mathcal{A}(k, x, l) = \hat{f}_k(x, l)$ satisfying Eq. (2.14). In the regime of Remark 2.16 – i.e. when both m and $\log q$ are at most $\text{poly}(n)$ – the runtime of $\mathcal{A}$ on input (k, x, l) is $\text{poly}(n, l)$.

We note that Hamiltonian phase states, and the IQP circuits which prepare them, were initially proposed as a class of states that is both easy to prepare and hard to simulate classically [SB09; BJS11]. However from our perspective, our primary interest in Hamiltonian phase states stems from the recently introduced assumptions, motivated by the apparent hardness of *learning* these states, that they can be used to construct both PRSG and OWSG. We start from the *Decision* HPS assumption:

Assumption 2.18: (Decision HPS assumption – Definition 2 in Ref. [BHH+25]) There exist functions $q(n) = 2^{O(\text{poly}(n))}$ and $m(n) = \text{poly}(n)$, and a classically efficiently sampleable distribution $\chi_{q,m,n}$ over $\mathbb{K}_{q,m,n}$, such that for all QPT algorithms $\mathcal{A}(|\phi\rangle^{\otimes t}) \rightarrow b$ that, on input t copies of an n -qubit state $|\psi\rangle$, output a bit $b \in \{0, 1\}$, one has

$$\left| \Pr_{k \leftarrow \chi_{q,m,n}} [\mathcal{A}(|\phi_k\rangle^{\otimes t}) = 1] - \Pr_{|\psi\rangle \leftarrow \text{Haar}(n)} [\mathcal{A}(|\psi\rangle^{\otimes t}) = 1] \right| \leq \text{negl}(n) \quad (2.18)$$

for all $t = \text{poly}(n)$, where $\text{Haar}(n)$ is the Haar distribution over n -qubit states.

Note that an immediate corollary of the above assumption is that, under the Decision HPS assumption, the following construction yields a PRSG.

Construction 2.19: (PRSG from Decision HPS) Given $m(n)$, $q(n)$ and a distribution $\chi_{q,m,n}$ which satisfy the Decision HPS assumption, define:

1. $\text{KeyGen}(1^n)$: Sample a key $k \leftarrow \chi_{q,m,n}$.
2. $\text{StateGen}(k \in \mathbb{K}_{q,m,n})$: Output an n -qubit state $|\phi_k\rangle \in \text{HPS}_{q,m,n}$ defined in Eqs. (2.16)-(2.17).

In particular, note that we have used n as the security parameter, and since $\log |\mathbb{K}_{q,m,n}| = m(n + \log q) = \text{poly}(n)$, KeyGen is efficient with respect to n .

With this established, we now introduce the *Search* HPS assumption.

Assumption 2.20: (Search HPS assumption – Definition 1 in Ref. [BHH+25]) There exist functions $q(n) = 2^{O(\text{poly}(n))}$ and $m(n) = \text{poly}(n)$, and a classically efficiently sampleable distribution $\chi_{q,m,n}$ over $\mathbb{K}_{q,m,n}$, such that for all QPT

algorithms $\mathcal{A}(|\phi\rangle^{\otimes t}) \rightarrow k'$ that, on input t copies of a n -qubit state, output a key $k' \in \mathbb{K}_{q,m,n}$, one has

$$\mathbb{E}_{\substack{k \leftarrow \chi_{q,m,n} \\ k' \leftarrow \mathcal{A}(|\phi_k\rangle^{\otimes t})}} \left[|\langle \phi_{k'} | \phi_k \rangle|^2 \right] \leq \text{negl}(n) \quad (2.19)$$

for all $t = \text{poly}(n)$, where $\mathbb{E}$ denotes the expectation value.

In this case, an immediate corollary of the Search HPS assumption is that the following construction yields a OWSG:

Construction 2.21: (OWSG from Search HPS) *Given $m(n)$, $q(n)$ and a distribution $\chi_{q,m,n}$ which satisfy the Search HPS assumption, define:*

1. **KeyGen**(1^n): *Sample a key $k \leftarrow \chi_{q,m,n}$.*
2. **StateGen**($k \in \mathbb{K}_{q,m,n}$): *Output an n -qubit state $|\phi_k\rangle \in \text{HPS}_{q,m,n}$.*
3. **Ver**($k, |\phi\rangle$): *Apply the projective measurement $\{\Pi_k^1, \Pi_k^0\} = \{|\phi_k\rangle\langle\phi_k|, \mathbb{I} - |\phi_k\rangle\langle\phi_k|\}$ on the state $|\phi\rangle$ and output the result. More specifically:*

$$\begin{aligned} \Pr[\text{Ver}(k, |\phi\rangle) \rightarrow 1] &= \text{tr}(\Pi_k^1 |\phi\rangle\langle\phi|) = |\langle\phi_k|\phi\rangle|^2, \\ \Pr[\text{Ver}(k, |\phi\rangle) \rightarrow 0] &= \text{tr}(\Pi_k^0 |\phi\rangle\langle\phi|) = 1 - |\langle\phi_k|\phi\rangle|^2. \end{aligned} \quad (2.20)$$

Remark 2.22: (Relation between HPS assumptions) Note that the Decision HPS assumption (Assumption 2.18) implies the Search HPS assumption (Assumption 2.20). Additionally, it can be that the functions $q(n)$ and $m(n)$ required to satisfy the Search HPS assumption (and therefore have a OWSG) are smaller than those required to satisfy the Decision HPS assumption (and therefore have a PRSG). Given that the Search HPS assumption is the weaker assumption, we focus from here on showing that the Search HPS assumption implies OWFs, which immediately yields the same conclusion for the Decision HPS assumption.

3 State certification

The notion of a "measure first, ask later" state certification protocol is central to this work. In order to define this rigorously, and to provide intuition for the definition, we start with the standard definition of a state certification protocol for an unknown target state [BOW17].

Definition 3.1: (State certification protocol for $|\phi\rangle$) We say that an algorithm $\mathcal{A}$ is a state certification protocol for an n -qubit state $|\phi\rangle$ from $t(n, \epsilon, \delta)$ copies if, for all states $|\psi\rangle$ and $(\epsilon, \delta) \in (0, 1)$:

1. If $|\psi\rangle = |\phi\rangle$ then

$$\Pr[\mathcal{A}(|\psi\rangle^{\otimes t}) = \text{Accept}] \geq 1 - \delta \quad (3.1)$$

2. If $|\langle\psi|\phi\rangle|^2 < 1 - \epsilon$ then

$$\Pr[\mathcal{A}(|\psi\rangle^{\otimes t}) = \text{Reject}] \geq 1 - \delta \quad (3.2)$$

We note that for *any* fixed state $|\phi\rangle$ there is a trivial state certification protocol with *constant* copy complexity with respect to n : Given $O(\epsilon^{-2} \log(\delta^{-1}))$ copies of $|\psi\rangle$ simply perform the projective measurement $\{|\phi\rangle\langle\phi|, \mathbb{I} - |\phi\rangle\langle\phi|\}$ on each copy, use this to estimate $|\langle\psi|\phi\rangle|^2$ sufficiently accurately, then make the appropriate decision. In other words, when one is allowed to measure in a $|\phi\rangle$ -dependent way – i.e., in a way which *depends on the target state* – then state certification is trivial. This observation makes clear that one could also consider the refined notion of a "measure first, ask later" state certification protocol. In such a protocol we insist that all measurements are *target state independent*, and that only the classical post-processing of measurement outcomes is target state dependent. More specifically, we have the following definition:

Definition 3.2: (“Measure first, ask later” state certification protocol for $|\phi\rangle$) Given an n -qubit state $|\phi\rangle$, consider a QPT algorithm $\mathcal{M}(|\psi\rangle^{\otimes t}) \rightarrow s \in \{0, 1\}^{g(t,n)}$ which does not depend on $|\phi\rangle$, and let $F_{|\phi\rangle} : \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\text{Accept}, \text{Reject}\}$ be a classical function which can depend on $|\phi\rangle$. We say that $(\mathcal{M}, F_{|\phi\rangle})$ is a “measure first, ask later” state certification protocol for $|\phi\rangle$ from $t(n, \epsilon, \delta)$ copies if, for all states $|\psi\rangle$ and $\epsilon, \delta \in (0, 1)$ with finite binary encodings:

1. If $|\psi\rangle = |\phi\rangle$ then

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [F_{|\phi\rangle}(s, \epsilon, \delta) = \text{Accept}] \geq 1 - \delta \quad (3.3)$$

2. If $|\langle\psi|\phi\rangle|^2 < 1 - \epsilon$ then

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [F_{|\phi\rangle}(s, \epsilon, \delta) = \text{Reject}] \geq 1 - \delta \quad (3.4)$$

We say that $(\mathcal{M}, F_{|\phi\rangle})$ is

1. *Copy efficient* if $t = \text{poly}(n, \epsilon^{-1}, \log \delta^{-1})$ is sufficient,
2. *Computationally efficient from O-access* if it is copy efficient and $F_{|\phi\rangle}$ is computationally efficient, given access to an oracle $\mathcal{O}(|\phi\rangle)$ which provides some notion of classical access to $|\phi\rangle$ (for example, an oracle which when queried on x provides the amplitude $\langle x|\phi\rangle$).
3. *Computationally efficient* if it is copy efficient and $F_{|\phi\rangle}$ is computationally efficient.
4. η -*simulable* if there exists a classical PPT algorithm $\mathcal{M}_C(1^n, 1^t) \rightarrow s \in \{0, 1\}^{g(t,n)}$ such that

$$d_{\text{TV}}(\mathcal{M}_C(1^n, 1^t), \mathcal{M}(|\phi\rangle^{\otimes t})) \leq \eta(n, t), \quad (3.5)$$

for all $t \in \mathbb{N}$, where $\mathcal{M}_C(1^n, 1^t)$ and $\mathcal{M}(|\phi\rangle^{\otimes t})$ are understood as distributions over $\{0, 1\}^{g(t,n)}$.

Remark 3.3: (Oracle access to the target state) We note that recent works on state certification assume different models of oracle access to the target state. For example, Ref. [HPS25] assumes an oracle which when queried on a bit string x , provides the amplitude $\langle x|\phi\rangle$, Ref. [CLS+26] considers an oracle which can provide amplitudes in either the computational or Hadamard basis, and Ref. [GHO25] considers an oracle which when queried on any sequence of single-qubit measurements and outcomes provides the probability of that specific outcome sequence. Whenever a “measure first, ask later” state certification protocol is efficient from O-access *and* that oracle can be efficiently classically simulated (with respect to n) for the specific target state, then the protocol will be computationally efficient.

We now want to generalize the definition above to a “measure first, ask later” state certification protocol for a *set* of states. To this end, it will be helpful to note that we can combine the two conditions given in the definition above into the single condition that

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} \left[\begin{cases} F_{|\phi\rangle}(s, \epsilon, \delta) = \text{Accept} \text{ if } |\psi\rangle = |\phi\rangle \\ F_{|\phi\rangle}(s, \epsilon, \delta) = \text{Reject} \text{ if } |\langle\psi|\phi\rangle|^2 < 1 - \epsilon \end{cases} \right] \geq 1 - \delta. \quad (3.6)$$

With this in hand, we can finally define the notion of a “measure first, ask later” state certification protocol for *simultaneously* certifying a *set* of states $\{|\phi_k\rangle \mid k \in [K]\}$ from a *single set* of measurement outcomes $s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})$.

Definition 3.4: (“Measure first, ask later” state certification protocol for $\{|\phi_k\rangle\}$) Let $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$ be a finite set of n -qubit states. Consider a QPT algorithm $\mathcal{M}(|\psi\rangle^{\otimes t}) \rightarrow s \in \{0, 1\}^{g(t,n)}$ and let $\mathcal{F} : \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\text{Accept}, \text{Reject}\}$ be such that $\mathcal{F}(\tilde{k}, \cdot, \cdot, \cdot) = \text{Reject}$ for all $\tilde{k} \notin \mathbb{K}$. We say that $(\mathcal{M}, \mathcal{F})$ is a “measure first, ask later” state certification protocol for $\{|\phi_k\rangle\}$ from $t = t(n, \epsilon, \delta, |\mathbb{K}|)$ copies if, for all states $|\psi\rangle$ and all $\epsilon, \delta \in (0, 1)$ with finite binary encodings:

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} \left[\forall k \in \mathbb{K} \begin{cases} \mathcal{F}(k, s, \epsilon, \delta) = \text{Accept} \text{ if } |\psi\rangle = |\phi_k\rangle \\ \mathcal{F}(k, s, \epsilon, \delta) = \text{Reject} \text{ if } |\langle\psi|\phi_k\rangle|^2 < 1 - \epsilon \end{cases} \right] \geq 1 - \delta. \quad (3.7)$$

We say that $(\mathcal{M}, \mathcal{F})$ is:

1. *Copy efficient* if $t = \text{poly}(n, \epsilon^{-1}, \log \delta^{-1}, \log |\mathbb{K}|)$ is sufficient.
2. *Computationally efficient from $\mathcal{O}$ -access* if it is copy efficient and the function $\mathcal{F}$ is computationally efficient given access to an oracle $\tilde{\mathcal{O}}$ which when queried on (k, x) returns the response from querying $\mathcal{O}(|\phi_k\rangle)$ on x for some oracle $\mathcal{O}(|\phi_k\rangle)$ providing some notion of classical access to $|\phi_k\rangle$.
3. *Computationally efficient* if it is copy efficient and $\mathcal{F}$ is computationally efficient.
4. η -simulable if there exists a classical PPT algorithm $\mathcal{M}_C(k, 1^n, 1^t) \rightarrow s \in \{0, 1\}^{g(t, n)}$ such that

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq \eta(|k|, n, t) \quad (3.8)$$

for all $k \in \mathbb{K}$ and $t \in \mathbb{N}$, where again $\mathcal{M}_C(k, 1^n, 1^t)$ and $\mathcal{M}(|\phi_k\rangle^{\otimes t})$ are understood as distributions over $\{0, 1\}^{g(t, n)}$.

Before we continue it is worth making a few remarks on the above definition.

Remark 3.5: (On the relation to classical shadows) We note that global Clifford shadows [HKP20] provides a paradigmatic *copy efficient* “measure first, ask later” state certification protocol for *any* set of states. However, for arbitrary sets of states this protocol will *not* be computationally efficient or η -simulable for meaningfully small values of η . Indeed, the term “measure first, ask later” has been deliberately borrowed from existing literature on randomized measurement protocols [EFH+23] to signify that Definition 3.4 is meant to provide a generalization and abstraction of “classical shadow type” state certification protocols.

Remark 3.6: (Relevant parameter ranges) Let us assume that the pair $(\mathcal{M}, \mathcal{F})$ is a measurement protocol for the set of n -qubit states $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$:

1. In many settings, one will typically be happy with both ϵ and δ being some small constants. However, in the cryptographic context we consider here, where we would like failure to occur with negligible probability, we will usually fix $\epsilon = 1/8$ and $\delta = 2^{-n}$. Additionally, we will work only with sets of states of size $|\mathbb{K}| = 2^{O(\text{poly}(n))}$. In this case, we have that $t = \text{poly}(n, \epsilon^{-1}, \log \delta^{-1}, \log |\mathbb{K}|)$ implies $t = O(\text{poly}(n))$ and that computational efficiency of $\mathcal{F}$ implies it can be evaluated in time $O(\text{poly}(n))$.
2. In order to be PPT, the algorithm $\mathcal{M}_C$ appearing in the definition of η -simulability should have runtime at most $\text{poly}(|k|, n, t)$ on input $(k, 1^n, 1^t)$. However, in the parameter ranges mentioned above – i.e. where $t = O(\text{poly}(n))$ and $K = 2^{O(\text{poly}(n))}$ – the runtime of $\mathcal{M}_C$ will be $O(\text{poly}(n))$.

Remark 3.7: (Refinements of “measure first, ask later” state certification protocols) We note that in Definition 3.4, we have placed no restrictions on the algorithm $\mathcal{M}$ beyond the fact that it is efficient. However, one can easily define the notion of a single-copy, or non-adaptive, “measure first, ask later” state certification protocol, by placing the appropriate restriction on $\mathcal{M}$. Intuitively, the more restrictions one places on $\mathcal{M}$, the harder it should be to construct such a state certification protocol for a given set of states. However, perhaps surprisingly, recent work has shown that (given access to an amplitude oracle) almost all quantum states can be certified by a “measure first, ask later” state certification protocol with *single qubit non-adaptive measurements* [HPS25].

Remark 3.8: (On η -simulability) We note that η -simulability means that for any $k \in \mathbb{K}$, one can efficiently sample (up to η accuracy) from the distribution over measurement outcomes obtained by measuring $|\phi_k\rangle^{\otimes t}$ via $\mathcal{M}$. While the utility or relevance of this property may not be as immediately clear as the efficiency properties, we will see in the following sections that simulability of a measurement protocol is what allows us to obtain a OWF from the OWP constructed via “measure first, ask later” state certification protocols.

4 Certifiable Microcrypt primitives

Equipped with the notion of a “measure first, ask later” state certification protocol for a set of states we can finally define *certifiable* variants of OWSG and PRSG.

Definition 4.1: (Certifiable one-way state generator (C-OWSG) and certifiable pseudorandom state generator (C-PRSG)) Consider a one-way state generator ($\text{KeyGen}, \text{StateGen}, \text{Ver}$) (or pseudorandom state generator ($\text{KeyGen}, \text{StateGen}$)), with $\mathbb{K}_\lambda := \text{supp}(\text{KeyGen}(1^\lambda))$ and $|\phi_k\rangle = \text{StateGen}(k)$ for all $k \in \mathbb{K}_\lambda$. Given a “measure first, ask later” state certification protocol $(\mathcal{M}, \mathcal{F})$ for the set of states $\{|\phi_k\rangle \mid k \in \mathbb{K}_\lambda\}$, we say that the tuple $(\text{KeyGen}, \text{StateGen}, \text{Ver}, (\mathcal{M}, \mathcal{F}))$ (or tuple $(\text{KeyGen}, \text{StateGen}, (\mathcal{M}, \mathcal{F}))$) is:

1. A *certifiable* one-way state generator (pseudorandom state generator) if $(\mathcal{M}, \mathcal{F})$ is copy efficient.
2. An *η -simulable certifiable* one-way state generator (pseudorandom state generator) if $(\mathcal{M}, \mathcal{F})$ is copy efficient and η -simulable.
3. An *efficiently certifiable* one-way state generator (pseudorandom state generator) if $(\mathcal{M}, \mathcal{F})$ is computationally efficient.
4. An *η -simulable efficiently certifiable* one-way state generator (pseudorandom state generator) if $(\mathcal{M}, \mathcal{F})$ is computationally efficient and η -simulable.

A-priori the utility of distinguishing carefully between all the certifiable OWSG variants defined above may not be clear. However, we show in the following sections that each of the above variants can be used to construct different cryptographic primitives. In particular:

1. Certifiable OWSGs can be used to construct OWPs (see Theorem 5.1).
2. Efficiently certifiable OWSGs can be used to construct EV-OWPs (see Corollary 5.3).
3. $1/3$ -simulable certifiable OWSGs can be used to construct OWFs via distributional OWFs (see Theorem 6.1).
4. η -simulable efficiently certifiable OWSGs can be used to directly construct OWFs (without going via distributional OWFs) whenever $\eta(|k|, n, t) = \text{negl}(\lambda)$ for $|k|, n$ and t which are all $O(\text{poly}(\lambda))$.

Given the above, we see that the same OWSGs can be used to construct a variety of different cryptographic primitives (in different cryptographic worlds), depending on the properties of the “measure first, ask later” state certification protocol with which it is equipped! As such, all the variants of certifiable OWSG defined above really should be considered as different primitives.

Remark 4.2: (All OWSG are certifiable) Recall from Remark 3.5 that global Clifford shadows provides a copy-efficient “measure first, ask later” state certification protocol for *any* set of states. Therefore, *any* OWSG (PRSG) equipped with the state certification protocol provided by global Clifford shadows is immediately a certifiable OWSG (PRSG). Using the results mentioned above (in particular Theorem 5.1) this means that one can construct a OWP from any OWSG, by using the global Clifford shadows state certification protocol. Indeed, this is precisely how Khurana and Tomer have proven that OWSG can be used to construct OWP [KT24].

Remark 4.3: (Further variants of certifiable primitives) As per Remark 3.7, it is clear that one can define restricted variations of “measure first, ask later” state certification protocols, and using these one could define variants of certifiable OWSG (PRSG) such as single-copy and non-adaptive certifiable and efficiently certifiable OWSG and PRSG. Given that the more restrictions one places on the state certification protocol the harder it is to construct the primitive, one could conjecture that it is possible to construct more powerful cryptographic protocols from restricted variants of the certifiable microcrypt primitives we introduce here.

Finally, with the above established, we make the observation that variants of certifiable *standard* PRSG (i.e., with super-logarithmic stretch) can be used to construct variants of certifiable OWSG. This follows directly from the fact that given a standard PRSG ($\text{KeyGen}, \text{StateGen}$) (i.e., with stretch $n(\lambda) = \omega(\log \lambda)$) the tuple $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ with the standard verification algorithm is a OWSG [CGG+25]. Specifically, as an immediate corollary of this prior result, together with the definitions of certifiable PRSG and OWSG, we have the following:

Corollary 4.4: (Certifiable PRSG imply certifiable OWSG) *Let $(\text{KeyGen}, \text{StateGen})$ be a standard pseudorandom state generator – i.e., with stretch $n(\lambda) = \omega(\log \lambda)$. Let Ver be as per Definition 2.10. Then:*

1. *If $((\text{KeyGen}, \text{StateGen}), (\mathcal{M}, \mathcal{F}))$ is a (efficiently) certifiable PRSG then $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ is a (efficiently) certifiable OWSG*
2. *If $((\text{KeyGen}, \text{StateGen}), (\mathcal{M}, \mathcal{F}))$ is an η -simulable (efficiently) certifiable PRSG then $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ is an η -simulable (efficiently) certifiable OWSG*

5 One-way puzzles from certifiable one-way state generators

Equipped with the tools of the previous section, we start here by proving that one can construct a OWP from *any* certifiable OWSG. This generalizes and abstracts the construction of OWP from OWSG by Khurana and Tomer [KT24], which utilized the *specific* “measure first, ask later” state certification protocol provided by global Clifford shadows. In particular, the theorem we provide below simply makes it clear that one can replace the global Clifford shadows protocol used in Ref. [KT24] with *any* copy-efficient “measure first, ask later” state certification protocol for the output states of the starting OWSG. As already mentioned, this is helpful because if the state certification protocol satisfies additional properties – such as being *computationally efficient* or η -*simulable* – then the OWP will also inherit these properties. As such, the result below provides a route to the construction of *variants* of OWPs – by utilizing *tailored* state certification protocols for the output states of the input OWSG – which one could not hope to obtain by only ever using global Clifford shadows to construct OWPs from OWSGs.

Theorem 5.1: (One-way puzzles from certifiable one-way state generators) *Given a certifiable one-way state generator, one can construct a one-way puzzle.*

Proof. Given a C-OWSG $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ with copy efficient “measure first, ask later” state certification protocol $(\mathcal{M}, \mathcal{F})$ from $t(n, \epsilon, \delta, |\mathbb{K}|) = \text{poly}(n, 1/\epsilon, \log(1/\delta), \log |\mathbb{K}|)$ copies, consider the following construction.

Construction 5.2: (OWP from C-OWSG) *Given a C-OWSG $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ with copy efficient “measure first, ask later” state certification protocol $(\mathcal{M}, \mathcal{F})$ from $t(n, \epsilon, \delta, |\mathbb{K}|) = \text{poly}(n, 1/\epsilon, \log(1/\delta), \log |\mathbb{K}|)$ copies*

1. *Define $\mathbb{K}_\lambda = \text{supp}(\text{KeyGen}(1^\lambda))$.*
2. *For all $k \in \mathbb{K}_\lambda$, define $|\phi_k\rangle$ as the $n = \text{poly}(\lambda)$ -qubit state $|\phi_k\rangle = \text{StateGen}(k)$.*

By definition, $(\mathcal{M}, \mathcal{F})$ is a copy efficient “measure first, ask later” state certification protocol for $\{|\phi_k\rangle\}_{k \in \mathbb{K}_\lambda}$. With this in hand, define the algorithms $(\text{SampPuzz}, \text{VerPuzz})$ as follows:

1. **SampPuzz.** *On input 1^λ :*
 - (a) *Run $\text{KeyGen}(1^\lambda)$ and obtain a key $k \in \mathbb{K}_\lambda$.*
 - (b) *Run $\text{StateGen}(k)$ $t = t(n, \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)$ times, and construct $|\phi_k\rangle^{\otimes t}$.*
 - (c) *Run $\mathcal{M}(|\phi_k\rangle^{\otimes t})$ and obtain s .*
 - (d) *Output (k, s) .*
2. **VerPuzz** *is defined via*

$$\text{VerPuzz}(k, s) = \begin{cases} 1, & \mathcal{F}(k, s, 1/8, 2^{-\lambda}) = \text{Accept}, \\ 0, & \mathcal{F}(k, s, 1/8, 2^{-\lambda}) = \text{Reject}. \end{cases}$$

Define the output OWP as the tuple $(\text{SampPuzz}, \text{VerPuzz})$.

We now prove that the algorithms $(\text{SampPuzz}, \text{VerPuzz})$ defined in Construction 5.2 constitute a OWP.

To this end, we start by proving that **SampPuzz** is a QPT algorithm with respect to λ . To do this note:

1. As a consequence of the fact that KeyGen is a QPT algorithm one has $|\mathbb{K}_\lambda| = 2^{\text{poly}(\lambda)}$ and $\log |\mathbb{K}_\lambda| = \text{poly}(\lambda)$.
2. As a consequence of the fact that StateGen is a QPT algorithm one has $n = \text{poly}(\lambda)$.
3. Therefore via the fact that $(\mathcal{M}, \mathcal{F})$ is a copy efficient measurement protocol from

$$t(n, \epsilon, \delta, |\mathbb{K}|) = \text{poly}(n, 1/\epsilon, \log(1/\delta), \log |\mathbb{K}|) \quad (5.1)$$

copies, we have that $t(n, 1/8, 2^{-\lambda}, |\mathbb{K}_\lambda|) = \text{poly}(\lambda)$.

4. Together with the assumption that $\mathcal{M}$ is a QPT algorithm (implicit in the assumption that $(\mathcal{M}, \mathcal{F})$ is a state certification protocol), this implies that SampPuzz is a QPT algorithm.

Now, let us establish the correctness and security of the pair $(\text{SampPuzz}, \text{VerPuzz})$. In what follows, for convenience we will drop the ϵ and δ arguments from $\mathcal{F}$ with the understanding that $\mathcal{F}(k, s) := \mathcal{F}(k, s, 1/8, 2^{-\lambda})$ for all (k, s) . With this in mind, we start by noting that with $t = t(n, \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)$ one has

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} \left[\forall k \in \mathbb{K}_\lambda \begin{cases} \mathcal{F}(k, s) = \text{Accept} \text{ if } |\psi\rangle = |\phi_k\rangle \\ \mathcal{F}(k, s) = \text{Reject} \text{ if } |\langle \psi | \phi_k \rangle|^2 < 7/8 \end{cases} \right] > 1 - 2^{-\lambda}. \quad (5.2)$$

We can now establish correctness and security.

Correctness: Using the properties of the state certification protocol $(\mathcal{M}, \mathcal{F})$ we have:

$$\begin{aligned} \Pr_{(k,s) \leftarrow \text{SampPuzz}(1^\lambda)} [\text{VerPuzz}(k, s) = 1] &= \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ s \leftarrow \mathcal{M}(|\phi_k\rangle^t)}} [\mathcal{F}(k, s) = \text{Accept}] \\ &= \sum_{k \in \mathbb{K}_\lambda} \left[\Pr(k) \left(\Pr_{s \leftarrow \mathcal{M}(|\phi_k\rangle^t)} [\mathcal{F}(k, s) = \text{Accept}] \right) \right] \\ &\geq \sum_{k \in \mathbb{K}_\lambda} \left[\Pr(k) \left(\Pr_{s \leftarrow \mathcal{M}(|\phi_k\rangle^t)} \left[\forall j \in \mathbb{K}_\lambda \begin{cases} \mathcal{F}(j, s) = \text{Accept} \text{ if } |\phi_k\rangle = |\phi_j\rangle \\ \mathcal{F}(j, s) = \text{Reject} \text{ if } |\langle \phi_k | \phi_j \rangle|^2 < 7/8 \end{cases} \right] \right) \right] \\ &> \sum_{k \in \mathbb{K}_\lambda} [\Pr(k)(1 - 2^{-\lambda})] \\ &= 1 - 2^{-\lambda}. \end{aligned}$$

Security: Note that

$$\begin{aligned} \Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [\text{VerPuzz}(k', s) = 1] &= \Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [\mathcal{F}(k', s) = \text{Accept}] \\ &= \Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [A(k, s, k') \vee B(k, s, k')], \end{aligned} \quad (5.3)$$

where

$$\begin{aligned} A(k, s, k') &:= [\mathcal{F}(k', s) = \text{Accept}] \wedge [|\langle \phi_{k'} | \phi_k \rangle|^2 < 7/8], \\ B(k, s, k') &:= [\mathcal{F}(k', s) = \text{Accept}] \wedge [|\langle \phi_{k'} | \phi_k \rangle|^2 \geq 7/8]. \end{aligned} \quad (5.4)$$

Now, using the guarantee of the state certification protocol $(\mathcal{M}, \mathcal{F})$, one has

$$\begin{aligned} \Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [A(k, s, k')] &= \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ s \leftarrow \mathcal{M}(|\phi_k\rangle^t) \\ k' \leftarrow \mathcal{A}(s)}}} [\mathcal{F}(k', s) = \text{Accept}] \wedge [|\langle \phi_{k'} | \phi_k \rangle|^2 < 7/8] \\ &\leq \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ s \leftarrow \mathcal{M}(|\phi_k\rangle^t)}} [\exists k' \in \mathbb{K}_\lambda ([\mathcal{F}(k', s) = \text{Accept}] \wedge [|\langle \phi_{k'} | \phi_k \rangle|^2 < 7/8])] \\ &\leq 2^{-\lambda}. \end{aligned} \quad (5.5)$$

Additionally,

$$\begin{aligned}
\Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [B(k, s, k')] &= \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ s \leftarrow \mathcal{M}(|\phi_k\rangle^t) \\ k' \leftarrow \mathcal{A}(s)}} \left[[\mathcal{F}(k', s) = \text{Accept}] \wedge [|\langle \phi_{k'} | \phi_k \rangle|^2 \geq 7/8] \right] \\
&\leq \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ s \leftarrow \mathcal{M}(|\phi_k\rangle^t) \\ k' \leftarrow \mathcal{A}(s)}} [|\langle \phi_{k'} | \phi_k \rangle|^2 \geq 7/8] \\
&:= \beta.
\end{aligned} \tag{5.6}$$

We now show that $\beta \leq \text{negl}(\lambda)$. To see this, recall from Definition 2.10 that $\text{Ver}(k, |\psi\rangle)$ corresponds to the outcome of the projective measurement $\{|\phi_k\rangle\langle\phi_k|, \mathbb{I} - |\phi_k\rangle\langle\phi_k|\}$ applied to the state $|\psi\rangle$. Then, for all QPT algorithms $\mathcal{A}$ one has

$$\text{negl}(\lambda) \geq \Pr_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ k' \leftarrow \mathcal{A} \circ \mathcal{M}(|\phi_k\rangle^{\otimes t})}} [\text{Ver}(k', |\phi_k\rangle) = 1] = \mathbb{E}_{\substack{k \leftarrow \text{KeyGen}(1^\lambda) \\ k' \leftarrow \mathcal{A} \circ \mathcal{M}(|\phi_k\rangle^{\otimes t})}} [|\langle \phi_{k'} | \phi_k \rangle|^2] \geq \frac{7}{8}\beta \tag{5.7}$$

where the first inequality follows from the security of the OWSG, and the final inequality follows from Markov's inequality. Therefore, $\beta \leq \text{negl}(\lambda)$. By a union bound, we then have

$$\Pr_{\substack{(k,s) \leftarrow \text{SampPuzz}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}}} [\text{VerPuzz}(k', s) = 1] \leq \text{negl}(\lambda). \tag{5.8}$$

□

With the above established, we obtain as promised an immediate corollary that if the input OWSG is *efficiently* certifiable, then Construction 5.2 actually yields an *efficiently verifiable* OWP. Specifically, we have the following.

Corollary 5.3: (Efficiently verifiable one-way puzzle from efficiently certifiable one-way state generator) *Given an efficiently certifiable one-way state generator one can construct an efficiently verifiable one-way puzzle.*

6 One-way functions from simulable certifiable one-way state generators

We have already seen in the previous section that given an (efficiently) certifiable OWSG, one can construct an (efficiently verifiable) OWP. In this section, we show the following:

1. If one is given a $1/3$ -simulable certifiable OWSG (with a classical KeyGen algorithm), then one can construct a quantum-secure OWF from the OWP obtained from the OWSG. More specifically, we will show that one can construct a suitably accurate classical sampling algorithm for the OWP, and therefore via the results of Corollary 2.8 one can construct a quantum-secure distributional OWF, from which via Theorem 2.5 one can construct a quantum-secure OWF.
2. If one is given an η -simulable *efficiently* certifiable OWSG (with a classical KeyGen algorithm), for a function η which is negligible in a sense to be made precise, then one can construct a simple quantum-secure OWF directly from the OWP obtained by the OWSG, *without* having to go via a distributional OWF. In particular, this OWF avoids the complexities incurred in the previous construction due to first compiling the distributional OWF into a weak OWF, and then compiling this weak OWF into a OWF. Essentially we show here that if one has a measurement protocol which is both efficiently certifiable, and η simulable for a “negligible” η , then one can obtain a simpler OWF.

6.1 OWFs from simulable certifiable OWSGs via distributional OWFs

We prove in this section that given a $1/3$ -simulable certifiable OWSG, with a classical PPT key generation algorithm KeyGen , one can construct a quantum-secure OWF.

Theorem 6.1: (OWF from 1/3-simulable certifiable OWSG with PPT KeyGen) *Given a 1/3-simulable certifiable OWSG $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$, in which KeyGen is a classical PPT algorithm, one can construct a quantum-secure one-way function.*

Proof. The idea of the proof is the following:

1. We know from Theorem 5.1 that given a certifiable OWSG $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$, one can construct a OWP $(\text{SampPuzz}, \text{VerPuzz})$ via Construction 5.2.
2. Therefore, if one can use the additional 1/3-simulability of $(\mathcal{M}, \mathcal{F})$ to construct a classical PPT algorithm SampPuzz_C satisfying

$$d_{\text{TV}}(\text{SampPuzz}(1^\lambda), \text{SampPuzz}_C(1^\lambda)) \leq 1/3, \quad (6.1)$$

then it follows from Corollary 2.8 that one can construct an explicit quantum-secure distributional OWF.

3. Given the quantum-secure distributional OWF, it then follows from Theorem 2.5 that one can construct a quantum-secure OWF.

As such, we show that one can construct a classical PPT algorithm satisfying Eq. (6.1). To this end, recall from Construction 5.2 that SampPuzz is defined as follows, on input 1^λ :

1. $k \leftarrow \text{KeyGen}(1^\lambda)$ (where KeyGen is assumed to be classical PPT).
2. $t \leftarrow t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)$, where $t(n, \epsilon, \delta, |\mathbb{K}_\lambda|)$ is the copy complexity of $(\mathcal{M}, \mathcal{F})$.
3. $|\phi_k\rangle \leftarrow \text{StateGen}(k)$
4. $s \leftarrow \mathcal{M}(|\phi_k\rangle^{\otimes t})$.
5. Output (k, s) .

The natural idea is then to construct SampPuzz_C by simply replacing steps 3 and 4 above with the implementation of $\mathcal{M}_C(k, 1^n, 1^t)$, where $\mathcal{M}_C$ is the η -accurate simulation of $\mathcal{M}$ guaranteed by η -simulability of $(\mathcal{M}, \mathcal{F})$. Specifically, define SampPuzz_C , on input 1^λ , via:

1. $k \leftarrow \text{KeyGen}(1^\lambda)$.
2. $t \leftarrow t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)$.
3. $s \leftarrow \mathcal{M}_C(k, 1^{n(\lambda)}, 1^t)$.
4. Output (k, s) .

Given this, we start by proving that SampPuzz_C is a PPT algorithm. To see this:

1. Note that we have assumed that KeyGen is a classical PPT algorithm, and recall from Section 2.2 that efficiency of both KeyGen and StateGen implies that $\mathbb{K}_\lambda \subseteq \{0, 1\}^{O(\text{poly}(\lambda))}$, and therefore both $|k| = O(\text{poly}(\lambda))$ and $n(\lambda) = O(\text{poly}(\lambda))$.
2. With the above, it follows from the copy efficiency of $(\mathcal{M}, \mathcal{F})$ that $t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|) = O(\text{poly}(\lambda))$.
3. Therefore, it follows from the fact that $\mathcal{M}_C$ is a PPT algorithm with a runtime on input $(k, 1^n, 1^t)$ of $O(\text{poly}(|k|, n, t)) = O(\text{poly}(\lambda))$.
4. As a result, the runtime of SampPuzz_C is $O(\text{poly}(\lambda))$, and it is also a PPT algorithm.

With this established, we now have to show that SampPuzz_C is sufficiently accurate. To do this, recall that 1/3-simulability of $(\mathcal{M}, \mathcal{F})$ enforces

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq 1/3, \quad (6.2)$$

for all (k, n, t) , from which it follows that

$$d_{\text{TV}}(\text{SampPuzz}(1^\lambda), \text{SampPuzz}_C(1^\lambda)) \leq 1/3, \quad (6.3)$$

by virtue of the fact that k is sampled identically in both SampPuzz and SampPuzz_C . $\square$

6.2 Direct OWF construction via efficient verifiability

In the previous section, we have shown that given a $1/3$ -simulable certifiable OWSG $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$, with classical PPT KeyGen , one can construct a quantum-secure OWF. However, the construction of this OWF involves multiple steps:

1. The construction of a OWP $(\text{SampPuzz}, \text{VerPuzz})$, via Construction 5.2.
2. The construction of a $1/3$ -accurate classical PPT sampling algorithm SampPuzz_C .
3. The construction of a quantum-secure *distributional* OWF $\{f_\lambda\}$ from SampPuzz_C (As per Corollary 2.8).
4. The construction of a quantum-secure one-way function $\{g_\lambda\}$ via the construction underlying Theorem 2.5. In particular, this construction proceeds via first constructing a quantum-secure *weak* OWF $\{h_\lambda\}$, and then constructing the quantum-secure OWF $\{g_\lambda\}$ from $\{h_\lambda\}$.

In this section, we show that if one starts from an η -simulable *efficiently certifiable* OWSG (with classical PPT KeyGen), for a function η which is negligible in a sense that we will make precise shortly, then a simple modification of the quantum-secure distributional OWF constructed in Step 3 above is immediately a quantum-secure OWF. In particular, one can obtain a relatively simple OWF without the chain of compilations

$$\text{distributional OWF} \rightarrow \text{weak OWF} \rightarrow \text{OWF} \quad (6.4)$$

that was required in the previous section. To do this:

1. We start by showing in Theorem 6.2 that given an *efficiently verifiable* OWP with a *negligibly* inaccurate classical sampling algorithm, a simple modification of the distributional OWF defined in Theorem 2.7 and Corollary 2.8 is immediately a quantum-secure OWF.
2. We already know from Corollary 5.3 that one can construct an EV-OWP from an efficiently certifiable OWSG. So, we simply show (using essentially the proof of Theorem 6.1) that when the input efficiently certifiable OWSG is also η -simulable, for appropriate η , then one can construct a negligibly inaccurate classical PPT sampling algorithm for the EV-OWP, and therefore by the previous point, directly construct a quantum-secure OWF.

To achieve point 1 above, it is helpful to understand why the arguments establishing *distributional* one-wayness of the function family $\{f_\lambda\}$ in Theorem 2.7, fail to prove that the same function is a OWF. To this end, recall from Theorem 2.7 that given a OWP $(\text{Samp}, \text{Ver})$ and an efficient deterministic *classical* algorithm Samp_C satisfying

$$d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq \frac{1}{3}, \quad (6.5)$$

with $\text{Samp}_C(1^\lambda, r) = (k_\lambda(r), s_\lambda(r))$, then the function family $\{f_\lambda\}$ considered by Khurana and Tomer (and shown to be a *distributional* OWF) is defined via $f_\lambda(r) = s_\lambda(r)$.

As an illustration, let's now imagine we have the even stronger assumption that

$$\text{Samp}(1^\lambda) = \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}, \quad (6.6)$$

so that *correctness* of the OWP becomes

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(k, s) = 1] = \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(k_\lambda(r), s_\lambda(r)) = 1] \geq 1 - \text{negl}, \quad (6.7)$$

and try to prove that $\{f_\lambda\}$ is in fact a standard OWF. The natural approach is to assume that there exists some successful adversary $\mathcal{A}$ for the OWF, and show that this can be used to construct a successful adversary $\mathcal{B}$ for the

underlying OWP. So, let's try to do this, and to make our job even easier, let's assume we have an extremely strong adversary for the OWF, which *always* succeeds, i.e., a PPT algorithm $\mathcal{A}$ satisfying

$$\Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, f_\lambda(r))}} [r' \in f_\lambda^{-1}(f_\lambda(r))] = \Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [r' \in s_\lambda^{-1}(s_\lambda(r))] = 1. \quad (6.8)$$

Then, we define our candidate OWP adversary $\mathcal{B}$ such that on input $(1^\lambda, s)$ with $(k, s) \leftarrow \text{Samp}(1^\lambda)$:

1. $r' \leftarrow \mathcal{A}(1^\lambda, s)$,
2. $k' \leftarrow k_\lambda(r')$.

Indeed, note that if $\mathcal{A}$ is such that for all $r \in \{0, 1\}^{m(\lambda)}$ one has $\mathcal{A}(1^\lambda, s_\lambda(r)) = r$, then using Eq. (6.7), one has

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(\mathcal{B}(1^\lambda, s), s) = 1] = \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(\mathcal{B}(1^\lambda, s_\lambda(r)), s_\lambda(r)) = 1] \quad (6.9)$$

$$= \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(k_\lambda(\mathcal{A}(1^\lambda, s_\lambda(r))), s_\lambda(r)) = 1] \quad (6.10)$$

$$= \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(k_\lambda(r), s_\lambda(r)) = 1] \quad (6.11)$$

$$\geq 1 - \text{negl}(\lambda), \quad (6.12)$$

i.e., $\mathcal{B}$ is a (very) succesful adversary for the OWP. As such, it seems intuitive/plausible that one could use Eq. (6.8) to show that

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(\mathcal{B}(1^\lambda, s), s)] = \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(\mathcal{B}(1^\lambda, s_\lambda(r)), s_\lambda(r))] \quad (6.13)$$

$$= \Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [\text{Ver}(k_\lambda(r'), s_\lambda(r))] \quad (6.14)$$

$$\geq \tilde{g}(\lambda) \quad (6.15)$$

for some non-negligible function $\tilde{g}$. Unfortunately, however, there is a problem! To see this, let's define the sets Good_λ and Bad_λ via

$$\text{Good}_\lambda = \{r \mid \text{Ver}(k_\lambda(r), s_\lambda(r)) = 1\}, \quad (6.16)$$

$$\text{Bad}_\lambda = \{r \mid \text{Ver}(k_\lambda(r), s_\lambda(r)) = 0\}. \quad (6.17)$$

Note that the correctness condition of the OWP then becomes

$$\Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [r \in \text{Good}_\lambda] \geq 1 - \text{negl}. \quad (6.18)$$

In particular, it can be the case that Bad_λ is non-empty. Let's assume that this is the case, and note that nothing prevents an adversary $\mathcal{A}$ which, as illustrated in Figure 2, in addition to satisfying Eq. (6.8) is fine-tuned so that with high probability it also outputs some $r' \in \text{Bad}_\lambda$, i.e.

$$\Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [r' \in \text{Bad}_\lambda] \geq 1 - \text{negl}. \quad (6.19)$$




Figure 2: An illustration of a “problematic” adversary $\mathcal{A}$ for the function $\{f_\lambda = s_\lambda\}$, which on input $s_\lambda(r)$ outputs some $r' \in f_\lambda^{-1}(f_\lambda(r)) \cap \text{Bad}_\lambda$.

Then, using both Eq (6.8) and Eq. (6.19) we have

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(\mathcal{B}(1^\lambda, s), s) = 1] = \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(\mathcal{B}(1^\lambda, s_\lambda(r)), s_\lambda(r)) = 1] \quad (6.20)$$

$$= \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(k_\lambda(\mathcal{A}(1^\lambda, s_\lambda(r))), s_\lambda(r)) = 1] \quad (6.21)$$

$$= \Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [\text{Ver}(k_\lambda(r'), s_\lambda(r)) = 1] \quad (6.22)$$

$$= \Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [\text{Ver}(k_\lambda(r'), s_\lambda(r')) = 1] \quad (6.23)$$

$$= \Pr_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ r' \leftarrow \mathcal{A}(1^\lambda, s_\lambda(r))}} [r' \in \text{Good}_\lambda] \quad (6.24)$$

$$< \text{negl}(\lambda), \quad (6.25)$$

i.e., $\mathcal{B}$ is *not* a successful adversary for the OWP, and our proof strategy fails! As such, we see that if we want to build an adversary for the OWP from an adversary for the OWF, it would be helpful if we could somehow “force” the OWF adversary to always output some $r' \in \text{Good}_\lambda$. Luckily, this can be achieved by simply considering the slightly modified function family $\{F_\lambda : \{0, 1\}^{m(\lambda)} \rightarrow \{0, 1\}^*\}_{\lambda \in \mathbb{N}}$ defined via

$$F_\lambda(r) = \begin{cases} (0, s_\lambda(r)) & r \in \text{Good}_\lambda, \\ (1, r) & r \in \text{Bad}_\lambda, \end{cases} \quad (6.26)$$

In particular, with this construction, if $r \in \text{Good}_\lambda$, then on input $(0, s_\lambda(r))$, a successful OWF adversary has to output some $r' \in \text{Good}_\lambda$, otherwise $F_\lambda(r) \neq F_\lambda(r')$. One might immediately object that on inputs $r \in \text{Bad}_\lambda$, the function is trivial to invert, but this is not problematic given that correctness enforces that this happens only with negligible probability. However, it is important to note that in order for $\{F_\lambda\}$ to be efficient to compute, we have to also insist that Ver is efficient – i.e., that the underlying OWP is an *efficiently verifiable* OWP.

Using these ideas, we then have the following:

Theorem 6.2: (One-way functions from classical efficiently verifiable one-way puzzles) *Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be some fixed polynomial. Given an efficiently verifiable one-way puzzle $(\text{Samp}, \text{Ver})$, assume there exists an efficient deterministic classical algorithm Samp_C such that*

$$d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq \text{negl}(\lambda) \quad (6.27)$$

where again $\text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}} = (k_\lambda(r), s_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}$ is understood as the distribution over outputs of $\text{Samp}_C(1^\lambda, r)$ with respect to bitstrings $r \in \{0, 1\}^{m(\lambda)}$ drawn uniformly at random. Then, the family of functions $\{F_\lambda : \{0, 1\}^{m(\lambda)} \rightarrow$

$\{0, 1\}^*_{\lambda \in \mathbb{N}}$ defined via

$$F_\lambda(r) = \begin{cases} (0, s_\lambda(r)) & r \in \text{Good}_\lambda, \\ (1, r) & r \in \text{Bad}_\lambda, \end{cases} \quad (6.28)$$

is a quantum-secure one-way function.

Proof. Let $(\text{Samp}, \text{Ver})$ be an EV-OWP. We start with a simple "robustness lemma", whose proof we defer to Appendix D, showing that replacing Samp with any negligibly close approximation in TV distance yields a new OWP.

Lemma 6.3: (Robustness of one-way puzzles) *Let $(\text{Samp}, \text{Ver})$ be a one-way puzzle, and let Samp' be any QPT or PPT algorithm satisfying*

$$\epsilon(\lambda) = d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}'(1^\lambda)) \leq \text{negl}(\lambda). \quad (6.29)$$

Then, $(\text{Samp}', \text{Ver})$ is a one-way puzzle.

With this in hand, it follows immediately that if, as per the theorem statement, Samp_C satisfies

$$d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq \text{negl}(\lambda) \quad (6.30)$$

then $(\text{Samp}_C, \text{Ver})$ is an (efficiently verifiable) OWP. Now, as already noted, the efficiency of F_λ follows from the efficiency of Samp_C and Ver . To obtain a contradiction, let's assume that $\{F_\lambda\}$ is *not* a quantum-secure OWF. In this case, there exists some QPT inversion algorithm $\mathcal{A}$ satisfying

$$\Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [F_\lambda(\mathcal{A}(1^\lambda, F_\lambda(r))) = F_\lambda(r)] \geq \epsilon(\lambda) \quad (6.31)$$

for some non-negligible ϵ . Alternatively, if we define the event

$$\text{Succ}_{\mathcal{A}} = \{F_\lambda(\mathcal{A}(1^\lambda, F_\lambda(r))) = F_\lambda(r)\}, \quad (6.32)$$

then

$$\Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Succ}_{\mathcal{A}}] \geq \epsilon(\lambda). \quad (6.33)$$

Using this, we now construct a QPT solver $\mathcal{B}$ for the OWP $(\text{Samp}_C, \text{Ver})$. In particular, define $\mathcal{B}$ as follows:

1. On input s and 1^λ , obtain $r' \leftarrow \mathcal{A}(1^\lambda, (0, s))$.
2. Output $k' \leftarrow k_\lambda(r')$.

Note that in the case that $s = s_\lambda(r)$ for some $r \in \text{Good}_\lambda$, then by construction of the OWF, if $\mathcal{A}$ succeeds in Step 1, then we know that $s_\lambda(r') = s_\lambda(r)$ and that $r' \in \text{Good}_\lambda$, and therefore that

$$\text{Ver}(\mathcal{B}(s_\lambda(r)), s_\lambda(r)) = \text{Ver}(k_\lambda(r'), s_\lambda(r)) = \text{Ver}(k_\lambda(r'), s_\lambda(r')) = 1. \quad (6.34)$$

With this in mind, note that using the correctness condition of $(\text{Samp}_C, \text{Ver})$, namely that

$$\Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(k_\lambda(r), s_\lambda(r)) = 1] = \Pr_{r \in \{0,1\}^{m(\lambda)}} [r \leftarrow \text{Good}_\lambda] \geq 1 - \delta(\lambda) \quad (6.35)$$

for some negligible δ , we have that

$$\begin{aligned} \Pr_{r \leftarrow \{0,1\}^{m(\lambda)}} [\text{Ver}(\mathcal{B}(s_\lambda(r)), s_\lambda(r)) = 1] &\geq \Pr[(r \in \text{Good}_\lambda) \wedge \text{Succ}_{\mathcal{A}}] \\ &= \Pr[\text{Succ}_{\mathcal{A}}] - \Pr[(r \notin \text{Good}_\lambda) \wedge \text{Succ}_{\mathcal{A}}] \\ &\geq \Pr[\text{Succ}_{\mathcal{A}}] - \Pr[(r \notin \text{Good}_\lambda)] \\ &\geq \epsilon(\lambda) - \delta(\lambda), \end{aligned} \quad (6.36)$$

which is non-negligible, contradicting the assumption that $(\text{Samp}_C, \text{Ver})$ is a OWP. $\square$

With this established, we can now prove the following result:

Theorem 6.4: (OWF directly from negl-simulable efficiently certifiable OWSG) *Given an η -simulable efficiently certifiable one-way state generator $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$, in which KeyGen is a classical PPT algorithm, if*

$$\sup_{k \in \mathbb{K}_\lambda} \eta(|k|, n(\lambda), t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)) = \text{negl}(\lambda), \quad (6.37)$$

then one can construct an efficient deterministic classical algorithm Samp_C , with $\text{Samp}_C(1^\lambda, r) = (k_\lambda(r), s_\lambda(r))$, such that

$$d_{\text{TV}}(\text{SampPuzz}(1^\lambda), \text{SampPuzz}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq \text{negl}(\lambda), \quad (6.38)$$

where $(\text{SampPuzz}, \text{VerPuzz})$ is the EV-OWP constructed from $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ via Construction 5.2. Then, the family of functions $\{F_\lambda : \{0,1\}^{m(\lambda)} \rightarrow \{0,1\}^*\}_{\lambda \in \mathbb{N}}$ defined via

$$F_\lambda(r) = \begin{cases} (0, s_\lambda(r)) & r \in \text{Good}_\lambda, \\ (1, r) & r \in \text{Bad}_\lambda, \end{cases} \quad (6.39)$$

is a quantum-secure one-way function.

Proof. Define the algorithm SampPuzz_C exactly as per the proof of Theorem 6.1. This algorithm is PPT by the same arguments given in the proof of Theorem 6.1. Also as per the proof of Theorem 6.1, it follows from the η -simulability of $(\mathcal{M}, \mathcal{F})$ that

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq \eta(|k|, n, t), \quad (6.40)$$

for all (k, n, t) . However, in this case, this implies that

$$d_{\text{TV}}(\text{SampPuzz}(1^\lambda), \text{SampPuzz}_C(1^\lambda)) \leq \mathbb{E}_{k \sim \text{KeyGen}(1^\lambda)} [\eta(|k|, n(\lambda), t(\lambda))] \quad (6.41)$$

$$\leq \sup_{k \in \mathbb{K}_\lambda} [\eta(|k|, n(\lambda), t(\lambda))] \quad (6.42)$$

$$= \text{negl}(\lambda), \quad (6.43)$$

where $t(\lambda) = t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)$ and the last inequality follows from our assumption on η . By pulling out the randomness in SampPuzz_C , one obtains an efficient classical deterministic algorithm which satisfies the claim of the theorem. The fact that the function family $\{F_\lambda\}$ is a quantum-secure OWF then follows immediately from Theorem 6.2. $\square$

7 Hamiltonian phase state assumptions imply one-way functions

In this section, we prove the following statement:

Theorem 7.1: (Search HPS assumption implies one-way functions) *If the Search HPS assumption (Assumption 2.20) is true, then one can construct a quantum-secure one-way function.*

As discussed earlier, the Decision HPS assumption implies the Search HPS assumption (constructively), and therefore a simple corollary of the above theorem is that one can also construct a quantum-secure OWF if the Decision HPS assumption is true.

Given the tools, of the previous sections, we note that the following logic would suffice for proving Theorem 7.1.

1. Under the Search HPS assumption, we know that we can construct (via Construction 2.21) a OWSG with output states in $\text{HPS}_{q,m,n}$ for some $q(n) = 2^{O(\text{poly}(n))}$ and $m(n) = O(\text{poly}(n))$. Note that this OWSG has a classical key generation algorithm KeyGen .

2. Given Theorem 6.1, we know that from any $1/3$ -simulable certifiable OWSG (with classical KeyGen), we can construct a quantum-secure OWF.
3. Therefore, to prove Theorem 7.1, it is sufficient to show that there exists a $1/3$ -simulable copy efficient "measure first, ask later" state certification protocol for $\text{HPS}_{q,m,n}$.

However, as discussed in Section 6.2, the OWF obtained via the above proof would be compiled from a distributional OWF. If we can prove that there exists an η -simulable *computationally* efficient "measure first, ask later" state certification protocol for $\text{HPS}_{q,m,n}$, where η is negligible in the sense of Theorem 6.4, then one can *directly* construct a OWF, without going via a distributional OWF. As such, instead of proving that there exists a $1/3$ -simulable copy efficient state certification protocol for $\text{HPS}_{q,m,n}$, we prove the stronger statement that there exists an η -simulable *computationally* efficient state certification protocol, for η sufficient for Theorem 6.4. More specifically:

1. In Section 7.1, we introduce the "measure first, ask later" state certification protocol from Ref. [HPS25] and prove that it is computationally efficient for any set of phase states with efficiently computable phase functions.
2. In Section 7.2, we prove that for any set of phase states with efficiently computable phase functions, the same state certification protocol is also η -simulable for a function η satisfying $\eta(|k|, n, t) = \text{negl}(\lambda)$ whenever $t = O(\text{poly}(n))$.
3. In Section 7.3, we then put the pieces together to prove Theorem 7.1, as per the high-level sketch above.

7.1 Computationally efficient certification of phase states

In this section, we present the recent "measure first, ask later" state certification protocol from Ref. [HPS25], and show that it is indeed computationally efficient for any set of phase states with efficiently computable phase functions (as a simple corollary of results already proven in Ref. [HPS25]). To present this certification protocol in the language of Definition 3.2, we start by defining an intermediate measurement protocol $\hat{\mathcal{M}}$, which acts on a single copy of an unknown state $|\psi\rangle$:

Algorithm 1 : $\hat{\mathcal{M}}$ – Single-copy measurement protocol implicit in Protocol 2 of Ref. [HPS25].

Input: n -qubit state $|\psi\rangle$

- 1: Choose a single index $i \in [n] := \{1, \dots, n\}$.
- 2: Measure all qubits *except* qubit i in the Z basis, and denote the outcome with $z \in \{0, 1\}^{n-1}$.
- 3: Choose $M \in \{X, Y, Z\}$ randomly, and measure qubit i in the M -basis. Denote the outcome as $b \in \{0, 1\}$.

Output: (i, z, M, b) .

To set conventions, we label the $+1$ Pauli eigenstate by $b = 0$ and the -1 eigenstate with $b = 1$. We stress that the above measurement protocol $\hat{\mathcal{M}}$ is *target* independent – i.e., it has no dependence on the target state $|\phi\rangle$ – and is therefore "measure first, ask later" in the sense of Section 3. Given this single-copy intermediate measurement protocol $\hat{\mathcal{M}}$, we then construct the measurement protocol $\mathcal{M}$ by applying algorithm $\hat{\mathcal{M}}$ independently on t copies of the unknown state $|\psi\rangle$ as follows:

Algorithm 2 : $\mathcal{M}$ – Full measurement protocol implicit in Protocol 2 in Ref. [HPS25]

Input: $|\psi\rangle^{\otimes t}$

- 1: Create an empty list $s = []$.
- 2: **for** $j \in [t] := \{1, \dots, t\}$, **do**
- 3: $s^{(j)} \leftarrow \hat{\mathcal{M}}(|\psi\rangle)$
- 4: Append $s^{(j)}$ to s

Output: $s = [s^{(1)}, \dots, s^{(t)}]$.

Now, for any fixed target state $|\phi\rangle$, we can define the classical post-processing function $F_{|\phi\rangle}$ required in Definition 3.2.

At a high level, $F_{|\phi\rangle}$ compares two objects extracted from each sample $s^{(j)}$: a single-qubit classical shadow $\sigma^{(j)}$, which is a linear (and target-independent) estimator of the true post-measurement state on qubit i , and a candidate single-qubit state $|\Psi^{(j)}\rangle$, reconstructed using the target state $|\phi\rangle$ via the amplitude oracle. The overlap between these two, averaged over t rounds, is then used to determine whether to accept or reject.

To make this precise, recall that each $s^{(j)}$ is in fact a tuple $s^{(j)} = (i, z, M, b) \in [n] \times \{0, 1\}^{n-1} \times \{X, Y, Z\} \times \{0, 1\}$. Furthermore, let's use the notation $z(j)$ for $j \in [n-1]$ to denote the j 'th element of $z \in \{0, 1\}^{n-1}$, and for any $i \in [n]$ and bit $b \in \{0, 1\}$, we define $z\|b_i \in \{0, 1\}^n$ as the bitstring defined from z via the insertion of bit b at index i – i.e., via

$$z\|b_i(j) = \begin{cases} z(j) & \text{if } j < i, \\ b & \text{if } j = i, \\ z(j-1) & \text{if } j > i. \end{cases} \quad (7.1)$$

With this in hand, we then define a single-qubit classical shadow $\sigma^{(j)}$ from $s^{(j)}$ as

$$s^{(j)} = (i, z, M, b) \rightarrow \sigma^{(j)} = \begin{cases} 3 H|b\rangle\langle b|H - \mathbb{I}_2, & M = X, \\ 3 SH|b\rangle\langle b|HS^\dagger - \mathbb{I}_2, & M = Y, \\ 3 |b\rangle\langle b| - \mathbb{I}_2, & M = Z, \end{cases} \quad (7.2)$$

and also the single qubit state $|\Psi^{(j)}\rangle$ from $s^{(j)}$ as:

$$s^{(j)} = (i, z, M, b) \rightarrow |\Psi^{(j)}\rangle = \frac{\alpha_0^{(j)}|0\rangle + \alpha_1^{(j)}|1\rangle}{\sqrt{|\alpha_0^{(j)}|^2 + |\alpha_1^{(j)}|^2}} \quad (7.3)$$

where $\alpha_b^{(j)} = \langle z\|b_i | \phi \rangle$ can be obtained from the amplitude oracle $O(|\phi\rangle)$. Finally, we then define the overlap

$$\omega^{(j)} = \text{Tr}(\sigma^{(j)} |\Psi^{(j)}\rangle\langle\Psi^{(j)}|). \quad (7.4)$$

With this, we can now define the function $F_{|\phi\rangle}$ formally in Algorithm 3:

Algorithm 3 : $F_{|\phi\rangle}$ implicit in Protocol 2 (with $r = 1$) in Ref. [HPS25]

Input: Amplitude oracle $O(|\phi\rangle)$, relaxation time τ , $\mathcal{M}$'s output $s = [s^{(1)}, \dots, s^{(t)}]$, and tolerance $\epsilon > 0$.

- 1: Create an empty list $\omega = []$.
- 2: **for** $j \in [t]$, **do**
- 3: Calculate the shadow $\sigma^{(j)}$ from $s^{(j)}$.
- 4: Using $O(|\phi\rangle)$ and $s^{(j)}$, calculate the state $|\Psi^{(j)}\rangle$.
- 5: Calculate the overlap $\omega^{(j)}$.
- 6: Append $\omega^{(j)}$ to ω .
- 7: Calculate the mean

$$\hat{\omega}[|\phi\rangle] := \frac{1}{t} \sum_{j=1}^t \omega^{(j)}$$

Output: If $\hat{\omega}[|\phi\rangle] \geq 1 - (3\epsilon)/(4\tau)$ output **Accept**, else output **Reject**.

Remark 7.2: (Efficiency of $F_{|\phi\rangle}$ for phase states) We emphasize that Algorithm 3 as stated requires amplitude-oracle access $O(|\phi\rangle)$ in Line 4. However, for any phase state $|\phi_f\rangle$ with efficiently computable phase function f , the oracle $O(|\phi_f\rangle)$ can itself be simulated in polynomial time, by directly evaluating $\langle x | \phi_f \rangle = 2^{-n/2} \exp(if(x))$.

Now, given the measurement protocol $(\mathcal{M}, F_{|\phi\rangle})$ for a *single* target state $|\phi\rangle$, we can easily construct the measurement protocol $(\mathcal{M}, \mathcal{F})$ for a set of states $\{|\phi_k\rangle \mid k \in \mathbb{K}\}$ (i.e., satisfying Definition 3.4) by defining

$$\mathcal{F}(k, \cdot, \cdot, \cdot) = \begin{cases} F_{|\phi_k\rangle}(\cdot, \cdot, \cdot) & \forall k \in \mathbb{K}, \\ \text{Reject} & \text{otherwise.} \end{cases} \quad (7.5)$$

For the case of phase states, we obtain the following as a straightforward corollary of Theorem 6 in Ref. [HPS25]:

Theorem 7.3: (Certifiability of phase states – Corollary of Theorem 6 in Ref. [HPS25]) *For any set of phase states $\text{PS}(\{f_k\}_{k \in \mathbb{K}}) = \{|\phi_k\rangle \mid k \in \mathbb{K}\}$, define $\mathcal{F}$ as per Eq. (7.5). For any state $|\phi_k\rangle$, denote by $\mathcal{O}(|\phi_k\rangle)$ the oracle which gives access to amplitudes of $|\phi_k\rangle$. Then, $(\mathcal{M}, \mathcal{F})$ with $\tau = n$ is a “measure first, ask later” state certification protocol for $\{|\phi_k\rangle\}$, which is computationally efficient from $\mathcal{O}$ -access, with copy complexity*

$$t(n, \epsilon, \delta, |\mathbb{K}|) = O\left(\frac{n^2}{\epsilon^2} \log\left(\frac{|\mathbb{K}|}{\delta}\right)\right). \quad (7.6)$$

Moreover, if the phase functions $\{f_{|\phi_k\rangle}\}_{k \in \mathbb{K}}$ are efficiently computable and membership in $\mathbb{K}$ can be efficiently decided, then $(\mathcal{M}, \mathcal{F})$ with $\tau = n$ is in fact a computationally efficient “measure first, ask later” state certification protocol for $\{|\phi_k\rangle\}$.

For completeness, we provide a proof of Theorem 7.3 in Appendix C.

7.2 Classical simulability of phase state certification

In the previous section, we established that the state certification protocol from Ref. [HPS25] is *computationally efficient* for any set of phase states with efficiently computable phase functions. In this section, we show that for any set of phase states with efficiently computable phase functions, this state certification protocol is also η -simulable, for a function $\eta(|k|, n, t)$ which is negligible with respect to n whenever t is at most a polynomial function of n .

Theorem 7.4: (negl-simulability of Ref. [HPS25]’s state certification protocol for phase states) *Let $\text{PS}(\{f_k\}_{k \in \mathbb{K}}) = \{|\phi_k\rangle \mid k \in \mathbb{K}\}$ be a set of phase states with efficiently computable phase functions $\{f_k\}_{k \in \mathbb{K}}$, and let $(\mathcal{M}, \mathcal{F})$ be the “measure first, ask later” state certification protocol from Theorem 7.3. Then, $(\mathcal{M}, \mathcal{F})$ is η -simulable for some function η satisfying $\eta(|k|, n, t) = \text{negl}(n)$ whenever $t = \text{poly}(n)$. More specifically, there exists a classical PPT algorithm $\mathcal{M}_C$ satisfying*

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq \eta(|k|, n, t) \quad (7.7)$$

for all $k \in \mathbb{K}$ and $n, t \in \mathbb{N}$, where η is such that whenever $t = O(\text{poly}(n))$, then $\eta(|k|, n, t) = \text{negl}(n)$.

Proof. Recall from Algorithm 2 that $\mathcal{M}$ works by repeatedly calling the *single-copy* measurement protocol $\hat{\mathcal{M}}$ defined in Algorithm 1. As such, we start by providing a classical randomized algorithm $\hat{\mathcal{M}}_C(k, n, l) \rightarrow s$ for approximately sampling within error $1/2^l$ from the distribution defined by $\hat{\mathcal{M}}(|\phi_k\rangle)$ for each n -qubit state $|\phi_k\rangle$ defined by $k \in \mathbb{K}$. Given this, we then define $\mathcal{M}_C$ by simply replacing the calls to $\hat{\mathcal{M}}$ (on $|\phi_k\rangle$) in Algorithm 2 with calls to $\hat{\mathcal{M}}_C$ (on input k).

To this end, we require some notation. In particular, given any bitstring $z \in \{0, 1\}^{n-1}$, index $i \in [n]$ and bit $b \in \{0, 1\}$, recall from Eq. (7.1) the definition of $z \| b_i \in \{0, 1\}^n$. Moreover, given any $k \in \mathbb{K}$, define

$$g_k(z, i) = f_k(z \| 1_i) - f_k(z \| 0_i), \quad (7.8)$$

where f_k is the phase function from Eq. (2.17). Using this, for any $k \in \mathbb{K}$, we define the distribution $\hat{\mathcal{M}}_C(k, n)$ as the distribution sampled from via the following algorithm, where we set the convention that $\text{Bern}(p)$ is the bit distribution satisfying $\Pr[b = 1] = p$:

Algorithm 4 Sampler for $\tilde{\mathcal{M}}_C(k, n)$

Input: $k \in \mathbb{K}, n \in \mathbb{N}$

- 1: Choose a single index $i \in [n]$ uniformly at random.
- 2: Sample $z \in \{0, 1\}^{n-1}$ from the uniform distribution.
- 3: Choose $M \in \{X, Y, Z\}$ randomly.
- 4: **if** $M = Z$, **then**
- 5: $b \leftarrow \text{Bern}(1/2)$
- 6: **else if** $M = X$, **then**
- 7: $b \leftarrow \text{Bern}\left[1 - \cos^2(g_k(z, i)/2)\right]$
- 8: **else if** $M = Y$, **then**
- 9: $b \leftarrow \text{Bern}\left[1 - \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right)\right]$

Output: (i, z, M, b) .

We now have the following lemma:

Lemma 7.5: (Exact classical simulation of the single-copy measurement distribution) *For all $k \in \mathbb{K}$, one has that*

$$\tilde{\mathcal{M}}_C(k, n) = \hat{\mathcal{M}}(|\phi_k\rangle), \quad (7.9)$$

where both $\tilde{\mathcal{M}}_C(k, n)$ and $\hat{\mathcal{M}}(|\phi_k\rangle)$ are understood as distributions.

Proof. Given that $|\phi_k\rangle$ is a phase state, we have that measuring any register of $|\phi_k\rangle$ in the computational basis yields a uniformly random bit. As such, the bitstring $z \in \{0, 1\}^{n-1}$ output by $\hat{\mathcal{M}}(|\phi_k\rangle)$ is uniformly random, as is the bitstring $z \in \{0, 1\}^{n-1}$ output by $\tilde{\mathcal{M}}_C(k, n)$. Now, note that the post-measurement state after obtaining i in Step 1, and $z \in \{0, 1\}^{n-1}$ in Step 2 of Algorithm $\tilde{\mathcal{M}}_C(k, n)$ is

$$|\phi_k(z)\rangle = \frac{1}{\sqrt{2}} (|0\rangle + e^{ig_k(z, i)}|1\rangle) \quad (7.10)$$

where as before $g_k(z, i) = f_k(z\|1_i) - f_k(z\|0_i)$ and f_k is the phase function from Eq. (2.17). Therefore, we have the following for the bit b output by Step 3 of $\tilde{\mathcal{M}}_C(k, n)$, conditioned on obtaining i and z in Steps 1 and 2:

1. If $M = Z$, then $\Pr[b = 0|z, i] = \Pr[b = 1|z, i] = 1/2$.
2. If $M = X$, then

$$\Pr[b|z, i] = \begin{cases} \cos^2(g_k(z, i)/2) & \text{if } b = 0, \\ 1 - \cos^2(g_k(z, i)/2) & \text{if } b = 1. \end{cases} \quad (7.11)$$

3. If $M = Y$, then

$$\Pr[b|z, i] = \begin{cases} \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right) & \text{if } b = 0, \\ 1 - \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right) & \text{if } b = 1. \end{cases} \quad (7.12)$$

But the above distributions are precisely the distributions sampled from in Step 4 of $\tilde{\mathcal{M}}_C(k, n)$, and both i and M are drawn identically in both algorithms. $\square$

Unfortunately, $\tilde{\mathcal{M}}_C(k, n)$ cannot be sampled from *exactly* in polynomial time due to the Bernoulli sampling steps in Lines 7 and 9 of Algorithm 4. However, using the efficient computability of the phase functions, together with standard sampling techniques, one can sample from a negligibly close approximation to the desired Bernoulli distributions in polynomial time. More specifically, we have the following lemma:

Lemma 7.6: (Efficient approximate Bernoulli sampling) *Assume the phase functions $\{f_k\}_{k \in \mathbb{K}}$ are efficiently computable. Then, there exist randomized algorithms $\mathcal{B}_1$ and $\mathcal{B}_2$, which on input $z \in \{0, 1\}^n$, $i \in [n]$, $k \in \mathbb{K}$ and l , run in time $O(\text{poly}(n, |k|, l))$ and satisfy*

$$d_{\text{TV}}(\mathcal{B}_1(z, i, k, l), \text{Bern}[1 - \cos^2(g_k(z, i)/2)]) \leq 2^{-l}, \quad (7.13)$$

$$d_{\text{TV}}\left(\mathcal{B}_2(z, i, k, l), \text{Bern}\left[1 - \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right)\right]\right) \leq 2^{-l}. \quad (7.14)$$

Proof. On input z, i, k, l , algorithm $\mathcal{B}_1$ does the following:

1. $\hat{f}_k(z\|1_i, l) \leftarrow \mathcal{A}(k, z\|1_i, l)$ and $\hat{f}_k(z\|0_i, l) \leftarrow \mathcal{A}(k, z\|0_i, l)$ where $\mathcal{A}$ is the algorithm implied via efficient computability of $\{f_k\}_{k \in \mathbb{K}}$.
2. Compute $\hat{g}_k(z, i, l) := \hat{f}_k(z\|1_i, l) - \hat{f}_k(z\|0_i, l)$.
3. Compute an l -bit dyadic rational $\hat{c}(z, i, k, l)$ satisfying

$$\left|\hat{c}(z, i, k, l) - \cos^2\left(\frac{\hat{g}_k(z, i, l)}{2}\right)\right| < \frac{1}{2^{l+1}}. \quad (7.15)$$

4. Output: sample $b \leftarrow \text{Bern}[1 - \hat{c}(z, i, k, l)]$.

Clearly, the output of the randomized algorithm $\mathcal{B}_1$ is that of a Bernoulli distribution $\mathcal{B}_1(z, i, k, l) = \text{Bern}(1 - \hat{c}(z, i, k, l))$ with parameter $1 - \hat{c}(z, i, k, l)$. To compute such a parameter, one (1st) reads k , (2nd) calls the corresponding $\mathcal{A}_k$ twice, and (3rd) uses efficiently computable operations (e.g., $\cos$ acting on a length- l binary fraction) to obtain the length- l binary fraction representation of $\hat{c}(z, i, k, l)$. Finally, having the length- l binary fraction representation of $\hat{c}$, one can sample from $\text{Bern}(1 - \hat{c})$ by flipping l fair coins. Therefore, $\mathcal{B}_1$ has runtime $\text{poly}(n, |k|, l)$.

Let us now show Eq. (7.13). We have

$$\begin{aligned} d_{\text{TV}}(\mathcal{B}_1(z, i, k, l), \text{Bern}[1 - \cos^2(g_k(z, i)/2)]) &= d_{\text{TV}}(\text{Bern}[1 - \hat{c}(z, i, k, l)], \text{Bern}[1 - \cos^2(g_k(z, i)/2)]) \\ &= \left|\hat{c}(z, i, k, l) - \cos^2\left(\frac{g_k(z, i)}{2}\right)\right|, \end{aligned} \quad (7.16)$$

and by the triangular inequality,

$$\left|\hat{c}(z, i, k, l) - \cos^2\left(\frac{g_k(z, i)}{2}\right)\right| \leq \left|\hat{c}(z, i, k, l) - \cos^2\left(\frac{\hat{g}_k(z, i, l)}{2}\right)\right| + \left|\cos^2\left(\frac{\hat{g}_k(z, i, l)}{2}\right) - \cos^2\left(\frac{g_k(z, i)}{2}\right)\right| \leq \frac{1}{2^l}, \quad (7.17)$$

where we have used Eq. (7.15) in the first term, the fact that $\cos^2$ has Lipschitz constant 1, and that $|\hat{g}_k(z, i, l) - g_k(z, i)| \leq 1/2^l$ by Eq. (2.14). That is, the second term is bounded as

$$\left|\cos^2\left(\frac{\hat{g}_k(z, i, l)}{2}\right) - \cos^2\left(\frac{g_k(z, i)}{2}\right)\right| \leq \frac{1}{2} |\hat{g}_k(z, i, l) - g_k(z, i)| \leq \frac{1}{2^{l+1}}. \quad (7.18)$$

The proof for $\mathcal{B}_2$ is essentially the same but replacing $\hat{c}$ with $\hat{d}(z, i, k, l)$ such that

$$\left|\hat{d}(z, i, k, l) - \cos^2\left(\frac{\hat{g}_k(z, i, l) - \pi/2}{2}\right)\right| < \frac{1}{2^{l+1}}. \quad (7.19)$$

Therefore, the runtime of $\mathcal{B}_2$ is again $\text{poly}(n, |k|, l)$, and its output is that of a Bernoulli distribution $\mathcal{B}_2(z, i, k, l) = \text{Bern}[1 - \hat{d}(z, i, k, l)]$ with parameter $1 - \hat{d}(z, i, k, l)$. Using the triangular inequality and same arguments above,

$$d_{\text{TV}}\left(\mathcal{B}_2(z, i, k, l), \text{Bern}\left[1 - \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right)\right]\right) = \left|\hat{d}(z, i, k, l) - \cos^2\left(\frac{g_k(z, i) - \pi/2}{2}\right)\right| \leq \frac{1}{2^l}, \quad (7.20)$$

and the lemma is proven. $\square$

With the above in hand, we can now straightforwardly define the algorithm $\hat{\mathcal{M}}_C(k, n, l)$ (see Algorithm 5) by replacing the exact Bernoulli sampling in $\tilde{\mathcal{M}}_C(k, n)$ with the approximate Bernoulli sampling via $\mathcal{B}_1$ and $\mathcal{B}_2$.

Algorithm 5 Sampler for $\hat{\mathcal{M}}_C(k, n, l)$

Input: $k \in \mathbb{K}, n, l \in \mathbb{N}$

- 1: Choose a single index $i \in [n]$ uniformly at random.
- 2: Sample $z \in \{0, 1\}^{n-1}$ from the uniform distribution.
- 3: Choose $M \in \{X, Y, Z\}$ randomly.
- 4: **if** $M = Z$, **then**
- 5: $b \leftarrow \text{Bern}(1/2)$
- 6: **else if** $M = X$, **then**
- 7: $b \leftarrow \mathcal{B}_1(z, i, k, l)$
- 8: **else if** $M = Y$, **then**
- 9: $b \leftarrow \mathcal{B}_2(z, i, k, l)$

Output: (i, z, M, b) .

In particular, we have the following lemma:

Lemma 7.7: (Efficient single-copy negligible error simulation) *For all $k \in \mathbb{K}$ and $l \in \mathbb{N}$, one has*

$$d_{\text{TV}}(\hat{\mathcal{M}}_C(k, n, l), \tilde{\mathcal{M}}_C(k, n)) \leq 2^{-l}. \quad (7.21)$$

Moreover, $\hat{\mathcal{M}}_C(k, n, l)$ runs in time $\text{poly}(n, l)$.

Proof. First, it is clear that the runtime of $\hat{\mathcal{M}}_C(k, n, l)$ is that of $\mathcal{B}_1$ and $\mathcal{B}_2$, that is, $\text{poly}(n, |k|, l)$. Moreover,

$$\begin{aligned} d_{\text{TV}}(\hat{\mathcal{M}}_C(k, n, l), \tilde{\mathcal{M}}_C(k, n)) &= \frac{1}{3} + \frac{1}{3} \frac{1}{n2^{n-1}} \sum_{i=1}^n \sum_{z \in \{0,1\}^{n-1}} d_{\text{TV}}(\mathcal{B}_1(z, i, k, l), \text{Bern}[1 - \cos^2(g_k(z, i)/2)]) \\ &\quad + \frac{1}{3} \frac{1}{n2^{n-1}} \sum_{i=1}^n \sum_{z \in \{0,1\}^{n-1}} d_{\text{TV}}(\mathcal{B}_2(z, i, k, l), \text{Bern}[1 - \cos^2(g_k(z, i)/2 - \pi/4)]) \\ &\leq \frac{2}{3} 2^{-l} \leq 2^{-l} \end{aligned} \quad (7.22)$$

where we have used the fact that both algorithms are the same when $M = Z$ for all i, z , and Eqs. (7.13) and (7.14) for the cases $M = X$ and $M = Y$, respectively. $\square$

Finally, we can put everything together and define the claimed algorithm $\mathcal{M}_C$ as follows:

Algorithm 6 : $\mathcal{M}_C(k, 1^n, 1^t)$ – PPT negl-approximate sampler for $\mathcal{M}(|\phi_k\rangle^{\otimes t})$

Input: $k, 1^n, 1^t$

- 1: Create an empty list $s = []$.
- 2: **for** $j \in [t] := \{1, \dots, t\}$, **do**
- 3: $s^{(j)} \leftarrow \hat{\mathcal{M}}_C(k, n, n)$
- 4: Append $s^{(j)}$ to s

Output: $s = [s^{(1)}, \dots, s^{(t)}]$.

By virtue of Lemma 7.5 and Lemma 7.7, we have that

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq t2^{-n}. \quad (7.23)$$

Therefore, for all $t = \text{poly}(n)$, we then immediately have

$$d_{\text{TV}}(\mathcal{M}_C(k, 1^n, 1^t), \mathcal{M}(|\phi_k\rangle^{\otimes t})) \leq \text{negl}(n) \quad (7.24)$$

as claimed. Moreover, $\mathcal{M}_C(k, 1^n, 1^t)$ just calls t times the algorithm $\hat{\mathcal{M}}_C(k, n, l)$ with $l = n$. Since the latter runs in time $\text{poly}(n, |k|, l)$, it follows that $\mathcal{M}_C(k, 1^n, 1^t)$ runs in time $\text{poly}(n, |k|, t)$, that is, polynomial in the size of its input. This concludes the proof of Theorem 7.4. $\square$

7.3 Proof of Theorem 7.1

Given Theorem 7.3 and Theorem 7.4, we can now straightforwardly prove Theorem 7.1.

Proof (Theorem 7.1). Set $n = \lambda$. Let the functions $m(\lambda) = \text{poly}(\lambda)$ and $q(\lambda) = 2^{O(\text{poly}(\lambda))}$ be as per the Search HPS assumption. Let $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ be the OWSG obtained from Construction 2.21 under the Search HPS assumption. Note that $\mathbb{K}_\lambda := \text{supp}(\text{KeyGen}(\lambda)) = \mathbb{K}_{q,m,\lambda}$. Let $(\mathcal{M}, \mathcal{F})$ be the state certification protocol from Theorem 7.3 for the set of phase states $\text{HPS}_{q,m,\lambda} = \{|\phi_k\rangle \mid k \in \mathbb{K}_\lambda\}$, with copy complexity

$$t(\lambda, \epsilon, \delta, |\mathbb{K}|) = O\left(\frac{\lambda^2}{\epsilon^2} \log\left(\frac{|\mathbb{K}|}{\delta}\right)\right). \quad (7.25)$$

We prove Theorem 7.1 by:

1. Showing that $(\mathcal{M}, \mathcal{F})$ is *efficiently computable*.
2. Showing that $(\mathcal{M}, \mathcal{F})$ is η -simulable for a function η satisfying

$$\sup_{k \in \mathbb{K}_\lambda} \left[\eta\left(|k|, n(\lambda), t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)\right) \right] = \text{negl}(\lambda). \quad (7.26)$$

Given both the above, it then follows from Theorem 6.4 that one can directly construct a quantum-secure OWF from the η -simulable efficiently certifiable OWSG $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$.

(1) Efficient computability: This follows immediately from Theorem 7.3, together with the observation in Remark 2.17 that the phase functions defining $\text{HPS}_{q,m,\lambda}$ are efficiently computable.

(2) η -simulability: From Theorem 7.4, it follows that $(\mathcal{M}, \mathcal{F})$ is η -simulable for η satisfying

$$\eta(|k|, \lambda, t) = \text{negl}(\lambda) \quad (7.27)$$

whenever $t = O(\text{poly}(\lambda))$. Now, using Eq. (7.25), the fact that $n = \lambda$ and the fact that $|\mathbb{K}_\lambda| = 2^{O(\text{poly}(\lambda))}$, we have

$$t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|) = O(\text{poly}(\lambda)), \quad (7.28)$$

and therefore,

$$\sup_{k \in \mathbb{K}_\lambda} \left[\eta\left(|k|, n(\lambda), t(n(\lambda), \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_\lambda|)\right) \right] = \text{negl}(\lambda). \quad (7.29)$$

$\square$

Remark 7.8: (Alternative state certification methods) We note that the protocol of [HPS25] is not the only computationally efficient and η -simulable state certification protocol available for (Hamiltonian) phase states. The recent work of [PCK+26] implicitly provides an alternative method that achieves *constant* copy complexity with respect to n —an improvement over the $O(n^2/\epsilon^2)$ scaling of Theorem 7.3—at the cost of requiring more elaborate, though still target-independent, measurements rather than single-qubit measurements alone. Using this state certification protocol in place of the one from [HPS25] would therefore also allow us to construct a OWF from the search HPS assumption, with different properties to the one we have obtained implicitly in our proof of Theorem 7.1. We note that more recent state certification protocols which work for *all* states using only single-qubit measurements [GHO25; LZ26] are *not* applicable to our setting, as they require target-state-dependent adaptive measurements—i.e., they are not “measure first, ask later” as required here.

Remark 7.9: (Explicit OWF) In the proof of Theorem 7.1, the OWF $\{F_\lambda\}$, whose existence has been proven, has been left implicit. In Appendix E, we provide an explicit specification of this OWF. We note here however that the domain of F_λ depends on the *copy-complexity* t of the state certification protocol $(\mathcal{M}, \mathcal{F})$ appearing above. In particular, this fact provides a motivation for the use of alternative state certification protocols with *constant* copy complexity, as discussed in Remark 7.8 above.

8 Towards concrete instantiations of efficiently verifiable one-way puzzles

The results of the previous sections give a general recipe for constructing EV-OWP from PRSG/OWSG together with “measure first, ask later” state certification protocols. Using this recipe, we see that we can construct an EV-OWP which provides a genuine instantiation of Quantumania from any family of states $\{|\phi_k\rangle\}_{k \in \mathbb{K}}$ satisfying the following three properties:

1. **PRSG/OWSG:** There exists a distribution over the family which yields either a standard PRSG or a OWSG (under a plausible assumption).
2. **Certifiability:** The family admits a computationally efficient “measure first, ask later” state certification protocol $(\mathcal{M}, \mathcal{F})$ in the sense of Definition 3.4.
3. **Independence from OWF:** The assumption under which the family is a PRSG or OWSG can be true independently of whether one-way functions exist. In particular:
 - (a) The assumption in Point 1 *cannot* be that one-way functions exist.
 - (b) The measurement protocol $(\mathcal{M}, \mathcal{F})$ should not be 1/3-simulable – i.e., the joint distribution over (k, s) induced by sampling a key k and running $\mathcal{M}$ on $|\phi_k\rangle^{\otimes t}$ should not be classically simulable to within constant total variation distance. If it is 1/3-simulable, then by Theorem 6.1, the EV-OWP built from the PRSG/OWSG together with the state certification protocol implies a one-way function.

We note that simply satisfying Properties 1 and 2 will yield an EV-OWP (via Corollary 5.3), however Property 3 is necessary for this EV-OWP to be genuinely in Microcrypt! In particular, our work highlights that if one is to instantiate Quantumania via this route, then one requires a family of states satisfying a delicate balance: It should be sufficiently complex to yield an OWSG (or a PRSG) without using OWFs, yet simple enough to admit a computationally efficient “measure first, ask later” state certification protocol, while at the same time remaining sufficiently intricate that the certification protocol itself cannot be efficiently classically simulated. In summary, they are families of states that are *hard enough to learn, structured enough to certify, yet not so structured that the certification protocol becomes classically simulable*. As such, the natural question raised by this work is the following:

Question 8.1: *Does there exist a family of states satisfying all three properties above?*

We stress that, prior to this work, Hamiltonian phase states were precisely such a candidate! However, we show here that they fail to satisfy Property 3. Given this, to the best of our knowledge, no known family of states currently satisfies all three properties simultaneously. Below, we discuss the shortcomings of plausible candidates, with a partial summary given in Table 1.

OWF phase states: The original PRSG construction of Ji, Liu and Song [JLS18] is obtained by considering a family of phase states with phases computed via a quantum-secure one-way function. As such, this family of states immediately cannot satisfy Property 3.

Hamiltonian phase states. As discussed in Remark 2.15, under either the Search (or Decision) HPS assumption, Hamiltonian phase states satisfy Property 1. Additionally, Theorem 7.3 (leveraging the state certification protocol of Ref. [HPS25]) shows that they satisfy Property 2. However, we have shown in Theorem 7.4 that they do not satisfy Property 3!

Low stabilizer rank states. Superpositions of few stabilizer states – i.e., states with low *stabilizer rank* – are

State family	PRSG/OWSG	Efficient measure-first certification	Independent of OWF
OWF phase states	Yes (assuming OWF)	Yes	No
Hamiltonian phase states	Assumed	Yes	No
Low-rank stabilizer states	Plausible*	Yes	No
Random brickwork circuits	Plausible	Unknown	Unknown

Table 1: A summary of the extent to which a variety of existing families of states satisfy the properties necessary for instantiating an EV-OWP which is genuinely in Microcrypt. Existing candidate state families populate different regions of this landscape, but no known family currently satisfies all of the desired properties. (*) Open problem listed as Question 2 in [AA23].

potentially plausible candidates for satisfying Property 1. In particular, despite recent work showing the learnability of states with bounded stabilizer *extent* [AD26], it remains an open question whether states of low stabilizer *rank* can be learned efficiently from copies (see Question 2 in Ref. [AA23]). Additionally, by exploiting the stabilizer structure of these states, one can show that classical shadows provides a computationally efficient “measure first, ask later” state certification protocol – i.e., that these states satisfy Property 2. Unfortunately, however, it appears that, as was the case for Hamiltonian phase states, the same structure that facilitates efficient certification *also* allows one to classically simulate the quantum part of the measurement process, and therefore, these states fail to satisfy Property 3. As such, states with low stabilizer rank provide a second example of a family of states that plausibly satisfies Properties 1 and 2, but not Property 3.

Random circuits. Output states of (log depth) random brickwork circuits are plausibly hard to learn or clone [FGS+25], giving good evidence for Property 1 (without assuming OWF in the construction). However, no computationally efficient “measure first, ask later” certification protocol is known for this family, and consequently, at present, they do not satisfy Property 2, and Property 3 cannot be assessed. We note that the lack of a computationally efficient “measure first, ask later” state certification protocol for these states is precisely why proposed QCCC digital signature schemes based on this class of states lack efficient classical certification [NLO+26].

PRSGs from Forrelation: Ref. [KQS+23] has constructed an oracle relative to which *t*-Forrelation states can be used to construct a *single-copy* PRSG and $P = NP$ – i.e., in this oracle world the single-copy PRSG does not imply OWF. Unfortunately, single-copy PRSG are only able to instantiate Nanocrypt, and to the best of our knowledge are too weak to construct OWPs and therefore do not suffice for our purpose. However, Ref. [KQS+23] has also outlined a path for constructing multi-copy PRSG, assuming a strong conjecture on properties of *t*-Forrelation states. In light of this, understanding the certifiability of such states, assuming the required property of *t*-Forrelation states provides an interesting route to better understanding whether such states could also provide an instantiation of EV-OWP in this oracle world.

Tensor network states: In one-dimension, quantum states with polynomial bond dimension are efficiently learnable [LMH+17], and therefore not suitable candidates for the construction of PRSG or OWSG. However, on arbitrary graphs this is no longer immediately the case, and such states may potentially be complex enough for the construction of PRSG or OWSG, while retaining sufficient structure for computationally efficient “measure first, ask later” state certification.

Remark 8.2: (On candidate ensembles for OWPs) In this section we have focused on candidate ensembles for the instantiation of EV-OWPs (and therefore Quantumania). However, if one is interested in constructing OWPs (and instantiating Countcrypt) the same list of criteria apply, but with only *copy* efficiency of the state certification protocol required (as opposed to *computational* efficiency). However, because *every* state ensemble admits a copy-efficient “measure first, ask later” state certification protocol (global Clifford shadows) the tools and techniques we develop here cannot offer more insight into the suitability of any particular state ensemble – i.e. tailored state certification protocols cannot add anything over classical shadows.

A Proof of Corollary 2.8

As mentioned in the main text, this proof proceeds identically to the proof of Theorem 2.7 (Claim D.1 in [KT25]), after establishing that from a successful quantum adversary for a distributional OWF (as per Definition 2.3) you can get the same guarantee you would get from a successful classical adversary (as per Definition 2.2).

To be more precise, let's assume that $\{f_\lambda\}_\lambda$ is *not* a quantum-secure distributional OWF. Under this assumption, one has that

$$\bar{F}_A(\lambda) > 1 - \frac{1}{\lambda^2} \quad (\text{A.1})$$

for infinitely many λ . If we use the notation $z \leftarrow \rho_{f_\lambda}^A(r)$ to denote z obtained by measuring $\rho_{f_\lambda}^A(r)$ in the computational basis, then it follows from the Fuchs-van de Graaf inequality and Jensen's inequality that

$$d_{\text{TV}} \left((r, f_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}, (z, f_\lambda(r))_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ z \leftarrow \rho_{f_\lambda}^A(r)}} \right) = \mathbb{E}_{r \leftarrow \{0,1\}^{m(\lambda)}} \left[d_{\text{TV}} \left(r' \leftarrow f_\lambda^{-1}(f_\lambda(r)), z \leftarrow \rho_{f_\lambda}^A(r) \right) \right] \quad (\text{A.2})$$

$$\leq \mathbb{E}_{r \leftarrow \{0,1\}^{m(\lambda)}} \left[\frac{1}{2} \|\sigma_{f_\lambda(r)} - \rho_{f_\lambda}^A\|_1 \right] \quad (\text{A.3})$$

$$\leq \mathbb{E}_{r \leftarrow \{0,1\}^{m(\lambda)}} \left[\sqrt{1 - F_{\text{sq}}(\sigma_{f_\lambda(r)}, \rho_{f_\lambda}^A)} \right] \quad (\text{A.4})$$

$$\leq \sqrt{1 - \bar{F}_A(\lambda)} \quad (\text{A.5})$$

$$< \frac{1}{\lambda}. \quad (\text{A.6})$$

Now, define $\tilde{\mathcal{A}}$ as the QPT algorithm which on input $(1^\lambda, y)$ does the following:

1. Obtain $\rho_y^A \leftarrow \mathcal{A}(1^\lambda, |y\rangle\langle y|)$.
2. Output $z \leftarrow \rho_y^A$.

Then, it follows from the above that

$$d_{\text{TV}} \left((r, f_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}, (z, f_\lambda(r))_{\substack{r \leftarrow \{0,1\}^{m(\lambda)} \\ z \leftarrow \tilde{\mathcal{A}}(1^\lambda, f_\lambda(r))}} \right) \leq \sqrt{1 - \bar{F}_A(\lambda)} < \frac{1}{\lambda}. \quad (\text{A.7})$$

The rest of the proof is now identical to the proof of Claim D.1 in [KT25], but we reproduce it for convenience. In particular, we use $\tilde{\mathcal{A}}$ to construct an adversary $\mathcal{B}$ for the OWP (Samp, Ver) from which $\{f_\lambda\}$ was constructed. To do this, recall that we have defined the existence of an efficient deterministic classical algorithm Samp_C satisfying

$$d_{\text{TV}}(\text{Samp}(1^\lambda), \text{Samp}_C(1^\lambda, r)_{r \leftarrow \{0,1\}^{m(\lambda)}}) \leq 1/3, \quad (\text{A.8})$$

where $\text{Samp}_C(1^\lambda, r) = (k_\lambda(r), s_\lambda(r))$. Using this, define the QPT puzzle adversary $\mathcal{B}$ on input $(1^\lambda, s)$ as follows:

1. $z \leftarrow \tilde{\mathcal{A}}(1^\lambda, s)$.
2. $(k_\lambda(z), s_\lambda(z)) \leftarrow \text{Samp}_C(1^\lambda, z)$
3. Output $k_\lambda(z)$.

Now, let $(k, s) \leftarrow \text{Samp}(1^\lambda)$. It follows from Eq. (A.7) that we have

$$d_{\text{TV}}((k_\lambda(r), s_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}, (\mathcal{B}(1^\lambda, s_\lambda(r)), s_\lambda(r))_{r \leftarrow \{0,1\}^{m(\lambda)}}) < \frac{1}{\lambda} \quad (\text{A.9})$$

and from the 1/3-correctness of Samp_C that

$$d_{\text{TV}}((k, s), (\mathcal{B}(1^\lambda, s), s)) \leq \frac{2}{3} + \frac{1}{\lambda}. \quad (\text{A.10})$$

Then, the correctness of the OWP implies

$$\Pr_{(k,s) \rightarrow \text{Samp}(1^\lambda)} [\text{Ver}(\mathcal{B}(1^\lambda, s), s) = 1] \geq \frac{1}{3} - \frac{1}{\lambda} - \text{negl}(\lambda) \quad (\text{A.11})$$

for infinitely many λ , which contradicts the assumed security of the OWP.

B Physical implementation of Hamiltonian phase states

Hamiltonian phase states are attractive as candidate cryptographic states not only because of their conjectured computational properties, but also because they can be prepared using simple commuting quantum circuits. In this appendix, we briefly discuss how the operations required to prepare such states can be realized in common experimental architectures. This physical accessibility makes the obstruction established in this work particularly relevant: it applies to a concrete and experimentally motivated family of quantum states rather than to a purely abstract cryptographic construction.

As discussed in Remark 2.15, Hamiltonian phase states can be prepared using *instantaneous quantum polynomial-time* (IQP) circuits. The matrix A specifies the structure and support of the Ising terms in the generating Hamiltonian. More precisely, for $A = (A_1, \dots, A_n) \in \mathbb{F}_2^n$, the corresponding operation is a multi-qubit Ising phase rotation of the form

$$U(\theta, A) := \exp \left(i\theta \bigotimes_{j=1}^n Z_j^{A_j} \right), \quad (\text{B.1})$$

where $\theta \in \mathbb{R}$ and the Hamming weight $|A|$ determines the number of qubits on which the operation acts non-trivially.

For $|A| = 2$, Eq. (B.1) is a two-qubit ZZ rotation. Such entangling phase operations are natural in Rydberg-atom platforms, where strong dipole–dipole interactions and the Rydberg-blockade mechanism enable programmable interactions between pairs of atoms [JCZ+00; SWM10; XXL+24; JP22; LKS+19; WGE+26]. Suitable pulse sequences can also generate multi-qubit controlled-phase operations within a blockade region, including three-qubit gates [JP22]. The precise range and connectivity of the available interactions depend on the geometry of the atomic array, the blockade radius, and the employed control protocol.

More general instances of Eq. (B.1) can be synthesized from two-qubit entangling gates and single-qubit rotations. Let $w = |A|$, and choose one of the w qubits as a parity accumulator. A sequence of $w - 1$ CNOT gates computes the parity of the participating qubits onto the accumulator. Applying a single-qubit Z rotation to the accumulator and subsequently reversing the CNOT sequence implements the desired w -qubit phase rotation. Thus, in the standard auxiliary-system-free construction, the operation requires $2(w - 1)$ CNOT gates, together with one single-qubit Z rotation. Equivalently, the construction can be expressed in terms of controlled- Z gates by conjugating the appropriate target qubits with Hadamard gates. The required parity ladders can be parallelized when the hardware connectivity permits, thereby reducing their circuit depth.

Rydberg-ion systems provide another possible architecture. Strong dipolar interactions between Rydberg-excited ions can facilitate multi-qubit phase operations, potentially allowing some of the required interactions to be implemented using short pulse sequences [MLL+08]. In superconducting-qubit platforms, single-qubit rotations and local two-qubit entangling gates are routinely available [KSB+20]. Higher-weight and geometrically non-local phase rotations must generally be compiled into the native gate set, with an overhead determined by the connectivity of the device.

Finally, the state certification protocol requires only single-qubit Pauli measurements after computational-basis measurements of the remaining qubits. These operations, as well as single-qubit readout, are standard capabilities of the architectures discussed above. We emphasize, however, that the role of this protocol in the present work is primarily conceptual: its measurement statistics on Hamiltonian phase states can be efficiently simulated classically, and it is precisely this simulability that allows us to prove that the Hamiltonian phase state assumptions imply one-way functions.

C Proof of Theorem 7.3

We start by stating and proving Theorem 6 from Ref. [HPS25], as the proof of Theorem 7.3 only requires a slight modification. To do this, consider the following Algorithm, which is Protocol 2 from Ref. [HPS25], expressed in the language of Section 3:

Algorithm 7 : Protocol 2 ($m = 1$) from Ref. [HPS25]

Input: Amplitude oracle $O(|\phi\rangle)$, t copies of unknown state $|\psi\rangle^{\otimes t}$ relaxation time τ and tolerance ϵ .

1: $s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})$

2: $o \leftarrow F_{|\phi\rangle}(s, \tau, \epsilon)$

Output: $o \in \{\text{Accept}, \text{Reject}\}$.

We then have the following:

Theorem C.1: (Theorem 6 from Ref. [HPS25]) *Given any target state $|\phi\rangle$ with relaxation time τ , error $\epsilon > 0$ and failure probability δ , running Algorithm 7 with*

$$t = 64 \frac{\tau^2}{\epsilon^2} \log \left(\frac{2}{\delta} \right) \quad (\text{C.1})$$

ensures that, with probability at least $1 - \delta$:

1. *If $|\langle \psi | \phi \rangle|^2 > 1 - \epsilon/(2\tau)$ then $o = \text{Accept}$,*
2. *If $|\langle \psi | \phi \rangle|^2 < 1 - \epsilon$ then $o = \text{Reject}$.*

Proof. We start by introducing some notation. In particular

1. We use the notation $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)$ to denote the *shadow overlap* between $|\psi\rangle$ and $|\phi\rangle$ defined via

$$\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle) = \langle \psi | L(|\phi\rangle) | \psi \rangle, \quad (\text{C.2})$$

where $L(|\phi\rangle)$ is the operator defined in Eq. (21) of Ref. [HPS25] (note that we use $|\phi\rangle$ to denote the target state here, but in [HPS25] the target state is denoted with $|\psi\rangle$).

2. We use the notation $\hat{\omega}_t(|\psi\rangle, |\phi\rangle)$ to denote the output of line 7 of Algorithm 3 for $F_{|\phi\rangle}$ when run on input $s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})$.

With the notation above, the proof of Theorem C.1 proceeds in two steps:

1. Show that given $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)$, one can decide whether $|\langle \psi | \phi \rangle|^2 > 1 - \epsilon/(2\tau)$ or $|\langle \psi | \phi \rangle|^2 < 1 - \epsilon$.
2. Bound the sample complexity necessary for $\hat{\omega}_t(|\psi\rangle, |\phi\rangle)$ to be sufficiently well concentrated around $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)$.

With this in mind, the first tool is the following:

Lemma C.2: (Adapted from Theorem 4 from Ref. [HPS25]) *Given a state $|\phi\rangle$ with relaxation time τ , one has the following:*

1. *If $|\langle \psi | \phi \rangle|^2 < 1 - \epsilon$ then $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle) < 1 - \epsilon/\tau$,*
2. *If $|\langle \psi | \phi \rangle|^2 > 1 - \epsilon/(2\tau)$ then $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle) \geq 1 - \epsilon/(2\tau)$,*

As such, if one wants to decide whether $|\langle \psi | \phi \rangle|^2 < 1 - \epsilon$ or $|\langle \psi | \phi \rangle|^2 > 1 - \epsilon/(2\tau)$ its sufficient to check whether $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle) < 1 - \epsilon/\tau$ or $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle) \geq 1 - \epsilon/(2\tau)$.

With this established, we now want to understand the concentration of $\hat{\omega}_t(|\psi\rangle, |\phi\rangle)$ around $\mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)$. To this end, note that Eq. (30) of Ref. [HPS25] can be written as

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [|\hat{\omega}_t(|\psi\rangle, |\phi\rangle) - \mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)| > \tilde{\epsilon}] \leq 2e^{-t\tilde{\epsilon}^2/2}, \quad (\text{C.3})$$

in the case $m = 1$. From the above, we then get the following lemma:

Lemma C.3: (Theorem 5 from Ref. [HPS25]) *For any fixed state $|\phi\rangle$, one has*

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [|\hat{\omega}_t(|\psi\rangle, |\phi\rangle) - \mathbb{E}[\omega](|\psi\rangle, |\phi\rangle)| > \tilde{\epsilon}] \leq \delta \quad (\text{C.4})$$

provided

$$t \geq 2 \frac{1}{\tilde{\epsilon}^2} \log \left(\frac{1}{\delta} \right). \quad (\text{C.5})$$

As a corollary of the above two lemmas, we see that if we set $\tilde{\epsilon} = \epsilon/(4\tau)$ then whenever

$$t \geq 2 \frac{1}{\tilde{\epsilon}^2} \log \left(\frac{1}{\delta} \right) = 64 \frac{\tau^2}{\epsilon^2} \log \left(\frac{1}{\delta} \right) \quad (\text{C.6})$$

one has, with probability at least $1 - \delta$, that:

1. If $|\langle \psi | \phi \rangle|^2 < 1 - \epsilon$ then $\hat{\omega}_t(|\psi\rangle, |\phi\rangle) < 1 - (3\epsilon)/(4\tau)$,
2. If $|\langle \psi | \phi \rangle|^2 > 1 - \epsilon/(2\tau)$ then $\hat{\omega}_t(|\psi\rangle, |\phi\rangle) \geq 1 - (3\epsilon)/(4\tau)$.

Now, the theorem statement follows from the fact that the output of $F_{|\phi\rangle}(s)$ on $s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})$ is determined precisely by whether $\hat{\omega}_t(|\psi\rangle, |\phi\rangle)$ is greater than or less than $1 - (3\epsilon)/(4\tau)$. $\square$

With this established, we note the following Lemma:

Lemma C.4: (Lemma 27 from Ref. [HPS25]) *For any phase state $|\phi\rangle$ the relaxation time is given by $\tau = n$.*

With the above, we are finally ready to prove Theorem 7.3;

Proof of Theorem 7.3. In order for the state certification protocol to work *simultaneously* for all states $|\phi_k\rangle$ one simply has to bound the sample complexity required for $\hat{\omega}_t(|\psi\rangle, |\phi_k\rangle)$ to be sufficiently well concentrated around $\mathbb{E}[\omega](|\psi\rangle, |\phi_k\rangle)$ for all $|\phi_k\rangle$ simultaneously. To do this, one takes a union bound using Eq. (C.3), and finds

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [\exists k : |\hat{\omega}_t(|\psi\rangle, |\phi_k\rangle) - \mathbb{E}[\omega](|\psi\rangle, |\phi_k\rangle)| > \tilde{\epsilon}] < |\mathbb{K}| 2e^{-t\tilde{\epsilon}^2/2}, \quad (\text{C.7})$$

from which it follows that

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} [\exists k : |\hat{\omega}_t(|\psi\rangle, |\phi_k\rangle) - \mathbb{E}[\omega](|\psi\rangle, |\phi_k\rangle)| > \tilde{\epsilon}] < \delta, \quad (\text{C.8})$$

provided that $t \geq \frac{2}{\tilde{\epsilon}^2} \log \left(\frac{2|\mathbb{K}|}{\delta} \right)$. Again setting $\tilde{\epsilon} = \epsilon/(4\tau)$ and using Lemma C.2 we then have

$$\Pr_{s \leftarrow \mathcal{M}(|\psi\rangle^{\otimes t})} \left[\forall k \in \mathbb{K} \begin{cases} \hat{\omega}_t(|\psi\rangle, |\phi_k\rangle) > 1 - (3\epsilon)/(4\tau) \text{ if } |\langle \psi | \phi_k \rangle|^2 > 1 - \epsilon/(2\tau) \\ \hat{\omega}_t(|\psi\rangle, |\phi_k\rangle) < 1 - (3\epsilon)/(4\tau) \text{ if } |\langle \psi | \phi_k \rangle|^2 < 1 - \epsilon \end{cases} \right] > 1 - \delta, \quad (\text{C.9})$$

provided that $t \geq 64 \frac{\tau^2}{\epsilon^2} \log \left(\frac{2|\mathbb{K}|}{\delta} \right)$. By using the fact that $\tau = n$ for phase states, the theorem statement follows. $\square$

D Proof of Lemma 6.3

Proof. We need to prove both correctness and security of $(\text{Samp}', \text{Ver})$. For both, we utilize the fact that for any function h with $h(k, s) \in [0, 1]$ for all (k, s) we have

$$\left| \mathbb{E}_{(k,s) \leftarrow \text{Samp}'(1^\lambda)} [h(k, s)] - \mathbb{E}_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [h(k, s)] \right| \leq \epsilon(\lambda). \quad (\text{D.1})$$

Correctness. Using the fact that $\text{Ver}(k, s) \in \{0, 1\}$ we have

$$\begin{aligned} \Pr_{(k,s) \leftarrow \text{Samp}'(1^\lambda)} [\text{Ver}(k, s) = 1] &= \mathbb{E}_{(k,s) \leftarrow \text{Samp}'(1^\lambda)} [\text{Ver}(k, s)] \\ &\geq \mathbb{E}_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(k, s)] - \epsilon(\lambda) \\ &= \Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [\text{Ver}(k, s) = 1] - \epsilon(\lambda) \\ &\geq 1 - \text{negl}(\lambda) - \epsilon(\lambda) \\ &\geq 1 - \text{negl}(\lambda). \end{aligned} \quad (\text{D.2})$$

Security: For any QPT adversary $\mathcal{A}(s) \rightarrow k'$ define the function

$$G_{\mathcal{A}}(k, s) = \Pr_{k' \leftarrow \mathcal{A}(s)} [\text{Ver}(k', s) = 1] \in [0, 1]. \quad (\text{D.3})$$

Then for any QPT adversary $\mathcal{A}$ we have

$$\begin{aligned} \Pr_{\substack{(k,s) \leftarrow \text{Samp}'(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}} [\text{Ver}(k', s) = 1] &= \mathbb{E}_{(k,s) \leftarrow \text{Samp}'(1^\lambda)} [G_{\mathcal{A}}(k, s)] \\ &\leq \mathbb{E}_{(k,s) \leftarrow \text{Samp}(1^\lambda)} [G_{\mathcal{A}}(k, s)] + \epsilon(\lambda) \\ &= \Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^\lambda) \\ k' \leftarrow \mathcal{A}(s)}} [\text{Ver}(k', s) = 1] + \epsilon(\lambda) \\ &\leq \text{negl}(\lambda) + \epsilon(\lambda) \\ &\leq \text{negl}(\lambda). \end{aligned} \quad (\text{D.4})$$

□

E The one-way function obtained from Search HPS

In the main text, the OWF $\{F_\lambda\}$ implied by Theorem 7.1 (under the Search HPS assumption) has been left implicit. In this appendix we make this OWF explicit. To this end:

1. We start from the Search HPS assumption, with functions q, m and distribution $\chi_{q,m,\lambda}$ over $\mathbb{K}_{q,m,\lambda}$.
2. Let $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ be the OWSG implied by the Search HPS assumption (i.e. as per Construction 2.21).
3. Let $(\mathcal{M}, \mathcal{F})$ be the η -simulable computationally efficient state certification protocol for $\{|\phi_k\rangle \mid k \in \mathbb{K}_{q,m,\lambda}\}$ from t copies given in Theorem 7.3.
4. Let $(\text{Samp}, \text{Ver})$ be the EV-OWP constructed from $((\text{KeyGen}, \text{StateGen}, \text{Ver}), (\mathcal{M}, \mathcal{F}))$ via Construction 5.2. In particular, recall that $\text{Samp}(1^\lambda) \rightarrow (k, s)$ with $k \in \mathbb{K}_{q,m,\lambda}$ and $s = [s^{(1)}, \dots, s^{(t(\lambda))}]$ where

$$(a) \quad t(\lambda) = t(\lambda, \epsilon = 1/8, \delta = 2^{-\lambda}, |\mathbb{K}_{q,m,\lambda}|)$$

$$(b) \quad s^{(j)} = (i^{(j)}, z^{(j)}, M^{(j)}, b^{(j)}) \text{ with } i^{(j)} \in [\lambda], z^{(j)} \in \{0, 1\}^{\lambda-1}, M^{(j)} \in \{X, Y, Z\} \text{ and } b^{(j)} \in \{0, 1\}$$

For convenience we define $S(\lambda) := \{[\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\} \times \{0, 1\}\}^{t(\lambda)}$ so that $s \in S(\lambda)$.

5. Let $\overline{\text{Samp}}_C(1^\lambda) \rightarrow (k, s) \in \mathbb{K}_{q,m,\lambda} \times S(\lambda)$ be the PPT algorithm for approximately sampling from $\text{Samp}(1^\lambda)$, which does the following on input 1^λ :

- (a) $k \leftarrow \chi_{q,m,\lambda}$

- (b) $s = [s^{(1)}, \dots, s^{(t(\lambda))}] \leftarrow \mathcal{M}_C(k, 1^\lambda, 1^{t(\lambda)})$ where $\mathcal{M}_C$ is given in Algorithm 6.

To define the OWF via the construction in Theorem 6.2 we now need to “pull out” the randomness from $\overline{\text{Samp}}_C$. To do this, we construct efficiently computable deterministic functions

1. $k_\lambda : \{0, 1\}^{p_1(\lambda)} \rightarrow \mathbb{K}_{q,m,\lambda}$
2. $s_\lambda : \{0, 1\}^{p_1(\lambda)} \times \{[\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\}\}^{t(\lambda)} \times \{\{0, 1\}^\lambda\}^{t(\lambda)} \rightarrow S(\lambda)$

which satisfy

$$(k_\lambda(r))_{r \leftarrow \{0,1\}^{p_1(\lambda)}} = \chi_{q,m,\lambda} \quad (\text{E.1})$$

and

$$\begin{aligned} \overline{\text{Samp}}_C(1^\lambda) = (k_\lambda(r_0), s_\lambda(r_0, r_1, r_2))_{\substack{r_0 \leftarrow \{0,1\}^{p_1(\lambda)} \\ r_1 \leftarrow \{[\lambda] \times \{0,1\}^{\lambda-1} \times \{X,Y,Z\}\}^{t(\lambda)} \\ r_2 \leftarrow \{\{0,1\}^\lambda\}^{t(\lambda)}}} \end{aligned} \quad (\text{E.2})$$

To this end, note that it follows from the assumption that $\chi_{q,m,\lambda}$ can be efficiently sampled from (with respect to λ) that there exists a polynomial p_1 and an efficiently computable deterministic function $\tilde{\chi}_{q,m,\lambda} : \{0, 1\}^{p_1(\lambda)} \rightarrow \mathbb{K}_{q,m,\lambda}$ satisfying

$$(\tilde{\chi}_{q,m,\lambda}(r))_{r \leftarrow \{0,1\}^{p_1(\lambda)}} = \chi_{q,m,\lambda}. \quad (\text{E.3})$$

Given this, we simply define $k_\lambda : \{0, 1\}^{p_1(\lambda)} \rightarrow \mathbb{K}_{q,m,\lambda}$ via $k_\lambda(r) = \tilde{\chi}_{q,m,\lambda}(r)$.

To define s_λ , we start by noting that $\mathcal{M}_C$ works by calling $\hat{\mathcal{M}}_C$ (Algorithm 5) iteratively, and therefore to pull the randomness out of $\mathcal{M}_C$ we have to pull the randomness out of $\hat{\mathcal{M}}_C$. To this, end let $\tilde{B}_j : \{0, 1\}^{\lambda-1} \times [\lambda] \times \mathbb{K}_{q,m,\lambda} \times \{0, 1\}^\lambda \rightarrow \{0, 1\}$ be the efficiently computable deterministic function satisfying

$$(\tilde{B}_j(z, i, k, r))_{r \leftarrow \{0,1\}^\lambda} = B_j(z, i, k, \lambda) \quad (\text{E.4})$$

where B_j is as defined in Lemma 7.6. With this, we then define

$$\tilde{s}_\lambda : \mathbb{K}_{q,m,\lambda} \times [\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\} \times \{0, 1\}^\lambda \rightarrow [\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\} \times \{0, 1\} \quad (\text{E.5})$$

via

$$\tilde{s}_\lambda(k, i, z, M, r) = \begin{cases} (i, z, M, r(1)) & \text{if } M = Z \\ (i, z, M, \tilde{B}_1(z, i, k, r)) & \text{if } M = X \\ (i, z, M, \tilde{B}_2(z, i, k, r)) & \text{if } M = Y \end{cases} \quad (\text{E.6})$$

where $r(1)$ denotes the first bit of the string r . With these definitions, it then follows from Eq. (E.4) and the definition of $\hat{\mathcal{M}}_C$ that

$$\hat{\mathcal{M}}_C(k, \lambda, \lambda) = (\tilde{s}_\lambda(k, i, z, M, r))_{\substack{(i,z,m) \leftarrow [\lambda] \times \{0,1\}^{\lambda-1} \times \{X,Y,Z\} \\ r \leftarrow \{0,1\}^\lambda}} \quad (\text{E.7})$$

In other words, $\tilde{s}_\lambda$ is the function obtained by pulling out the randomness from $\hat{\mathcal{M}}_C$. With this, for any

$$(r_0, r_1, r_2) \in \{0, 1\}^{p_1(\lambda)} \times \{[\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\}\}^{t(\lambda)} \times \{\{0, 1\}^\lambda\}^{t(\lambda)} \quad (\text{E.8})$$

with

$$r_1 = ((i^{(1)}, z^{(1)}, M^{(1)}), \dots, (i^{(t(\lambda))}, z^{(t(\lambda))}, M^{(t(\lambda))})) \quad (\text{E.9})$$

$$r_2 = (r^{(1)}, \dots, r^{(t(\lambda))}) \quad (\text{E.10})$$

we can finally define the function $s_\lambda : \{0, 1\}^{p_1(\lambda)} \times \{[\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\}\}^{t(\lambda)} \times \{\{0, 1\}^\lambda\}^{t(\lambda)} \rightarrow S(\lambda)$ via

$$s_\lambda(r_0, r_1, r_2) = \left(\tilde{s}_\lambda(k_\lambda(r_0), i^{(1)}, z^{(1)}, M^{(1)}, r^{(1)}), \dots, \tilde{s}_\lambda(k_\lambda(r_0), i^{(t(\lambda))}, z^{(t(\lambda))}, M^{(t(\lambda))}, r^{(t(\lambda))}) \right), \quad (\text{E.11})$$

which can be checked to satisfy Eq. (E.2). With this established, let's define

$$R(\lambda) = \{0, 1\}^{p_1(\lambda)} \times \{[\lambda] \times \{0, 1\}^{\lambda-1} \times \{X, Y, Z\}\}^{t(\lambda)} \times \{\{0, 1\}^\lambda\}^{t(\lambda)}. \quad (\text{E.12})$$

Using the construction of Theorem 6.2, the OWF implied by Theorem 7.1 is then $\{F_\lambda : R(\lambda) \rightarrow \{0, 1\}^*\}$ with

$$F_\lambda(r) = \begin{cases} (0, s_\lambda(r)) & r \in \text{Good}_\lambda \\ (1, r) & r \in \text{Bad}_\lambda \end{cases} \quad (\text{E.13})$$

where

$$\text{Good}_\lambda = \{r \mid \text{Ver}(k_\lambda(r), s_\lambda(r)) = F(k_\lambda(r), s_\lambda(r), 1/8, 2^{-\lambda}) = 1\}, \quad (\text{E.14})$$

$$\text{Bad}_\lambda = \{r \mid \text{Ver}(k_\lambda(r), s_\lambda(r)) = F(k_\lambda(r), s_\lambda(r), 1/8, 2^{-\lambda}) = 0\}. \quad (\text{E.15})$$

and F is the efficiently computable post-processing function of the state certification protocol, defined in Eq. (7.5). As mentioned in Remark 7.9, we note that the domain $R(\lambda)$ depends on the copy complexity $t(\lambda)$, and that the length of strings in the co-domain could be reduced by considering state certification protocols with improved copy-complexity.

References

- [AA23] Anurag Anshu and Srinivasan Arunachalam. *A survey on the complexity of learning quantum states*. 2023. arXiv: [2305.20069](https://arxiv.org/abs/2305.20069) [quant-ph]. URL: <https://arxiv.org/abs/2305.20069> (page 41).
- [AD26] Srinivasan Arunachalam and Arkopal Dutt. *Tomography of quantum states with bounded extent*. 2026. arXiv: [2606.07425](https://arxiv.org/abs/2606.07425). URL: <https://arxiv.org/abs/2606.07425> (page 41).
- [ADD+21] Myrto Arapinis, Mahshid Delavar, Mina Doosti, and Elham Kashefi. “Quantum Physical Unclonable Functions: Possibilities and Impossibilities”. In: *Quantum* 5 (2021), p. 475. ISSN: 2521-327X. DOI: [10.22331/q-2021-06-15-475](https://doi.org/10.22331/q-2021-06-15-475). URL: <http://dx.doi.org/10.22331/q-2021-06-15-475> (page 9).
- [AGQ+22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. “Pseudorandom (function-like) quantum state generators: New definitions and applications”. In: *Theory of Cryptography*. Ed. by Eike Kiltz and Vinod Vaikuntanathan. Cham: Springer Nature Switzerland, 2022, pp. 237–265 (pages 3, 17).
- [ALY23] Prabhanjan Ananth, Yao-Ting Lin, and Henry Yuen. *Pseudorandom strings from pseudorandom quantum states*. 2023. arXiv: [2306.05613](https://arxiv.org/abs/2306.05613) (pages 3, 8, 11, 17).
- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. “Cryptography from pseudorandom quantum states”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 208–236 (page 3).
- [BCN24] John Bostanci, Boyang Chen, and Barak Nehoran. *Oracle separation between quantum commitments and quantum one-wayness*. Cryptology ePrint Archive, Paper 2024/1568. 2024. URL: <https://eprint.iacr.org/2024/1568> (page 11).
- [BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. “On the computational hardness needed for quantum cryptography”. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Ed. by Yael Tauman Kalai. Vol. 251. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 24:1–24:21. DOI: [10.4230/LIPIcs.ITCS.2023.24](https://doi.org/10.4230/LIPIcs.ITCS.2023.24). arXiv: [2209.04101](https://arxiv.org/abs/2209.04101) (pages 3, 8).
- [BDL+24] Simon Becker, Nilanjana Datta, Ludovico Lami, and Cambyse Rouze. “Classical shadow tomography for continuous variables quantum systems”. In: *IEEE Transactions on Information Theory* 70.5 (May 2024), pp. 3427–3452. ISSN: 1557-9654. DOI: [10.1109/TIT.2024.3357972](https://doi.org/10.1109/TIT.2024.3357972). URL: <http://dx.doi.org/10.1109/TIT.2024.3357972> (page 10).

- [BHH+24] Christian Bertoni, Jonas Haferkamp, Marcel Hinsche, Marios Ioannou, Jens Eisert, and Hakop Pashayan. “Shallow shadows: Expectation estimation using low-depth random Clifford circuits”. In: *Phys. Rev. Lett.* 133 (2024), p. 020602. DOI: [10.1103/PhysRevLett.133.020602](https://doi.org/10.1103/PhysRevLett.133.020602) (page 10).
- [BHH+25] John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. *Efficient quantum pseudorandomness from Hamiltonian phase states*. 2025. arXiv: [2410.08073](https://arxiv.org/abs/2410.08073) (pages 1, 4, 5, 9, 18, 19).
- [BJS11] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy”. In: *Proc. Roy. Soc. A* 467.2126 (2011), pp. 459–472 (page 19).
- [BMM+25] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. “A new world in the depths of Microcrypt: Separating OWSGs and quantum money from QEFID”. In: *Advances in Cryptology – EUROCRYPT 2025*. Springer Nature Switzerland, 2025, pp. 23–52. ISBN: 9783031910982. DOI: [10.1007/978-3-031-91098-2_2](https://doi.org/10.1007/978-3-031-91098-2_2). URL: http://dx.doi.org/10.1007/978-3-031-91098-2_2 (page 11).
- [BOW17] Costin Bădescu, Ryan O’Donnell, and John Wright. *Quantum state certification*. 2017. arXiv: [1708.06002](https://arxiv.org/abs/1708.06002) [quant-ph]. URL: <https://arxiv.org/abs/1708.06002> (page 20).
- [BS19] Zvika Brakerski and Omri Shmueli. *(Pseudo) random quantum states with binary phase*. 2019. arXiv: [1906.10611](https://arxiv.org/abs/1906.10611) (page 9).
- [BS20] Zvika Brakerski and Omri Shmueli. *Scalable pseudorandom quantum states*. 2020. arXiv: [2004.01976](https://arxiv.org/abs/2004.01976) (pages 8, 17).
- [CCH+22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. “Exponential separations between learning with and without quantum memory”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 574–585 (page 11).
- [CDX+24] Chi-Fang Chen, Jordan Docter, Michelle Xu, Adam Boulund, Fernando G. S. L. Brandao, and Patrick Hayden. “Efficient unitary designs from random sums and permutations”. In: *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2024, pp. 476–484. DOI: [10.1109/focs61266.2024.00037](https://doi.org/10.1109/focs61266.2024.00037) (page 9).
- [CEF26] Jonathan Conrad, Jens Eisert, and Steven T. Flammia. “Chasing shadows with Gottesman-Kitaev-Preskill codes”. In: *Quantum* 10 (Jan. 2026), p. 1973. ISSN: 2521-327X. DOI: [10.22331/q-2026-01-19-1973](https://doi.org/10.22331/q-2026-01-19-1973). URL: <http://dx.doi.org/10.22331/q-2026-01-19-1973> (page 10).
- [CGG+25] Bruno Cavalar, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. “On the computational hardness of quantum one-wayness”. In: *Quantum* 9 (2025), p. 1679. DOI: [10.22331/q-2025-03-27-1679](https://doi.org/10.22331/q-2025-03-27-1679) (pages 17, 23).
- [CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. *On central primitives for quantum cryptography with classical communication*. Cryptology ePrint Archive, Paper 2024/356. 2024 (pages 3, 8, 11, 17).
- [CL26] Alexandru Cojocaru and Laura Lewis. *Equivalence Between Average-Case Hardness of Learning and Cryptography for Mixed Quantum States*. 2026. arXiv: [2608.14331](https://arxiv.org/abs/2608.14331) [quant-ph]. URL: <https://arxiv.org/abs/2608.14331> (page 9).
- [CLS+26] Andrea Coladangelo, Jerry Li, Joseph Sloate, and Ellen Wu. *The Power of Two Bases: Robust and copy-optimal certification of nearly all quantum states with few-qubit measurements*. 2026. arXiv: [2602.11616](https://arxiv.org/abs/2602.11616) [quant-ph]. URL: <https://arxiv.org/abs/2602.11616> (pages 9, 10, 21).
- [CLS26] Andrea Coladangelo, Jerry Li, and Joseph Sloate. *Robust quantum state certification and uncertainty principles for total influence*. 2026. arXiv: [2607.27184](https://arxiv.org/abs/2607.27184). URL: <https://arxiv.org/abs/2607.27184> (pages 9, 10).
- [EFH+23] Andreas Elben, Steven T. Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoit Vermersch, and Peter Zoller. “The randomized measurement toolbox”. In: *Nature Rev. Phys.* 5.1 (2023), pp. 9–24 (pages 10, 22).
- [FGS+25] Bill Fefferman, Soumik Ghosh, Makrand Sinha, and Henry Yuen. *The hardness of learning quantum circuits and its cryptographic applications*. 2025. arXiv: [2504.15343](https://arxiv.org/abs/2504.15343) (pages 9, 12, 41).
- [GHO25] Meghal Gupta, William He, and Ryan O’Donnell. *Few single-qubit measurements suffice to certify any quantum state*. 2025. arXiv: [2506.11355](https://arxiv.org/abs/2506.11355) [quant-ph]. URL: <https://arxiv.org/abs/2506.11355> (pages 7, 9, 10, 21, 39).

- [GMM+25] Eli Goldin, Tomoyuki Morimae, Saachi Mutreja, and Takashi Yamakawa. *CountCrypt: Quantum Cryptography between QCMA and PP*. 2025. arXiv: [2410.14792](https://arxiv.org/abs/2410.14792) [quant-ph]. URL: <https://arxiv.org/abs/2410.14792> (pages 3, 4, 8, 16).
- [Gol+01] Oded Goldreich et al. *Foundations of cryptography*. Vol. 1. 3. Cambridge university press Cambridge, 2001 (pages 7, 14).
- [GPS24] Daniel Grier, Hakop Pashayan, and Luke Schaeffer. “Sample-optimal classical shadows for pure states”. In: *Quantum* 8 (2024), p. 1373. DOI: [10.22331/q-2024-06-17-1373](https://doi.org/10.22331/q-2024-06-17-1373) (page 10).
- [GY25] Alex B. Grilo and Álvaro Yáñez. “Quantum pseudoresources imply cryptography”. In: (2025). arXiv: [2504.15025](https://arxiv.org/abs/2504.15025) (page 3).
- [HE23] Dominik Hangleiter and Jens Eisert. “Computational advantage of quantum random sampling”. In: *Rev. Mod. Phys.* 95.3 (2023). DOI: [10.1103/revmodphys.95.035001](https://doi.org/10.1103/revmodphys.95.035001) (page 9).
- [HH24] Taiga Hiroka and Min-Hsiu Hsieh. *Computational complexity of learning efficiently generatable pure states*. 2024. arXiv: [2410.04373](https://arxiv.org/abs/2410.04373) (page 9).
- [HHM25] Taiga Hiroka, Min-Hsiu Hsieh, and Tomoyuki Morimae. *Hardness of quantum distribution learning and quantum cryptography*. 2025. arXiv: [2507.01292](https://arxiv.org/abs/2507.01292) (page 9).
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Predicting many properties of a quantum system from very few measurements”. In: *Nature Physics* 16.10 (2020), pp. 1050–1057. DOI: [10.1038/s41567-020-0932-7](https://doi.org/10.1038/s41567-020-0932-7) (pages 6, 7, 10, 22).
- [HPS25] Hsin-Yuan Huang, John Preskill, and Mehdi Soleimanifar. “Certifying almost all quantum states with few single-qubit measurements”. In: *Nat. Phys.* 21 (2025). Also in Proc. FOCS 2024, p. 1834. DOI: [10.1038/s41567-025-03025-1](https://doi.org/10.1038/s41567-025-03025-1). arXiv: [2404.07281](https://arxiv.org/abs/2404.07281) (pages 4, 5, 7–10, 12, 21, 22, 33–35, 39, 40, 44, 45).
- [HW23] Jonas Helsen and Michael Walter. “Thrifty shadow estimation: Reusing quantum circuits and bounding tails”. In: *Phys. Rev. Lett.* 131.24 (2023), p. 240602. DOI: [10.1103/PhysRevLett.131.240602](https://doi.org/10.1103/PhysRevLett.131.240602) (page 10).
- [IL89] Russell Impagliazzo and Michael Luby. “One-way functions are essential for complexity based cryptography”. In: *30th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society. 1989, pp. 230–235 (pages 7, 14, 15).
- [Imp92] Russell Impagliazzo. “Pseudo-random generators for cryptography and for randomized algorithms”. PhD thesis. PhD thesis, University of California, Berkeley, 1992. [http://cseweb.ucsd.edu/...](http://cseweb.ucsd.edu/~imp/papers/thesis.pdf), 1992 (page 15).
- [JCZ+00] Dieter Jaksch, J. Ignacio Cirac, Peter Zoller, Steven L. Rolston, Robin Côté, and Mischa D. Lukin. “Fast quantum gates for neutral atoms”. In: *Phys. Rev. Lett.* 85 (10 2000), pp. 2208–2211. DOI: [10.1103/PhysRevLett.85.2208](https://doi.org/10.1103/PhysRevLett.85.2208) (page 43).
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. “Pseudorandom quantum states”. In: *Advances in Cryptology – CRYPTO 2018*. Ed. by Hovav Shacham and Alexandra Boldyreva. Cham: Springer International Publishing, 2018, pp. 126–152 (pages 3, 8, 17, 40).
- [JP22] Sven Jandura and Guido Pupillo. “Time-optimal two- and three-qubit gates for Rydberg atoms”. In: *Quantum* 6 (2022), p. 712. DOI: [10.22331/q-2022-05-13-712](https://doi.org/10.22331/q-2022-05-13-712) (page 43).
- [KG22] Dax Enshan Koh and Sabee Grewal. “Classical shadows with noise”. In: *Quantum* 6 (2022), p. 776. DOI: [10.22331/q-2022-08-16-776](https://doi.org/10.22331/q-2022-08-16-776) (page 10).
- [KK07] Elham Kashefi and Iordanis Kerenidis. *Statistical Zero Knowledge and quantum one-way functions*. 2007. arXiv: [quant-ph/0511266](https://arxiv.org/abs/quant-ph/0511266) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/0511266> (pages 7, 14, 15).
- [KLP+25] Andrey Boris Khesin, Jonathan Z. Lu, Alexander Poremba, Akshar Ramkumar, and Vinod Vaikuntanathan. *Average-case complexity of quantum stabilizer decoding*. 2025. arXiv: [2509.20697](https://arxiv.org/abs/2509.20697) (page 9).
- [KNY23] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. *Publicly Verifiable Deletion from Minimal Assumptions*. Cryptology ePrint Archive, Paper 2023/538. 2023. URL: <https://eprint.iacr.org/2023/538> (page 8).
- [KQS+23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. “Quantum Cryptography in Algorithmica”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. STOC ’23. ACM, 2023, pp. 1589–1602. DOI: [10.1145/3564246.3585225](https://doi.org/10.1145/3564246.3585225) (pages 3, 11, 41).

- [KQT25] William Kretschmer, Luowen Qian, and Avishay Tal. “Quantum-computable one-way functions without one-way functions”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC ’25. ACM, 2025, pp. 189–200. DOI: [10.1145/3717823.3718144](https://doi.org/10.1145/3717823.3718144) (pages 3, 8, 11, 16, 17).
- [Kre21] William Kretschmer. “Quantum Pseudorandomness and Classical Complexity”. In: *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*. Ed. by Min-Hsiu Hsieh. Vol. 197. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 2:1–2:20. ISBN: 978-3-95977-198-6. DOI: [10.4230/LIPIcs.TQC.2021.2](https://doi.org/10.4230/LIPIcs.TQC.2021.2). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2021.2> (pages 3, 11, 17).
- [KSB+20] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I-Jan Wang, Simon Gustavsson, and William D. Oliver. “Superconducting qubits: Current state of play”. In: *Ann. Rev. Cond. Matt. Phys.* 11 (2020), pp. 369–395. DOI: [10.1146/annurev-conmatphys-031119-050605](https://doi.org/10.1146/annurev-conmatphys-031119-050605) (page 43).
- [KT24] Dakshita Khurana and Kabir Tomer. *Commitments from quantum one-wayness*. 2024. arXiv: [2310.11526](https://arxiv.org/abs/2310.11526) (pages 3–8, 10, 15, 17, 23, 24).
- [KT25] Dakshita Khurana and Kabir Tomer. “Founding quantum cryptography on quantum advantage, or, towards cryptography from #P hardness”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025, pp. 178–188 (pages 3, 4, 7, 9, 15–17, 42).
- [LKS+19] Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T. Wang, Sepehr Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletić, Hannes Pichler, and Mikhail D. Lukin. “Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms”. In: *Phys. Rev. Lett.* 123 (17 2019), p. 170503. DOI: [10.1103/PhysRevLett.123.170503](https://doi.org/10.1103/PhysRevLett.123.170503). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.170503> (page 43).
- [LMH+17] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos. “Efficient tomography of a quantum many-body system”. In: *Nature Physics* 13.12 (2017), pp. 1158–1162. ISSN: 1745-2481. DOI: [10.1038/nphys4244](https://doi.org/10.1038/nphys4244). URL: <http://dx.doi.org/10.1038/nphys4244> (page 41).
- [LPQ+26] Jonathan Z. Lu, Alexander Poremba, Yihui Quek, and Akshar Ramkumar. *Post-quantum cryptography from quantum stabilizer decoding*. 2026. arXiv: [2603.19110](https://arxiv.org/abs/2603.19110) (pages 4, 5, 9).
- [LZ26] Yunting Li and Huangjun Zhu. “Universal and efficient quantum state verification via Schmidt decomposition and mutually unbiased bases”. In: *Quantum* 10 (Mar. 2026), p. 2011. ISSN: 2521-327X. DOI: [10.22331/q-2026-03-04-2011](https://doi.org/10.22331/q-2026-03-04-2011). URL: <http://dx.doi.org/10.22331/q-2026-03-04-2011> (pages 7, 39).
- [MH25] Fermi Ma and Hsin-Yuan Huang. *How to construct random unitaries*. 2025. arXiv: [2410.10116](https://arxiv.org/abs/2410.10116) (page 8).
- [MLL+08] Markus Mueller, Lin-Mei Liang, Igor Lesanovsky, and Peter Zoller. “Trapped Rydberg ions: From spin chains to fast quantum gates”. In: *New J. Phys.* 10 (2008), p. 093009. DOI: [10.1088/1367-2630/10/9/093009](https://doi.org/10.1088/1367-2630/10/9/093009) (page 43).
- [MPS+24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. *Simple constructions of linear-depth t -designs and pseudorandom unitaries*. 2024. arXiv: [2404.12647](https://arxiv.org/abs/2404.12647) (page 9).
- [MY22] Tomoyuki Morimae and Takashi Yamakawa. “Quantum commitments and signatures without one-way functions”. In: *Advances in Cryptology – CRYPTO 2022*. Springer Nature Switzerland, 2022, pp. 269–295. DOI: [10.1007/978-3-031-15802-5_10](https://doi.org/10.1007/978-3-031-15802-5_10) (pages 3, 8, 16).
- [MY24] Tomoyuki Morimae and Takashi Yamakawa. *One-wayness in quantum cryptography*. 2024. arXiv: [2210.03394](https://arxiv.org/abs/2210.03394) (pages 3, 17).
- [NLO+26] Pradeep Niroula, Minzhao Liu, Sivaprasad Omanakuttan, David Amaro, Shouvanik Chakrabarti, Soumik Ghosh, Zichang He, Yuwei Jin, Fatih Kaleoglu, Steven Kordonowy, Rohan Kumar, Michael A. Perlin, Akshay Seshadri, Matthew Steinberg, Joseph Sullivan, Jacob Watkins, Henry Yuen, and Ruslan Shaydulin. *Digital signatures with classical shadows on near-term quantum computers*. 2026. arXiv: [2602.04859](https://arxiv.org/abs/2602.04859) (pages 12, 41).
- [PCK+26] Guedong Park, Jaekwon Chang, Yosep Kim, Yong Siah Teo, and Hyunseok Jeong. *Sample- and hardware-efficient fidelity estimation by stripping phase-dominated magic*. 2026. arXiv: [2602.09710](https://arxiv.org/abs/2602.09710) (pages 10, 12, 39).

- [PQS25] Alexander Poremba, Yihui Quek, and Peter Shor. *The learning stabilizers with noise problem*. 2025. arXiv: [2410.18953](#) (pages [4](#), [5](#), [9](#)).
- [Reg09] Oded Regev. “On lattices, learning with errors, random linear codes, and cryptography”. In: *J. ACM* 56 (2009). doi: [10.1145/1568318.1568324](#) (pages [9](#), [12](#)).
- [RS19] Roy Radian and Or Sattath. “Semi-Quantum Money”. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. AFT ’19. ACM, Oct. 2019, pp. 132–146. doi: [10.1145/3318041.3355462](#). URL: <http://dx.doi.org/10.1145/3318041.3355462> (pages [7](#), [14](#), [15](#)).
- [Sat24] Or Sattath. *MicroCrypt Zoo: An interactive visualization of quantum cryptographic primitives*. <https://sattath.github.io/microcrypt-zoo/>. Accessed: 2026-04-04. 2024 (pages [3](#), [4](#), [8](#)).
- [SB09] Dan Shepherd and Michael J. Bremner. “Temporally unstructured quantum computation”. In: *Proc. Roy. Soc. A* 465.2105 (2009), pp. 1413–1439 (page [19](#)).
- [SWM10] Mark Saffman, Thad G. Walker, and Klaus Mølmer. “Quantum information with Rydberg atoms”. In: *Rev. Mod. Phys.* 82 (3 2010), pp. 2313–2363. doi: [10.1103/RevModPhys.82.2313](#) (page [43](#)).
- [WGE+26] Clara Wassner, Tommaso Guaita, Jens Eisert, and Jose Carrasco. “Holonomic quantum computation: a scalable adiabatic architecture”. In: *Quantum* 10 (2026), p. 2080. doi: [10.22331/q-2026-04-23-2080](#) (page [43](#)).
- [WHL+23] Kianna Wan, William J. Huggins, Joonho Lee, and Ryan Babbush. “Matchgate shadows for fermionic quantum simulation”. In: *Commun. Math. Phys.* 404 (2023), pp. 629–700. doi: [10.1007/s00220-023-04844-0](#). arXiv: [2207.13723](#) (page [10](#)).
- [XXL+24] Ming Xue, Shijie Xu, Xinwei Li, and Xiangliang Li. “High-fidelity and robust controlled-Z gates implemented with Rydberg atoms via echoing rapid adiabatic passage”. In: *Phys. Rev. A* 110 (3 2024), p. 032619. doi: [10.1103/PhysRevA.110.032619](#) (page [43](#)).
- [ZRM21] Andrew Zhao, Nicholas C. Rubin, and Akimasa Miyake. “Fermionic partial tomography via classical shadows”. In: *Phys. Rev. Lett.* 127.11 (2021), p. 110504. doi: [10.1103/PhysRevLett.127.110504](#) (page [10](#)).