

A New World in the Depths of Microcrypt: Separating OWSGs and Quantum Money from QEFID

Amit Behera¹, Giulio Malavolta², Tomoyuki Morimae³, Tamer Mour⁴, Takashi Yamakawa^{5,3}

¹Department of Computer Science, Ben-Gurion University of the Negev, Beersheba, Israel

behera@post.bgu.ac.il

²Bocconi University, Milan, Italy

giulio.malavolta@hotmail.it

³Yukawa Institute for Theoretical Physics, Kyoto University, Kyoto, Japan

tomoyuki.morimae@yukawa.kyoto-u.ac.jp

⁴Bocconi University, Milan, Italy

tamer.mour@unibocconi.it

⁵NTT Social Informatics Laboratories, Tokyo, Japan

takashi.yamakawa@ntt.com

Abstract

While in classical cryptography one-way functions (OWFs) are widely regarded as the “minimal assumption”, the situation in quantum cryptography is less clear. Recent works have put forward two concurrent candidates for the minimal assumption in quantum cryptography: One-way state generators (OWSGs), postulating the existence of a hard *search* problem with an efficient verification algorithm, and EFI pairs, postulating the existence of a hard *distinguishing* problem. Two recent papers [Khurana and Tomer STOC’24; Batra and Jain FOCS’24] showed that OWSGs imply EFI pairs, but the reverse direction remained open.

In this work, we give strong evidence that the opposite direction does not hold: We show that there is a quantum unitary oracle relative to which EFI pairs exist but OWSGs do not. In fact, we show a slightly stronger statement that holds also for EFI pairs that output classical bits (QEFID pairs).

As a consequence, we separate, via our oracle, QEFID pairs and one-way puzzles from OWSGs and several other Microcrypt primitives, including efficiently verifiable one-way puzzles and unclonable state generators. In particular, this solves a problem left open in [Chung, Goldin, and Gray Crypto’24].

Using similar techniques, we also establish a fully black-box separation (which is slightly weaker than an oracle separation) between private-key quantum money schemes and QEFID pairs.

One conceptual implication of our work is that the existence of an efficient verification algorithm may lead to qualitatively stronger primitives in quantum cryptography.

Contents

1	Introduction	3
1.1	Our Results	5
1.2	Overview of Techniques	6
1.3	On the Worst-Case Simulatability of Our Oracles	10
1.4	Related and Concurrent Works	11
1.5	Open Questions	12
2	Preliminaries	12
2.1	Quantum Cryptographic Primitives	13
2.2	On The Complexity of The Reflection Oracle	16
3	The Oracles	18
4	Existence of QEFID Pairs	20
5	Impossibility of OWSGs	24
5.1	Gentle Search for QPSPACE-aided POVMs	24
5.2	Proof of Theorem 5.1	26
6	Further Implications of Theorem 1.1	27
7	Separating Private-key Quantum Money from QEFID	29
7.1	Proof of Theorem 7.2	31
7.2	Proof of Lemma 7.4	33
A	Defining Fully Black-Box Separation of Private-key Quantum Money Schemes from QEFID pairs	39
B	From Absolute-Gap Distinguisher to Positive-Gap Distinguisher	40
C	Proof of Generalized Reflection Emulation	41

1 Introduction

In cryptography, a central question is to determine the minimal assumptions needed to construct various cryptographic primitives. In classical cryptography, the existence of one-way functions (OWFs) is regarded as the *minimal* computational assumption: On the one hand, OWFs imply the existence of a broad range of basic cryptographic primitives [LR86, IL89, ILL89], including pseudorandom generators (PRGs) [HILL99], pseudorandom functions (PRFs) [GGM86], commitments [Nao90], symmetric key encryption (SKE), and digital signatures [Rom90]. On the other hand, OWFs are implied by essentially any non-trivial cryptographic primitive with computational security. This equivalence gives strong evidence that OWFs may be the most basic building block in the classical cryptographic landscape.

In contrast, the situation in quantum cryptography appears to be fundamentally different. Morimae and Yamakawa [MY22] and Ananth, Qian, and Yuen [AQY22] independently initiated the study of quantum cryptography from assumptions potentially weaker than OWFs. Specifically, these works demonstrate constructions of various quantum cryptographic primitives, including commitments, SKE, and digital signatures, from pseudorandom state generators (PRSGs) [JLS18], a quantum analog of PRGs. While PRSGs can be constructed from OWFs, they are unlikely to imply OWFs [Kre21, KQST23, LMW24]. We use the term *Microcrypt* to denote the set of cryptographic primitives that are potentially weaker than OWFs.

Despite the large body of work in the subject, our understanding of the relation between different cryptographic primitives in Microcrypt is still limited. In particular, it is currently considered an open problem to determine what is the *minimal* assumption in quantum cryptography.

EFI pairs and one-way state generators. In this vein, Brakerski, Canetti, and Qian [BCQ23] introduced the concept of EFI pairs, which are efficiently samplable pairs of quantum states that are statistically far but computationally indistinguishable. They showed that the existence of EFI pairs is equivalent to that of many quantum cryptographic primitives, such as commitments, zero-knowledge proofs, and multi-party computation. In a parallel line of work, Morimae and Yamakawa [MY22, MY24] proposed one-way state generators (OWSGs), which are a quantum analogue of OWFs, as another potential minimal assumption in quantum cryptography. They demonstrated that the existence of OWSGs is implied by the existence of many primitives, including (pure-state) private-key quantum money schemes, SKE, and digital signatures. Roughly, a OWSG is a quantum polynomial-time (QPT) algorithm that maps a classical string x to a (possibly mixed) quantum state ψ_x satisfying the following requirements:¹

- **Efficient verifiability:** There is a QPT verification algorithm \mathcal{V} that accepts (x, ψ_x) with an overwhelming probability over uniform x .
- **One-wayness:** Given polynomially many copies of ψ_x , no QPT adversary can find x' such that $\mathcal{V}(x', \psi_x)$ accepts, except with negligible probability.

We can think of EFI pairs and OWSGs as postulating two different types of hardness, that in the classical world are equivalent [Gol90]: EFI pairs postulate hardness of decision, namely that of a distinguishing problem. On the other hand, OWSGs postulate hardness of search, namely that of inverting an easy-to-compute function, where solutions are efficiently verifiable. Classically, these two hardness notions are equivalent due to the Goldreich-Levin theorem [GL89]. Furthermore, in the classical world the existence of an efficient verification algorithm is without loss of generality, since one can verify the validity of a pair $(x, f(x))$, by simply recomputing f .

¹OWSGs are first defined in [MY22] for the case of pure state outputs and then generalized to the case of mixed state outputs in [MY24]. In this work, OWSGs mean the mixed state output version, unless otherwise specified.

A recent breakthrough by Khurana and Tomer [KT24] connected these two primitives by showing that OWSGs with pure state outputs imply EFI pairs. Soon after that, Batra and Jain [BJ24] generalized their result to the case of mixed state outputs. In fact, they show equivalence between EFI pairs and inefficiently-verifiable OWSGs (IV-OWSGs), which are a weaker variant of OWSGs where the verification algorithm is allowed to run in unbounded-time.² As a consequence, we now understand that most quantum cryptographic primitives imply EFI pairs (or, equivalently, IV-OWSGs) suggesting that they currently represent the minimal assumptions in quantum cryptography. However, at present, nothing is known about the *reverse direction*: Could it be that EFI pairs imply OWSGs, or is it the case that EFI pairs characterize a separate world in the depths of Microcrypt? Note that, by [BJ24], this is equivalent to asking whether IV-OWSGs imply OWSGs. In other words:

Does efficient verification come “for free” in quantum cryptography?

The objective of our work is to make progress on this question. As we shall discuss next, the phenomenon of efficient vs inefficient verification is quite common in quantum cryptography, and many of the known open problems revolve around this question.

The QCCC model. A well-studied model in the literature is the setting of quantum computation with classical communication (QCCC), where quantum computations are performed locally, but all communication is restricted to be classical. In this model, [KT24] postulated the existence of a primitive, called one-way puzzles (OWPuzzs), which directly generalizes OWFs. A OWPuzz consists of a QPT sampling algorithm S and an unbounded-time verification algorithm \mathcal{V} , where S produces a pair of classical strings: a key k and a puzzle s such that $\mathcal{V}(k, s)$ accepts. The one-wayness requires that no QPT adversary, given s , can find a key k' such that $\mathcal{V}(k', s)$ accepts, except with negligible probability. In [KT24] it is shown that OWPuzzs are implied by many primitives in the QCCC model, including public and symmetric key encryption, digital signatures, and commitments.

This idea was further developed in [CGG24], who introduced an efficiently verifiable variant of OWPuzzs (EV-OWPuzzs) where \mathcal{V} is a QPT algorithm, and showed that many (but not all) of the aforementioned primitives in the QCCC model also imply EV-OWPuzzs. For instance, the QCCC version of EFI pairs (QEFID pairs), where the distributions are classical bits, are not known to imply EV-OWPuzzs. This question was left open in [CGG24].

The techniques developed in [CGG24], which are in turn based on [Kre21], crucially rely on the fact that EV-OWPuzzs capture the hardness of problems with classical inputs and classical outputs, but do not seem to extend to the context of *quantum* inputs (and classical outputs), such as QEFID pairs.

Unclonability. A defining property of quantum mechanics is that of *unclonability of quantum states*. In cryptography, this property is used in the context of (private-key) quantum money schemes [Wie83, JLS18]. A private-key quantum money scheme allows one to mint banknotes $\$_k$ in the form of quantum state, and it should be hard to clone banknotes, i.e., to create more than t valid banknotes, given $\$_k^{\otimes t}$. Importantly, given the secret key k , one can efficiently verify the validity of a banknote state.

Although quantum money schemes are a central cryptographic primitive, that arguably started the field of quantum cryptography, very little is known on the relation with other Microcrypt primitives. For the special

²Batra and Jain [BJ24] refer to IV-OWSGs as statistically-verifiable OWSGs (sv-OWSGs). The term “IV-OWSG” was introduced by Malavolta, Morimae, Walter, and Yamakawa [MMWY24] who concurrently showed a similar result but with an exponential security loss.

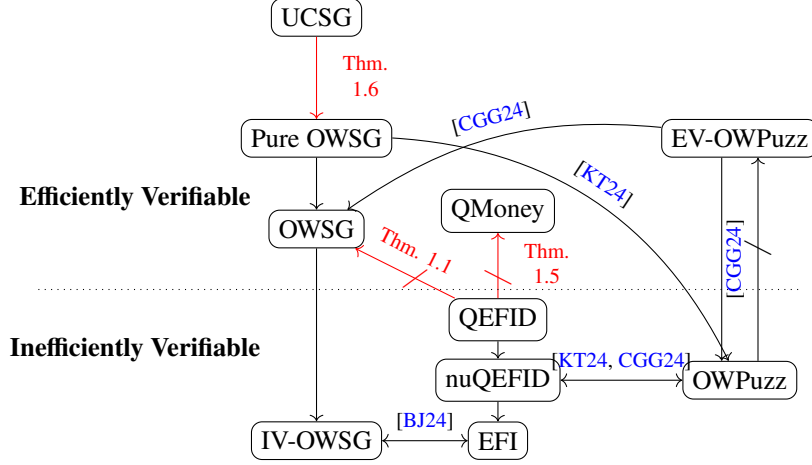


Figure 1: Implications among Microcrypt primitives, where pure OWSG means OWSGs with pure state outputs and nuQEFID means non-uniform QEFID pairs. We regard EFI pairs and (nu)QEFID pairs as inefficiently verifiable due to their equivalence (up to non-uniformity) to IV-OWSGs and OWPuzzs, respectively. Arrows without references indicate trivial implications.

case of pure money states, it is known that quantum money schemes imply OWSGs [MY24], but no general relation is known to hold either way.

1.1 Our Results

In this paper, we make progress on the above questions, and we show strong evidence that there is a qualitative difference between efficient verification and one-wayness in quantum cryptography. We provide a schematic overview of our results in Figure 1, and we discuss our main theorems in the following.

Our main result shows a separation between OWSGs and EFI pairs (Theorems 4.1 and 5.1). In fact, our main theorem is slightly more general.

Theorem 1.1 (OWSGs vs QEFID Pairs). *There exists a unitary oracle relative to which exponentially-hard QEFID pairs exist but OWSGs do not.*

QEFID pairs are a variant of EFI pairs where the outputs are not quantum states but classical bit strings (note that the generation algorithm is still QPT). Since QEFID pairs trivially imply EFI pairs, we have the following corollary, which solves the first open question negatively.

Corollary 1.2 (OWSGs vs EFI pairs). *There exists a unitary oracle relative to which exponentially-hard EFI pairs exist but OWSGs do not.*

Moreover, our main theorem leads to further implications, which we hereby summarize .

Since EV-OWPuzzs imply OWSGs [CGG24, Theorem 32, Section 10], we obtain a separation between the analogous primitives of EFI pairs and OWSGs, in the QCCC model. This gives a negative answer to a possibility raised in [CGG24].

Corollary 1.3 (EV-OWPuzzs vs QEFID pairs). *There exists a unitary oracle relative to which QEFID pairs exist, but EV-OWPuzzs do not.*

Finally, since QEFID pairs imply OWPuzzs [CGG24, Lemma 8], our main theorem also separates OWSGs from OWPuzzs.

Corollary 1.4 (OWPuzzs vs OWSGs). *There exists a unitary oracle relative to which OWPuzzs exist, but OWSGs do not.*

Overall, our results provide a clearer picture of quantum one-wayness, by showing a separation between efficiently and inefficiently verifiable versions of the different cryptographic primitives.

Next, we turn our attention to the relation between unclonable cryptography and other primitives in Microcrypt. As our main theorem, we show the first separation between quantum money and QEFID pairs.

Theorem 1.5 (Quantum Money vs QEFID). *There exists no fully black-box construction of private-key quantum money schemes from exponentially-hard QEFID pairs.*

On the other hand, using similar techniques devised to prove Theorem 1.1, we show that unclonable state generators (UCSGs) imply OWSGs. UCSGs (implicitly defined in [JLS18]) have the same syntax as quantum money, but they satisfy a slightly stronger unclonability guarantee.

Theorem 1.6 (UCSGs vs OWSGs). *UCSGs imply pure OWSGs.*

Finally, combining Theorems 1.1 and 1.6, we obtain the following implication.

Corollary 1.7 (UCSGs vs QEFID pairs). *There exists a unitary oracle relative to which QEFID pairs exist, but UCSGs do not.*

1.2 Overview of Techniques

To obtain our main result, namely, the separation of OWSGs from QEFID pairs, we aim to find a unitary oracle that provides us with the right amount of “hardness” on which we may base the construction of an QEFID pair yet, at the same time, gives enough “computational power” to break OWSGs, deeming them non-existent. As is often the case in oracle separations, it is thus easier to split our task and regard it as building two oracles: one that models the hardness sufficient for QEFID pairs and another that models a powerful inverter that breaks OWSGs.

Before describing our oracles, let us give some underlying intuition. Statistical-vs-computation gap between distributions (as implied by QEFID pairs) necessitates producing randomness. In the classical setting, where computation is always deterministic, such randomness may be generated only by applying a deterministic procedure over a random seed. Such a procedure is *derandomizable*. Additionally, one can efficiently verify whether an output was generated using a given seed, by repeating the computation. Our construction of QEFID pairs cannot be derandomizable, i.e., it cannot possibly rely solely on such a source of randomness since it cannot be secure in a world where any OWSG is broken.³ Our QEFID pairs therefore must attain their entropy from a different source of randomness, namely from quantum uncertainty. In the quantum world, we are able to produce randomness through non-deterministic quantum computation that involves measurements. Importantly, such a procedure is *underandomizable* and does not generally induce efficient verifiability of its outcome.

³Such a construction of QEFID pairs, by the same proof as in [Gol90], would imply (quantum-evaluatable) PRGs that trivially imply OWSGs.

Building QEFID pairs relative to an oracle. A most basic underrandomizable quantum procedure that generates a random outcome is the following: prepare a uniform superposition $\sum_{x \in X} |x\rangle$ then measure in the standard basis. This will sample a uniformly random element in X . A natural way to produce a pseudorandom distribution over λ bits (which implies, in particular, QEFID pairs), then, is to somehow create the state $\sum_{s \in S} |s\rangle$, where S is a large enough uniformly random subset of $\{0, 1\}^\lambda$, then measure to obtain a uniformly random $s \leftarrow S$. This brings us to define our first oracle: The oracle O is a controlled swap between $|0^\lambda\rangle$ and $|S\rangle = \sum_{s \in S} |s\rangle$, where S is a uniformly random subset of size $2^{\lambda/2}$ over $\{0, 1\}^\lambda \setminus \{0^\lambda\}$. That is, given control bit 1, the oracle maps $|0^\lambda\rangle$ to $|S\rangle$, $|S\rangle$ to $|0^\lambda\rangle$, and acts as identity on the subspace orthogonal to the span of $|0^\lambda\rangle$ and $|S\rangle$. On control bit 0, the oracle is identity. Our QEFID pairs under O simply call the oracle O with input $|1\rangle |0^\lambda\rangle$ to obtain $|S\rangle = \sum_{s \in S} |s\rangle$, then measure to get a uniformly random element from S .⁴

To argue why our QEFID pairs give a pseudorandom distribution under the oracle O (remember there is one more oracle that we ought to introduce), we apply a *reflection emulation* technique, first introduced by Ji, Liu and Song [JLS18]. In their work, they show that access to a reflection oracle about a state $|\psi\rangle$, i.e. the unitary oracle $R_\psi = I - 2|\psi\rangle\langle\psi|$, may be simulated given sufficiently many copies of $|\psi\rangle$. This is useful to us due to the key observation that the oracle O is, in fact, the reflection unitary about the state $|1\rangle |S-\rangle$, where $|S-\rangle = \frac{1}{\sqrt{2}}(|S\rangle - |0^\lambda\rangle)$. Consequently, any algorithm \mathcal{A} with access to O that is successful in breaking our QEFID pairs may be simulated by an algorithm \mathcal{B} that is not given access to O , but takes as input sufficiently many copies of $|S-\rangle$, and breaks QEFID pairs as successfully.

Therefore, we have reduced our task to showing that any such \mathcal{B} is incompetent in breaking the QEFID construction. Recall such \mathcal{B} takes as input many copies of the state $|S-\rangle$, for a uniformly random subset S of size $2^{\lambda/2}$, and a string $s \in \{0, 1\}^\lambda$, which is sampled either uniformly at random or uniformly at random from the subset S . To show indistinguishability between these two cases, we bound the trace distance between the mixed states corresponding to the two distributions. Importantly, this argument is statistical and holds even given any additional oracle that is independent of O .

Breaking OWSGs. Having shown the existence of QEFID pairs relative to the oracle O , it remains to prove the other half of our statement: OWSGs do not exist. This requires complementing the oracle O with an additional oracle that gives the necessary computational power for such a task.

The question of identifying an oracle under which there is no efficiently verifiable quantum one-wayness has been studied in the literature. Cavalari et al. [CGG⁺23] show that, given quantum access to a **PP** oracle, there exists a generic attack that breaks any pure OWSGs. Besides the fact that it is not clear how to extend their attack to any OWSGs, there is a more inherent limitation to adapting their approach to our setting. More specifically, their attack breaks OWPuzzs, which in turn implies that pure OWSGs are broken (due to [KT24], see Figure 1). Since QEFID pairs too imply OWPuzzs, any attempt to apply their attack (or any strengthening thereof) to our world, consisting of the oracle O , will result in an attack against our QEFID construction if successful. Therefore, we cannot hope to use their approach to establish a separation by, for instance, appending the oracle **PP** ^{O} to our world.

A different method to generically break efficiently verifiable one-wayness is proposed by the framework of *shadow tomography* [Aar19, HKP20], specifically the task of *gentle search* [Aar19]. A gentle search algorithm takes, as input, a collection of POVM elements $\{\Pi_k\}_{k \in \mathcal{K}}$ (indexed by an arbitrary key space \mathcal{K}) and a state $|\phi\rangle$. Its goal is to output a key $k \in \mathcal{K}$ such that Π_k accepts $|\phi\rangle$ with good probability, i.e., $\text{Tr}(\Pi_k |\phi\rangle\langle\phi|)$ is large, assuming such a k exists. Breaking OWSGs via gentle search is straightforward: Letting the POVM Π_k denote verification under key k , gentle search finds a key under which the input state is verified, hence successfully

⁴We could potentially define our oracle to perform the measurement as well, outputting a random $s \leftarrow S$ at every query. This would not be, however, a unitary oracle as we have promised.

inverting OWSGs. As observed by [CCS24], the gentle search procedure from [Aar19] can be performed by a **QPSPACE** computation⁵ when the POVMs are implemented by polynomial-size quantum circuits, even when these circuits themselves have access to a **QPSPACE** oracle. Consequently, by augmenting our world with a **QPSPACE** oracle,⁶ we may break any OWSG construction that uses the **QPSPACE** oracle. This is, however, not entirely sufficient for our separation result to hold: We must rule out any OWSG construction that possibly uses the oracle O as well.

A straightforward attempt would be to equip the attack with the oracle $\mathbf{QPSPACE}^O$, i.e., now the oracle can implement any poly-space quantum circuit that has access to O . Such an oracle, however, is too strong as it breaks our QEFID construction as well due to [CGG⁺23] (since the attack using the **PP** oracle can be performed using the **QPSPACE** oracle). Instead, we observe that the reflection emulation technique, which we used previously to prove the security of our QEFID pairs, comes in handy here as well. By reflection emulation, we can replace the verification circuit of OWSGs, which possibly makes queries to O and **QPSPACE**, with a circuit that makes queries only to **QPSPACE**, yet is given in addition copies of the state $|S-\rangle$. We then apply the gentle search procedure described above w.r.t. the POVM elements $\{\Pi_k\}_k$, where Π_k implements the new **QPSPACE**-aided verification circuit with key k .

Of course, now the POVM elements expect as input, besides the state $|\phi\rangle$ to be verified, a polynomial amount of copies of $|S-\rangle$. Hence, we must provide the gentle search with sufficiently many such copies. To generate these copies, we apply phase estimation for the unitary $R_{S-} := I - 2|S-\rangle\langle S-|$ on the state $|0^\lambda\rangle$ using O : Recall O is a reflection about $|1\rangle|S-\rangle$ and hence a controlled-version of R_{S-} , and $|0^\lambda\rangle$ is a uniform superposition of $|S-\rangle = \frac{1}{\sqrt{2}}(|0^\lambda\rangle - |S\rangle)$, which has eigenvalue -1 under R_{S-} , and $|S+\rangle = \frac{1}{\sqrt{2}}(|0^\lambda\rangle + |S\rangle)$, with eigenvalue 1 .

To conclude, our attack against any OWSG construction with efficient verification circuit \mathcal{V} , under oracles O and **QPSPACE**, takes as input sufficiently many copies of the challenge state $|\phi\rangle$ to be inverted, and performs the following steps: 1. Compute the **QPSPACE**-aided circuit \mathcal{V}' , which emulates \mathcal{V} while replacing the calls to the oracle O with taking as input copies of the state $|S-\rangle$. 2. Generate sufficiently many copies of $|S-\rangle$ via phase estimation for O . 3. Using the **QPSPACE** oracle, perform gentle search w.r.t. the POVM elements $\{\Pi_k\}_k$, where Π_k implements \mathcal{V}' under key k , and input constituting of $|\phi\rangle$ and copies of $|S-\rangle$ (at this point, we have sufficiently many copies of this input as required by gentle search).

The attack against any OWSGs under the oracles O and **QPSPACE**, together with the existence of a QEFID pair under these oracles, completes the oracle separation between OWSGs and QEFID pairs.

We now proceed to discuss the techniques underlying our second main result, namely separating private-key quantum money schemes from QEFID pairs.

Separating Private-key Quantum Money schemes from QEFID pairs. Recall that a private-key quantum money scheme consists of three QPT algorithms, KeyGen that samples a key k , Mint that mints a quantum money state \mathbb{S}_k (which can be generally mixed), and Verify that on receiving a key k and an alleged money state ρ either accepts or rejects. The money scheme is called unforgeable if it is hard, for any QPT adversary who is given money states minted under $k^* \leftarrow \mathit{KeyGen}$, to produce new money states that are valid under the same key k^* . That is, for any polynomial m , the adversary is given $\mathbb{S}_{k^*}^{\otimes m}$ sampled by running $\mathit{Mint}(k^*)$ for m times, where $k^* \leftarrow \mathit{KeyGen}$. The adversary is given access to the verification oracle w.r.t. k^* and it breaks the scheme if it outputs a quantum state that contains at least $m + 1$ valid money states that pass verification

⁵Note that this **QPSPACE** is not the class of classical decision problems that can be decided by quantum polynomial-space computing. This is a unitary operation over polynomial number of qubits.

⁶As we have mentioned, this is not an oracle that solves classical decision problems, but a unitary operation. Therefore, this oracle takes quantum states as input and outputs quantum states.

under k^* .

When quantum money schemes produce pure money states they imply OWSGs [MY24]. However, for general mixed money states, we do not know whether private-key quantum money schemes imply OWSGs or not. This gap becomes evident when we try to break private-key quantum money schemes by using the same attack as the one used to break OWSGs in the proof of Theorem 1.1.

Indeed, assume we try to apply gentle search to find the key k^* using which the given money states were minted. Clearly, given such a key we may reproduce as many valid money states as we wish. The only guarantee given by gentle search, however, is that the key it finds accepts the input money state $\$_{k^*}$ with good probability. That is, it is possible that the gentle search returns a key $k \neq k^*$ such that $\$_{k^*}$ is valid also under k , yet producing a new money state using k will result in a money state $\$_k$ that does not pass verification under k^* . For example, consider a quantum money scheme where there is a dummy key k_0 under which any state passes verification. Gentle search, in this case, might always return k_0 and is hence useless.⁷

A different strategy is then required to attack arbitrary quantum money schemes that make use of the oracle O . While several such statistical attacks have been proposed in the literature [Aar16, Aar19], none seem to be applicable in our oracle-relative world: An extension of the gentle search procedure [Aar16] which succeeds in breaking any private-key quantum money scheme in the plain model requires calling the verification circuit an exponential number of times, consequently requiring exponential number of calls to the oracle O in our setting (or preparing exponentially many copies of $|S-\rangle$ if we wish to evoke reflection emulation). Another known attack, in some ways more efficient, is based on the stronger tool of *shadow tomography* [Aar19]. Unfortunately, as we discuss next, the attack from [Aar19] also fails in our setting, but we manage to use shadow tomography in a different attack to break any private-key quantum money scheme relative to O .

Shadow tomography may be seen as a strengthening of gentle search. In shadow tomography, we are again given a collection of POVM elements $\{\Pi_k\}_{k \in \mathcal{K}}$ and an input state ρ , and the goal is to output, up to some precision error, the acceptance probability of ρ under *all given POVM elements*, i.e. $\text{Tr}(\Pi_k \rho)$ for all k (recall gentle search outputs a key where such probability high enough). As noted by [Aar19], shadow tomography immediately gives an attack against any private-key quantum money scheme in the plain model. Letting the POVM element Π_k correspond to the acceptance of the verification circuit under key k , shadow tomography allows us to (approximately) estimate the probability that a given state $\$$ passes verification under each key, namely, $\Pr[\top \leftarrow \text{Verify}(k, \$)]$ for each k . Then, given $\$_{k^*} \leftarrow \text{Mint}(k^*)$, an attacker can find, by brute-force, a different state $\$'$ that has close-enough acceptance probabilities under all keys and output it. The attacker can produce as many copies of $\$'$ as he wants. In particular, since $\$_{k^*}$ passes verification under k^* with good probability, then so does $\$'$.

The above approach completely fails when verification has access to the oracle O , even with reflection emulation in hand. Consider carrying the above attack w.r.t. a verification circuit that, instead of querying O , takes as input polynomially many copies of $|S-\rangle$ besides $\$$ (as obtained by the reflection emulation technique from [JLS18]). Via brute-force, we are able to find all states that exhibit acceptance probabilities similar to the input $\$ \otimes |S-\rangle^{\otimes t}$ (where t is the number of copies required by reflection emulation). These may include states of the form $\$' \otimes |S'-\rangle^{\otimes t}$, for $S' \neq S$. Since the number of possible subsets $S' \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$ of size $2^{\lambda/2}$ is doubly-exponential, it is impossible to identify those satisfying $S = S'$ using polynomially many queries to the oracle O (as this would imply an efficient algorithm for learning S and consequently breaking

⁷This being said, we note that, when the money states are pure, there is a successful alternative way to apply gentle search. Rather than considering the POVMs that implement the verification circuit under k , perform gentle search w.r.t. the POVMs that implement a swap test between their input and the pure money state minted under the corresponding key k (essentially, performing a projection into the latter).

our QEFID pairs).

Alternatively, one may attempt to perform the brute force over the money state register alone; to any possible value it may take, attach the copies $|S-\rangle^{\otimes t}$ that we take as input, then invoke shadow tomography thereon. For every such possible money state, we obtain the acceptance probabilities under all keys, which we compare to those of the given money states until a match is found. Every invocation of shadow tomography, however, disturbs the copies of $|S-\rangle$. Such disturbance, even if in the slightest, limits the number of states we can brute-force over to an insufficient number of trials.

While the above brute-force idea does not work for us, it serves as inspiration to our ultimate solution. We observe that we need not brute-force over all possible states that the money state register may take. Rather, it is sufficient to look at states that can be produced by the minting algorithm of the private-key money scheme, i.e., $\$k \leftarrow \text{Mint}(k)$ for any key k . Roughly speaking, since these states have classical description (i.e. the corresponding keys), we are able to “brute-force” over them via a single invocation of shadow tomography, thus requiring only a polynomial number of copies of $|S-\rangle$! In more details, just as before, we apply shadow tomography to estimate the probability of acceptance of the input state $\$k^*$ under any key k , which we denote by b_k^* . Next, we apply shadow tomography w.r.t. the POVM elements $\{(\Pi_{k',k}, I - \Pi_{k',k})\}_{k,k'}$ and the all-zero state, where, for any pair of keys k', k , $\Pi_{k',k}$ corresponds to the following event: 1. Generate $\$k' \leftarrow \text{Mint}(k')$ by running $\text{Mint}(k')$ (which can be implemented as a unitary since k' is classical) on many copies of the state $|S-\rangle^{\otimes t}$. 2. Run $c \leftarrow \text{Verify}(k, \$k')$ and the verification accepts.

For any key k' , we obtain the acceptance probability of $\$k'$ under all possible keys k , which we denote by $\{b_{k',k}\}_k$. We choose a key k' such that $\{b_{k',k}\}_k \approx \{b_k^*\}_k$, and output sufficiently many states minted using k' . Note that such a k' exists since, for $k' = k^*$, we have $\{b_{k',k}\}_k \approx \{b_k^*\}_k$ as $\$k'$ distributes like the input $\$k^*$. In fact, it is evident that any key k' satisfying $\{b_{k',k}\}_k \approx \{b_k^*\}_k$ will be good enough for forfeiting purpose as money minted using k' will have acceptance probability under k^* close to that of $\$k^*$ (which is close to 1 by correctness), and hence the attack is successful.

We note that our attack against private-key quantum money schemes is an (O -aided) statistical attack with inefficient runtime and space complexity. However, it is query-efficient. It is not clear whether there exists an oracle O' under which such an attack can be made efficient and still be able to break any private-key quantum money scheme, possibly using the oracle O' . Thus, in contrast to the first separation result of OWSGs from QEFID pairs, our separation of private-key quantum money schemes from QEFID pairs is, formally speaking, a fully-black-box separation and does not satisfy the stronger notion of an oracle separation between the two primitives. We leave the question of establishing an oracle separation between private-key quantum money schemes and QEFID pairs to future work.

1.3 On the Worst-Case Simulatability of Our Oracles

Existing separations involving Microcrypt primitives [Kre21, CCS24, CM24, AGL24] heavily rely on the concentration property of the Haar measure, such as Lévy’s lemma [Wat18, Theorem 7.37]. In contrast, the oracles with respect to which we achieve Theorems 1.1 and 1.5 are reflections about random subset states. This is the first example of separations between Microcrypt primitives in an oracular world that does not involve common Haar random states or unitaries. We elaborate on what we think is a feature of our oracles which, despite seemingly technical, we believe offers the potential for further applications of our approach to new separation results.

In both our separation and separations relative to common Haar random states or unitaries, the proofs involve “de-oraclicizing” certain algorithms, i.e., simulating oracle-aided algorithms by algorithms that do not query any oracles, albeit via completely different and incomparable simulation techniques.

In our proof, we invoke a simulation technique from the work of [JLS18], which we refer to as *reflection emulation* (see Section 1.2). Via reflection emulation we are able to simulate our reflection oracle *in the worst-case*, independently of the distribution on the oracle. In sharp contrast, the simulation of Haar random states or unitaries, as performed in previous works [Kre21, CCS24, CM24] crucially relies on the concentration property of the Haar measure, and hence the simulation guarantees are *distribution-sensitive*, i.e., may not hold for a different non-Haar distribution over the same support.

This being said, due to the Haar concentration measure, the guarantee from Haar state simulation is stronger in terms of precision, as the concentration of Haar measure provides precision inverse-exponential in the dimension. In comparison, simulating reflection about a state $|\psi\rangle$ via reflection emulation provides precision inverse-polynomial in the number of copies of the state $|\psi\rangle$ that are given.

To demonstrate the advantage of worst-case simulation over a distribution-sensitive one, consider the following example. Given a Haar random unitary U of dimension D , let us look at a QMA verifier V^U , that on any n -bit instance, takes as input an alleged witness state $|\phi\rangle \in \mathbb{C}^D$, queries U on $|0\dots 0\rangle$ n times to get the state $|\psi_U\rangle^{\otimes n}$, where $|\psi_U\rangle = U|0\dots 0\rangle$, then performs a swap test on the given state $|\phi\rangle$ and $|\psi_U\rangle^{\otimes n}$. Clearly, if we apply standard Haar random emulation [Kre21], then we would simulate queries to U by sampling a different unitary U' and running $V^{U'}$. In such a case, the largest possible acceptance probability of the obtained simulation $V^{U'}$, for any fixed oracle U , is $\frac{1}{2} + 1/\binom{2^n+D-1}{2^n-1}$. In contrast, the original verifier V^U has maximum acceptance probability 1 as it always accepts the state $|\psi_U\rangle^{\otimes n}$. On the other hand, if we consider a reflection unitary U and the same verifier as above, it can be easily shown that the worst-case reflection emulation results in the same maximum acceptance probability for the emulated verification, up to emulation precision.

1.4 Related and Concurrent Works

Separations inside Microcrypt have been studied before. [CGG24] separates EV-OWPuzzs from PRSGs by ruling out EV-OWPuzzs in the oracular world of [Kre21]. Separations of QCCC key agreement protocols and QCCC commitments from Pseudorandom function-like state generators (PRFSGs) in the common Haar random state model were shown in [AGL24]. Recently, Chen, Coladangelo, and Sattath [CCS24] showed a separation between 1-PRSGs and PRSGs. In a concurrent work [BCN24], the authors prove a related result to Theorem 1.1 by separating OWSGs from OWPuzzs and 1-PRSGs. We outline a few differences between [CCS24, BCN24] and our work.

The results in [CCS24, BCN24] bear resemblances but are incomparable to Theorem 1.1 since [CCS24] rules out PRSGs while constructing 1-PRSGs (and hence, quantum commitments), and [BCN24] rules out OWSGs while constructing 1-PRSGs (and hence, quantum commitments) and OWPuzzs whereas in Theorem 1.1, we construct QEFID pairs (and hence OWPuzzs, and quantum commitments) while ruling out OWSGs in the respective oracular worlds. Notably, since one of the distributions in our construction of QEFID pairs (see Construction 4.2) is the uniform distribution, our works and [CCS24, BCN24] construct pseudorandom (but statistically far from the uniform) distributions but in different regimes. Namely, we construct a classical pseudorandom distribution but in a keyless regime, i.e., the quantum sampler simply samples a classical bit string, whereas [CCS24, BCN24] construct 1-PRSGs, which can be viewed as quantum pseudorandom distributions in a keyed regime, i.e., first, a classical key of size λ is sampled, and then a pure quantum state of size larger than λ is generated using the key. Note that it is not possible to construct a classical pseudorandom distribution in a keyed regime and also hope to rule out OWSGs, since such a keyed classical pseudorandom distribution will constitute a quantum-evaluable PRGs, which trivially implies OWSGs.

Second, our oracles differ from that of [CCS24, BCN24] as both works rely on the common Haar random state model, whereas we consider an oracle corresponding to random subset states. Even though random subset states of appropriate size are known to be statistically close to Haar random states [JMW23, GTB23], this does not imply our oracle separation of OWSGs from QEFID pairs in Theorem 1.1 also holds under the oracles considered in [CCS24, BCN24], and vice-versa. This is because the statistical indistinguishability in [JMW23, GTB23] is an average-case guarantee, which implies that the two oracle distributions are indistinguishable only on average. Hence, even if there exists an oracle in the support of one distribution relative to which every construction of the primitive (say OWSGs) is broken by an adversary, there might not exist such an oracle, and such an adversary in the second distribution.

1.5 Open Questions

Our work leaves the following questions open:

1. While our oracles do not involve sampling Haar random states and unitaries, our oracles are still quantum. The natural next step would be to upgrade the oracles to a classical oracle with quantum superposition access, which we leave open for future works.
2. We show that there are no EV-OWPuzzs in the oracular world we consider, which by [CGG24] suggests that many quantum primitives in the QCCC model, including non-interactive QCCC commitments, cannot exist in this world. This still leaves open the question of whether QCCC commitments exist in this world. Any answer to this question would greatly affect our understanding of quantum commitments. A positive answer to this question would imply that the gap between OWSGs and commitments is independent of the quantum communication in the commitments protocol. On the other hand, a negative answer to the question would mean that QCCC commitments are strictly stronger than quantum commitments, implying that quantum communication is indeed a crucial resource for quantum commitments.
3. In the oracular world that we consider, QEFID pairs exist but OWSGs do not. Can we show an oracle separation in the opposite direction? That is, can we give an oracle relative to which OWSGs exist but QEFID pairs do not? A slightly stronger question would be: can we give an oracle relative to which OWSGs exist but OWPuzzs (which are implied by QEFID pairs) do not? Note that the construction of OWPuzzs from OWSGs by [KT24] crucially requires that the OWSGs are pure, and hence a positive answer to the question mentioned above does not contradict [KT24]. Due to the construction of OWPuzzs from pure OWSG by [KT24], an oracle separation between arbitrary mixed OWSGs and OWPuzzs would imply a separation between pure and arbitrary mixed OWSGs.
4. We show a statistical attack on any private-key quantum money schemes in our oracular world. Can we improve it to an efficient attack, relative to some unitary oracle, thereby improving the separation between private-key quantum money schemes and QEFID pairs to an oracle separation?
5. In our oracular world, both EFI pairs and QEFID pairs exist: Is there an oracle separation between EFI pairs and QEFID pairs? What is the relation between the two notions?

2 Preliminaries

Notation and Terminology. We use standard notations of cryptography and quantum information. We use λ and κ to denote security parameters. For a set S , we use $x \leftarrow S$ to denote that x is sampled uniformly at

random from S . For a distribution D , we use $x \leftarrow D$ to denote that x is sampled according to the distribution D . For a distribution D , we use $x \in D$ to denote that x is in the support of D . For two distributions $P = \{p_x\}_x$ and $Q = \{q_x\}_x$, we denote by $\text{SD}(P, Q) := \frac{1}{2} \sum_x |p_x - q_x|$ the statistical distance between P and Q . For two quantum states ρ and σ , $\text{TD}(\rho, \sigma)$ is their trace distance.

A quantum oracle is an oracle that realizes arbitrary (possibly inefficient) quantum channel. A unitary oracle is given by an oracle where the channel is a unitary. An oracle-aided (quantum) algorithm is a quantum algorithm that has access to a quantum oracle, i.e., the circuit description of the algorithm can have the oracle as a gate. We say that an algorithm is O -aided if it presumes access to the oracle O . We naturally extend this notation to a set \mathcal{O} , or distribution, of oracles when access to $O \in \mathcal{O}$ is given. We say that an oracle-aided algorithm is polynomial-query if there exist polynomials $q, p : \mathbb{N} \rightarrow \mathbb{N}$ such that the algorithm, on input of length n , makes at most $q(n)$ queries to its oracle, where each query is of length at most $p(n)$ (qu)bits. I is the two-dimensional identity operator. We often write $I^{\otimes n}$ as I if the dimension is clear from the context. For any pure quantum state $|\phi\rangle$, we denote the reflection unitary about ϕ by $R_\phi = I - 2|\phi\rangle\langle\phi|$. We denote the controlled reflection unitary by R_ϕ^c . It holds that $R_\phi^c = R_{1,\phi} = I - 2|1\rangle|\phi\rangle\langle\phi|\langle 1|$.

2.1 Quantum Cryptographic Primitives

We hereby provide formal definitions of the main quantum cryptographic notions considered in this work: QEFID pairs, OWSGs, and private-key quantum money schemes. Additionally, we give definitions of related notions to which implications of our separation results extend.

We begin with the definition of QEFID pairs.

Definition 2.1 (QEFID pairs). *Let $T : \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon : \mathbb{N} \rightarrow [0, 1]$. A (T, ϵ) -QEFID pair is a pair of distribution ensembles $(D_0, D_1) := (\{D_0(1^\lambda)\}_\lambda, \{D_1(1^\lambda)\}_\lambda)$ over classical bit strings that satisfy the following properties:*

- **Efficiently samplable:** *There exists a QPT algorithm which, on input security parameter 1^λ and a bit $b \in \{0, 1\}$ samples from $D_b(1^\lambda)$.*
- **Statistically far:** *There exists a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that, for any $\lambda \in \mathbb{N}$,*

$$\text{SD}(D_0(1^\lambda), D_1(1^\lambda)) \geq 1/p(\lambda).$$

- **Computationally indistinguishable:** *For any non-uniform quantum algorithm \mathcal{A} that, on input 1^λ and a bit string z , runs in time $T(\lambda)$, and any $\lambda \in \mathbb{N}$,*

$$\left| \Pr_{z \leftarrow D_0(1^\lambda)} [\mathcal{A}(1^\lambda, z) = 1] - \Pr_{z \leftarrow D_1(1^\lambda)} [\mathcal{A}(1^\lambda, z) = 1] \right| \leq \epsilon(\lambda).$$

We say that (D_0, D_1) is a (T, ϵ) -QEFID pair relative to an oracle O , if the distributions are sampled by a polynomial-query O -aided algorithm and computational indistinguishability holds against any O -aided adversary \mathcal{A} that, on input 1^λ and a bit string, runs in time at most $T(\lambda)$.

Lastly, we say that such a pair is simply an QEFID pair if T is a polynomial and ϵ is negligible. We say that it is an exponentially-hard QEFID pair if there exists a constant $c \in \mathbb{N}$ such that $T(\lambda) = \Omega(2^{\frac{\lambda}{c}})$ and $\epsilon(\lambda) = O(2^{-\frac{\lambda}{c}})$. These notions immediately extend to the oracle-relative setting.

Next, we recall the definition of one-way state generators (OWSGs) from [MY22]. Note that [BJ24] shows equivalence between *inefficiently-verifiable* OWSGs and EFI pairs. Here, we consider *efficiently-verifiable* OWSGs.

Definition 2.2 (OWSGs [MY24, MY22]). A one-way state generator, or OWSG for short, with associated key space $\mathcal{K} = \{\mathcal{K}_\kappa \subseteq \{0, 1\}^\kappa\}_\kappa$, is a couple $(\mathcal{G}, \mathcal{V})$ of QPT algorithms with the following syntax:

- $\phi \leftarrow \mathcal{G}(k)$: On input a key $k \in \mathcal{K}_\kappa$, the generation algorithm outputs a (possibly mixed) quantum state ϕ ,
- $\{0, 1\} \leftarrow \mathcal{V}(k, \phi)$: On input a key $k \in \mathcal{K}_\kappa$ and a state ϕ , the verification algorithm outputs 1 (accepts) or 0 (rejects),

and that satisfies the following properties:

- **Correctness:** There exists a negligible function⁸ $\text{negl} : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Pr[1 \leftarrow \mathcal{V}(k, \phi); k \leftarrow \mathcal{K}_\kappa, \phi \leftarrow \mathcal{G}(k)] \geq 1 - \text{negl}(\kappa).$$

- **One-wayness:** For any non-uniform QPT algorithm \mathcal{A} and any polynomial $t : \mathbb{N} \rightarrow \mathbb{N}$, there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow [0, 1]$ such that, for any $\kappa \in \mathbb{N}$,

$$\Pr[1 \leftarrow \mathcal{V}(k', \phi); k \leftarrow \mathcal{K}_\kappa, \phi \leftarrow \mathcal{G}(k), k' \leftarrow \mathcal{A}(1^\kappa, \phi^{\otimes t(\kappa)})] < \epsilon(\kappa).$$

We say that $(\mathcal{G}, \mathcal{V})$ is a OWSG relative to an oracle O if the algorithms \mathcal{G} and \mathcal{V} are O -aided and satisfy correctness and one-wayness w.r.t any non-uniform polynomial-query O -aided quantum adversary \mathcal{A} .

In Appendix A, we further define a black-box construction of private-key quantum money schemes from QEFID pairs.

We now give definitions for one-way puzzles and the special case of efficiently verifiable one-way puzzles.

Definition 2.3 (OWPuzzs [KT24] and Efficiently-Verifiable OWPuzzs [CGG24]). A one-way puzzle (OWPuzz) is a pair $(\text{Samp}, \text{Ver})$ of algorithms with Samp being QPT, having the following syntax:

1. $(\text{ans}, \text{puzz}) \leftarrow \text{Samp}(1^\kappa)$ takes as input the security parameter κ , and outputs two bit strings, an answer and a puzzle.
2. $b \leftarrow \text{Ver}(\text{ans}', \text{puzz})$ takes an (alleged) answer, ans' , and the puzzle, and outputs 1 (accept) or 0 (reject).

We require the following two properties.

- **Correctness:** There exists a negligible function $\text{negl} : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Pr[1 \leftarrow \text{Ver}(\text{ans}, \text{puzz}); (\text{ans}, \text{puzz}) \leftarrow \text{Samp}(1^\kappa)] \geq 1 - \text{negl}(\kappa).$$

- **Security:** For any non-uniform QPT algorithm \mathcal{A} , there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Pr[1 \leftarrow \text{Ver}(\text{ans}', \text{puzz}); (\text{ans}, \text{puzz}) \leftarrow \text{Samp}(1^\kappa), \text{ans}' \leftarrow \mathcal{A}(1^\kappa, \text{puzz})] \leq \text{negl}(\kappa).$$

We say that $(\text{Samp}, \text{Ver})$ is an efficiently-verifiable one-way puzzle (EV-OWPuzz) if Ver is also a QPT algorithm.

⁸We note that this is equivalent to a definition with inverse poly correctness as correctness can be amplified by standard repetition.

Lastly, we define *unclonable state generators*, which were implicitly considered in the work of [JLS18].

Definition 2.4 (Unclonable State Generators (UCSGs) ([BEM⁺23], Implicit in [JLS18])). An unclonable state generator (UCSG) Samp with a key space $\mathcal{K} = \{\mathcal{K}_\kappa \subseteq \{0, 1\}^\kappa\}_\kappa$ is a QPT algorithm that, on input a key k , outputs a pure state $|\phi_k\rangle$. We require the following property, which we call **unclonability**: For any QPT adversary \mathcal{A} ,

$$\mathbb{E}_{k \leftarrow \mathcal{K}_\kappa} \left[\text{Tr} \left(|\phi_k\rangle \langle \phi_k|^{\otimes m+1} \mathcal{A}(1^\kappa, |\phi_k\rangle^{\otimes m}) \right) \right] \leq \text{negl}(\kappa). \quad (1)$$

We say that a UCSG is a **strong UCSG** if the security also holds against any adversary \mathcal{A} who has polynomial oracle access to the controlled reflection unitary $R_{\phi_k}^c$.

Lastly, we say that a UCSG is a (strong) UCSG relative to an oracle O , if Samp is a polynomial-time O -aided generation algorithm, and satisfies the security w.r.t any non-uniform polynomial-time O -aided (and, resp., $R_{\phi_k}^c$ -aided) quantum adversary \mathcal{A} .

Remark 2.5. We note that despite the similarities between unclonability and unforgeability, there is a subtle difference in the definition of UCSGs (Definition 2.4) and private-key quantum money schemes (Definition 2.6) even when restricted to pure states. In the cloning game (Definition 2.4), the cloner can ask for any (polynomially large) m copies of the state $|\phi_k\rangle$, and the winning condition can be interpreted as that the cloner wins if it submits exactly $m + 1$ registers, such that all the registers should pass the rank-1 projection into $|\phi_k\rangle$, whereas in the forging game (Definition 2.6 restricted to pure states), the forger after getting any m copies of the money state $|\phi_k\rangle$ can output any (polynomial) number of registers and wins if least $m + 1$ passes verification. This difference of passing at least $m + 1$ verifications instead of submitting exactly $m + 1$ registers and passing all the rank-1 projections can be bridged if the cloning adversary gets access to the rank-1 projection into the state $|\phi_k\rangle$, and hence strong UCSGs and pure private-key quantum money schemes are equivalent.⁹

Lastly, we define private-key quantum money schemes.

Definition 2.6 (Private-Key Quantum Money Schemes [JLS18]). A private-key quantum money scheme is a tuple $(\text{KeyGen}, \text{Mint}, \text{Verify})$ of QPT algorithms having the following syntax.

1. $k \leftarrow \text{KeyGen}(1^\kappa)$ takes a security parameter κ and outputs a classical secret key, $k \in \{0, 1\}^\kappa$.
2. $\$k \leftarrow \text{Mint}(k)$ takes the secret key k and outputs a quantum money state $\$k$.
3. $b \leftarrow \text{Verify}(k, \$)$ receives the secret key k and an (alleged) quantum money state $\$$, and outputs 1 (accept) or 0 (reject).

We require the following properties.

- **Correctness:** A private-key quantum money scheme is called μ -correct, for $\mu : \mathbb{N} \rightarrow [0, 1]$, if for all $\kappa \in \mathbb{N}$,

$$\Pr[\text{Verify}(k, \$k) = 1; k \leftarrow \text{KeyGen}(1^\kappa), \$k \leftarrow \text{Mint}(k)] \geq \mu(\kappa).$$

We say that the scheme is simply correct if μ is inverse-polynomial and we say it is perfectly correct if $\mu = 1$.

⁹Technically, the equivalence holds between strong UCSGs and pure private-key quantum money schemes with a rank-1 projection into the money state, but the rank-1 verification is not an extra assumption, since any pure private-key quantum money scheme can be attached with a rank-1 verification.

- **Unforgeability:** For any polynomials $m, m' : \mathbb{N} \rightarrow \mathbb{N}$ where $m' > m$, and for any non-uniform QPT algorithm \mathcal{A} , there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Pr[\exists S \subseteq [m'(\kappa)], |S| > m, \text{Verify}(k, \$_i) = 1 \forall i \in S; \\ k \leftarrow \text{KeyGen}(1^\kappa), \$_k \leftarrow \text{Mint}(k), \$_{1, \dots, m'(\kappa)} \leftarrow \mathcal{A}^{\text{Verify}(k, \cdot)}(1^\kappa, \$_k^{\otimes m(\kappa)})] \leq \text{negl}(\kappa),$$

where $\$_{1, \dots, m'(\kappa)}$ is a state on $m'(\kappa)$ registers and $\$_i$ denotes its i^{th} register.

We say that the private-key quantum money scheme is statistically unforgeable if unforgeability holds against any arbitrary (and not necessarily QPT) adversary that makes polynomially many queries to the verification oracle.

Further, we say that $(\text{KeyGen}, \text{Mint}, \text{Verify})$ is a private-key quantum money scheme relative to an oracle O if the algorithms KeyGen , Mint and Verify are O -aided and satisfy correctness and unforgeability w.r.t any non-uniform polynomial-time O -aided quantum adversary \mathcal{A} . The oracle-relative notion extends to statistical unforgeability in the natural way.

2.2 On The Complexity of The Reflection Oracle

The reflection unitary $R_\phi = I - 2|\phi\rangle\langle\phi|$ and its controlled variant $R_\phi^c = I - 2|1\rangle|\phi\rangle\langle\phi|\langle 1|$ play a significant role in our proofs. In particular, the latter is a component in the oracle under which we prove our separation results. In this section, we recall a couple of known facts regarding the power of the reflection oracle and its complexity.

We begin with the following folklore application of the controlled reflection oracle about a state $|\phi\rangle$ for projecting onto $|\phi\rangle$.

Proposition 2.7 (Controlled Reflection to Projection). *Let $|\phi\rangle$ be a quantum state and let $R_\phi^c = I - 2|1\rangle|\phi\rangle\langle\phi|\langle 1|$ be the corresponding controlled reflection unitary. There exists an R_ϕ^c -aided quantum algorithm that takes as input a quantum state $|\psi\rangle$, makes a single query to its oracle, outputs $|\phi\rangle$ with probability $|\langle\psi|\phi\rangle|^2$ and otherwise, with probability $1 - |\langle\psi|\phi\rangle|^2$, declares failure.*

Proof. The required R_ϕ^c -aided quantum algorithm on an input quantum state $|\psi\rangle$ in input register W does the following.

1. Initializes a qubit-register T to the state $|-\rangle_T$.
2. Queries the oracle R_ϕ^c with W as the target register and T as the control register.
3. Runs the Hadamard gate H on register T .
4. Measures T in the computational basis.
5. Output the input register W and declare success if the measurement outputs 0 else declare failure.

For the analysis, note that for any input state $|\psi\rangle$, there exists a pure state $|\phi\rangle^\perp$ which is orthogonal to $|\phi\rangle$, and complex numbers a, b such that $|a|^2 + |b|^2 = 1$ and $|\psi\rangle = a|\phi\rangle + b|\phi\rangle^\perp$. Therefore,

$$|\langle\psi|\phi\rangle|^2 = |a|^2, 1 - |\langle\psi|\phi\rangle|^2 = |b|^2.$$

Hence, note that the state on registers W, T after the oracle query is

$$\begin{aligned} R_\phi^c(|\psi\rangle_W \otimes |-\rangle_T) &= a \cdot R_\phi^c(|\phi\rangle_W \otimes |-\rangle_T) + b \cdot R_\phi^c(|\phi\rangle_W^\perp \otimes |-\rangle_T) \\ &= a \cdot |\phi\rangle_W \otimes |+\rangle_T + b \cdot |\phi\rangle_W^\perp \otimes |-\rangle_T. \end{aligned}$$

Therefore the state on registers W, T before the measurement is

$$a \cdot |\phi\rangle_W \otimes |0\rangle_T + b \cdot |\phi\rangle_W^\perp \otimes |1\rangle_T.$$

Hence with probability $|a|^2 = |\langle\psi|\phi\rangle|^2$, the measurement outcome is 0, i.e. the algorithm succeeds, and the resulting post-measured state on W is $|\phi\rangle_W$, and with probability $|b|^2 = 1 - |\langle\psi|\phi\rangle|^2$, the measurement outcome is 1, i.e. the algorithm fails. \square

Next, we recall the following theorem from [JLS18], which states that the reflection oracle about $|\psi\rangle$ may be efficiently emulated given copies of the state $|\psi\rangle$.

Theorem 2.8 (Reflection Emulation [JLS18, Theorem 4]). *Let Q be a quantum oracle. Let $|\psi\rangle$ be a quantum state and let $R_\psi = I - 2|\psi\rangle\langle\psi|$ be the corresponding reflection unitary. Let $|\phi\rangle$ be a state not necessarily independent of $|\psi\rangle$. Let \mathcal{A} be a (Q, R_ψ) -aided quantum circuit that makes q queries to the oracle R_ψ . Then, there exists a Q -aided quantum circuit \mathcal{B} such that, for any $\ell \in \mathbb{N}$,*

$$\text{TD} \left(\mathcal{A}^{Q, R_\psi}(|\phi\rangle) \otimes |\psi\rangle^{\otimes \ell}, \mathcal{B}^Q(|\phi\rangle \otimes |\psi\rangle^{\otimes \ell}) \right) \leq \frac{2q}{\sqrt{\ell+1}}.$$

Further, if \mathcal{A} is of polynomial size then so is \mathcal{B} .

The above theorem is, in fact, slightly stronger than Theorem 4 of [JLS18] in the following two aspects. First, we allow both \mathcal{A} and \mathcal{B} access to an additional oracle Q . This is possible because in the proof of Theorem 4 of [JLS18], the emulation circuit \mathcal{B} uses \mathcal{A} as a black-box. Second, our version of the theorem bounds the distance between the outcome states including the registers where the copies of $|\psi\rangle$ reside. That is, compared to Theorem 4 of [JLS18], we replace $\mathcal{A}^{Q, R_\psi}(|\phi\rangle)$ with $\mathcal{A}^{Q, R_\psi}(|\phi\rangle) \otimes |\psi\rangle^{\otimes \ell}$. This is also possible, because, looking at the proof of Theorem 4 of [JLS18], \mathcal{B} does not disturb $|\psi\rangle^{\otimes \ell}$ when it uses them to emulate R_ψ . We briefly recall the construction of \mathcal{B} from the proof by [JLS18].

Let $S = \sqrt{\ell+1}\mathbb{C}^N$ denote the symmetric subspace over $\ell+1$ registers, where \mathbb{C}^N represents the Hilbert space in which the N -dimensional state $|\psi\rangle$ resides.¹⁰ Let R_S denote the reflection about the symmetric subspace. The circuit \mathcal{B} is defined the same as \mathcal{A} except that any query to the oracle R_ψ is replaced by the application of R_S over the input state (i.e. the register containing the query) and the ℓ registers containing the copies of $|\psi\rangle$.

For completeness, we provide a proof of Theorem 2.8 in Appendix C.

The following is an extension of Theorem 2.8 to the case where multiple reflection oracles are used.

Corollary 2.9. *Let Q be a quantum oracle. Let $|\psi_1\rangle, \dots, |\psi_m\rangle$ be quantum states and let, for $j \in [m]$, $R_{\psi_j} = I - 2|\psi_j\rangle\langle\psi_j|$ be the corresponding reflection unitary. Let $|\phi\rangle$ be a state not necessarily independent of $|\psi_1\rangle, \dots, |\psi_m\rangle$. Let \mathcal{A} be a $(Q, R_{\psi_1}, \dots, R_{\psi_m})$ -aided quantum circuit that makes q queries to its oracles $\{R_{\psi_j}\}_j$ (in total). Then, there exists a Q -aided quantum circuit \mathcal{B} such that, for any $\ell \in \mathbb{N}$,*

$$\begin{aligned} \text{TD} \left(\mathcal{A}^{Q, R_{\psi_1}, \dots, R_{\psi_m}}(|\phi\rangle) \otimes |\psi_1\rangle^{\otimes \ell} \otimes \dots \otimes |\psi_m\rangle^{\otimes \ell}, \right. \\ \left. \mathcal{B}^Q(|\phi\rangle \otimes |\psi_1\rangle^{\otimes \ell} \otimes \dots \otimes |\psi_m\rangle^{\otimes \ell}) \right) \leq \frac{2q}{\sqrt{\ell+1}}. \end{aligned}$$

Further, if \mathcal{A} is of polynomial size then so is \mathcal{B} .

¹⁰The symmetric subspace contains all states that are invariant to a permutation over their registers.

Proof. The corollary is obtained by applying Theorem 2.8 m times to replace the reflection oracles with copies of the corresponding states, one at a time. Let q_j be the number of queries made by \mathcal{A} to the oracle R_{ψ_j} (it holds $\sum_j q_j = q$). We denote by \mathcal{B}_j a quantum algorithm that is given access to $R_{\psi_{j+1}}, \dots, R_{\psi_m}$ and takes $|\psi_1\rangle^{\otimes \ell}, \dots, |\psi_j\rangle^{\otimes \ell}$ as input (besides $|\phi\rangle$). We define \mathcal{B}_0 to behave identically to \mathcal{A} . For $j > 0$, we let \mathcal{B}_j be the algorithm obtained by applying Theorem 2.8 to \mathcal{B}_{j-1} , to replace R_{ψ_j} with the input $|\psi_j\rangle^{\otimes \ell}$, where the input state to \mathcal{B}_{j-1} (denoted by $|\phi\rangle$ in the theorem's statement) is $|\phi\rangle \otimes |\psi_1\rangle^{\otimes \ell} \otimes \dots \otimes |\psi_{j-1}\rangle^{\otimes \ell}$. Due to the theorem, we know that the trace distance between the outcome of \mathcal{B}_j and that of \mathcal{B}_{j-1} is $q_j 2/\sqrt{\ell + 1}$ and, hence, the corollary holds for $\mathcal{B} = \mathcal{B}_m$. \square

3 The Oracles

We begin by defining the unitary oracle under which the separation between OWSGs and QEFID pairs from Theorem 1.1 holds. Our oracle is actually composed of two oracles: First, an oracle O , which captures the hardness required for building a QEFID pair. In fact, we define a distribution \mathcal{O} over oracles, and we show that a ‘‘hard’’ oracle $O \in \mathcal{O}$ exists via a probabilistic argument. The distribution \mathcal{O} will be useful also to establish the separation between private-key quantum money schemes and QEFID pairs (Theorem 1.5). Second, an oracle **QSPACE** [CCS24], which allows simulating any QSPACE computation¹¹ and provides the necessary power to break any OWSG candidate (while preserving the security of our QEFID pair).

We begin by defining a useful special case of the controlled reflection unitary.

Definition 3.1 (Unitary U_W). For any $\lambda \in \mathbb{N}$ and $W \subseteq \{0, 1\}^\lambda \setminus \{0^\lambda\}$, we define the states

$$|W\rangle = \frac{1}{\sqrt{|W|}} \sum_{x \in W} |x\rangle \quad \text{and} \quad |W-\rangle = \frac{1}{\sqrt{2}} (|W\rangle - |0^\lambda\rangle).$$

We define the $(\lambda + 1)$ -qubit unitary U_W as follows

$$U_W = I - 2|1\rangle|W-\rangle\langle W-|\langle 1|.$$

Namely, U_W is the controlled reflection unitary about $|W-\rangle$, i.e. R_{W-}^c .

Note that, R_{W-} maps $|W\rangle$ to $|0^\lambda\rangle$ and $|0^\lambda\rangle$ to $|W\rangle$, and acts as identity on the subspace orthogonal to $\text{Span}(|0^\lambda\rangle, |W\rangle)$.

We now define the distribution \mathcal{O} over oracles O where, roughly speaking, a random oracle $O \leftarrow \mathcal{O}$ applies the unitary U_S over its $(\lambda + 1)$ -qubit input, for a uniformly random subset S of size $2^{\lambda/2}$.

Definition 3.2 (The Oracle \mathcal{O}). We denote by \mathcal{O} the distribution over quantum oracles where

- **Randomness:** A random $O_S \leftarrow \mathcal{O}$ is defined by an ensemble $S = \{S_\lambda\}_{\lambda \in \mathbb{N}}$ where, for any $\lambda \in \mathbb{N}$, S_λ is a uniformly random subset in $\{0, 1\}^\lambda \setminus \{0^\lambda\}$ of size $|S_\lambda| = 2^{\frac{\lambda}{2}}$.
- **Query:** For any $S = \{S_\lambda\}_{\lambda \in \mathbb{N}}$, $\lambda \in \mathbb{N}$ and $(\lambda + 1)$ -qubit input state ρ , we define $O_S(\rho) = U_{S_\lambda} \rho U_{S_\lambda}^\dagger$.

¹¹Note that this QSPACE computation does not mean the classical computation that can solve decision problems that are decided by quantum polynomial-space computing. It means a unitary operation over polynomial number of qubits. Therefore, our **QSPACE** oracle is a quantum oracle that takes a quantum state as input and outputs a quantum state.

A key ingredient in our proofs is the application of the reflection emulation technique from [JLS18] to our oracle \mathcal{O} . Since any $O \in \mathcal{O}$ essentially computes a reflection unitary about the state $|S_\lambda-\rangle$ on any input of length $\lambda + 1$ (where S_λ is the random subset underlying O), reflection emulation (namely Theorem 2.8) tells us that we are able to simulate the oracle given copies of the corresponding state $|S_\lambda-\rangle$. The following proposition follows from Theorem 2.8 and Definition 3.1.

Proposition 3.3 (Emulating \mathcal{O}). *Let $\epsilon > 0$, $m \in \mathbb{N}$ and $\lambda_1, \dots, \lambda_m \in \mathbb{N}$. Let $q \in \mathbb{N}$ and define $t = 2q^2/\epsilon^2 - 1$. For every $i \in [m]$, let S_i be a subset of strings in $\{0, 1\}^{\lambda_i} \setminus \{0^{\lambda_i}\}$, and U_{S_i}, \dots, U_{S_m} be as defined in Definition 3.1 and Q be any other quantum oracle. Let $\mathcal{A}^{Q, U_{S_1}, \dots, U_{S_m}}$ be an oracle-aided quantum algorithm that makes at most q queries in total to the oracles U_{S_1}, \dots, U_{S_m} . Then, there exists a quantum algorithm \mathcal{B} such that, for any input state ρ ,*

$$\text{TD} \left[\left(\mathcal{A}^{Q, U_{S_1}, \dots, U_{S_m}}(\rho) \bigotimes_{i=1}^m |S_{i-}\rangle^{\otimes t} \right), \mathcal{B}^Q \left(\rho, \bigotimes_{i=1}^m |S_{i-}\rangle^{\otimes t} \right) \right] \leq \epsilon.$$

Moreover, the running time of \mathcal{B} is polynomial in that of \mathcal{A} and ℓ .

Lastly, we recall the definition of the QPSPACE oracle from [CCS24].

Definition 3.4 (QPSPACE Oracle [CCS24]). *We define the QPSPACE machine oracle, which we denote by QPSPACE , as follows. The oracle takes as input an ℓ -qubit state ρ , a description of a classical Turing machine M and an integer $t \in \mathbb{N}$. The oracle runs M for t steps to obtain the description of a quantum circuit C that operates on exactly ℓ qubits. If M does not terminate after t steps, or if the output circuit C does not operate on ℓ qubits, the oracle returns \perp . Otherwise, the oracle applies C on ρ and returns the output without measurement.*

We stress that we restrict ourselves to a “unitary world”. That is, our oracles define unitary operations and, additionally, we always provide any quantum algorithm in our oracle-relative world with access to a unitary *and its inverse*. In our case, however, this is w.l.o.g. since the oracle O is equivalent to its inverse and the inverse of QPSPACE can be simulated by a single query to QPSPACE .

Proposition 3.5. *There exists a QPSPACE -aided QPT algorithm which, on any QPSPACE -input (ρ, M, t) where ρ is an ℓ -qubit state, M is the description of a Turing machine and $t \in \mathbb{N}$ (see Definition 3.4), outputs $\text{QPSPACE}^{-1}(\rho, M, t) = C^{-1}(\rho)$, where C is the quantum circuit whose description is output by the machine M after running it for t steps.*

Proof. The algorithm computes the description of the Turing machine M' that acts as follows. M' runs M for t steps to obtain the gates of the circuit C and then inverts them one by one starting from the last gate (without printing or storing all the intermediate gates), to obtain the inverse circuit C^{-1} . Let t' denote the runtime of M' . The algorithm then calls the QPSPACE oracle with (ρ, M', t') .

The runtime of M' is at most $O(t + |C|^2)$ (note that inversion is applied, w.l.o.g., to constant-size gates), where, $|C|$ is the size of C , i.e., the number of gates. Since $|C|$ is obviously bounded by t , it holds that $t' = O(t^2)$ and, in particular, its description as an integer is at most a constant factor by that of t . The description size of M' is larger than that of M by an additive $O(\log |C| + \log t)$. Hence, our simulation algorithm requires a single polynomial-size query. \square

4 Existence of QEFID Pairs

As a first step towards proving Theorem 4.1, we demonstrate that an exponentially-hard QEFID pair (as per Definition 2.1) exists relative to the oracles we defined in Section 3, namely a random $O \leftarrow \mathcal{O}$ and **QPSPACE**.

Theorem 4.1 (QEFID pairs exist under $(\mathcal{O}, \text{QPSPACE})$). *Let the oracles $\mathcal{O}, \text{QPSPACE}$ be as defined in Definitions 3.2 and 3.4, respectively. Then, with probability 1 over the choice of $O \leftarrow \mathcal{O}$, exponentially-hard QEFID pairs, more precisely $(2^{\lambda/50}, 2^{-\lambda/50})$ -QEFID pairs, exist relative to $(O, \text{QPSPACE})$.*

We start by describing our oracle-relative candidate QEFID pair.

Construction 4.2 (QEFID Pair under \mathcal{O}). *We construct an \mathcal{O} -aided QEFID pair (D_0^O, D_1^O) which, given $O \in \mathcal{O}$, acts as follows:*

$D_0^O(1^\lambda)$: *Query O with $|1\rangle|0^\lambda\rangle$ to get $|1\rangle|S_\lambda\rangle$. Measure the second register in the computational basis to obtain a λ -bit string $s \in S_\lambda$. Output it.*

$D_1^O(1^\lambda)$: *Output a uniformly random $u \leftarrow \{0, 1\}^\lambda$.*

It is immediate, by construction, that the QEFID pair from Construction 4.2 is efficiently samplable for any $O \in \mathcal{O}$. It remains, then, to show that it satisfies statistical fairness and computational indistinguishability relative to $(O, \text{QPSPACE})$ (see Definition 2.1). Let us start with the former.

Proposition 4.3 (Statistical Fairness). *For any fixed $O \in \mathcal{O}$, the distributions $D_0^O(1^\lambda)$ and $D_1^O(1^\lambda)$ defined in Construction 4.2, are statistically far.*

Proof. Let $S = \{S_\lambda\}_\lambda$ be the fixed randomness corresponding to the oracle O (see Definition 3.2). Fix $\lambda \in \mathbb{N}$ and let, for any $x \in \{0, 1\}^\lambda$, p_x^0 and p_x^1 denote the probabilities that $D_0^O(1^\lambda)$ and, respectively, $D_1^O(1^\lambda)$ output x . Then,

$$\begin{aligned} \text{SD}(D_0^O(1^\lambda), D_1^O(1^\lambda)) &= \frac{1}{2} \sum_{x \in \{0, 1\}^\lambda} |p_x^0 - p_x^1| \\ &= \frac{1}{2} \sum_{x \in S} |p_x^0 - p_x^1| + \frac{1}{2} \sum_{x \notin S} |p_x^0 - p_x^1| = \frac{1}{2} \sum_{x \in S} \left| \frac{1}{2^{\lambda/2}} - \frac{1}{2^\lambda} \right| + \frac{1}{2} \sum_{x \notin S} \frac{1}{2^\lambda} \\ &= \frac{1}{2} \cdot 2^{\lambda/2} \cdot \frac{2^{\lambda/2} - 1}{2^\lambda} + \frac{1}{2} \cdot (2^\lambda - 2^{\lambda/2}) \cdot \frac{1}{2^\lambda} = 1 - \frac{1}{2^{\lambda/2}}, \end{aligned}$$

which is overwhelming in the security parameter λ . □

Next, we argue computational indistinguishability of our construction through the following lemma.

Lemma 4.4 (Computational Indistinguishability). *With probability 1 over the choice of $O \leftarrow \mathcal{O}$, for any non-uniform $(O, \text{QPSPACE})$ -aided adversary \mathcal{A} , that on any security parameter λ makes at most $O(2^{\lambda/50})$ calls to the oracle O (and is otherwise unbounded), it holds that*

$$\left| \Pr_{z \leftarrow D_0(1^\lambda)} [\mathcal{A}^{O, \text{QPSPACE}}(1^\lambda, z) = 1] - \Pr_{z \leftarrow D_1(1^\lambda)} [\mathcal{A}^{O, \text{QPSPACE}}(1^\lambda, z) = 1] \right| = O(2^{-\lambda/50}),$$

for any $\lambda \in \mathbb{N}$.

Note that we show computational indistinguishability even against adversaries that are unbounded in their runtime or queries to the **QPSpace** oracle. While this is stronger than what we need for the oracle separation between OWSGs and QEFID pairs, it will become useful for our separation concerning private-key quantum money schemes (Theorem 1.5).

The proof of Theorem 4.1 follows immediately by combining Proposition 4.3 and Lemma 4.4. In what comes next, we prove Lemma 4.4.

By reflection emulation (see Proposition 3.3), we may emulate any $(O, \mathbf{QPSpace})$ -aided adversary by an adversary that does not have access to the oracle yet requires copies of the state $|S_\lambda-\rangle$, where S_λ is the random subset underlying O . This reduces our task to showing computational indistinguishability against such adversaries. By replacing access to the oracle O with copies of the state $|S_\lambda-\rangle$, our goal becomes showing that it is hard to break our QEFID given such copies. At the core of such an argument, then, is the following statistical lemma.

Lemma 4.5. *Let $\lambda \in \mathbb{N}$ and let $S, W \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$ be two uniformly random subsets of size $2^{\frac{\lambda}{2}}$, $s \leftarrow S$, and $u \leftarrow \{0, 1\}^\lambda$. Then, it holds that*

$$\text{TD} \left[\left\{ |S-\rangle^{\otimes t} \otimes |s\rangle \right\}_{S,s}, \left\{ |W-\rangle^{\otimes t} \otimes |u\rangle \right\}_{W,u} \right] < 2^{-\frac{\lambda}{2}+1} + \sqrt{t \cdot 2^{-\frac{\lambda}{2}}}.$$

Proof. Observe that the distribution on the r.h.s. in the trace distance satisfies

$$\text{TD} \left[\left\{ |W-\rangle^{\otimes t} \otimes |u\rangle \right\}_{W,u}, \left\{ |W'-\rangle^{\otimes t} \otimes |u'\rangle \right\}_{W',u'} \right] < 2^{-\frac{\lambda}{2}+1}, \quad (2)$$

where r.h.s distribution (W', u') is sampled by first sampling u' uniformly at random from $\{0, 1\}^\lambda$, and then sampling W' as a uniformly random subset over $\{0, 1\}^\lambda \setminus \{0^\lambda\}$ of size $2^{\frac{\lambda}{2}}$ conditioned on $u' \notin W'$. Equation (2) holds since a random subset of size $2^{\frac{\lambda}{2}}$ does not contain a fixed string u' except with probability $2^{\lambda/2}/(2^\lambda - 1) < 2^{1-\lambda/2}$. It is enough to bound, then, the trace distance between $\{|S-\rangle \otimes |s\rangle\}_{S,s}$ and $\{|W'-\rangle \otimes |u'\rangle\}_{W',u'}$.

Next, observe that the l.h.s. distribution in the lemma's statement can be sampled as follows: Sample a uniform $s \leftarrow \{0, 1\}^\lambda$, then a uniform subset $S \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$ of size $2^{\frac{\lambda}{2}}$ that contains s . On the other hand, u' and W' may be sampled by choosing u' uniformly, then a uniformly random subset W'' containing u' , and then defining $W' = W'' \cup \{w\} \setminus \{u'\}$ for a uniformly random $w \leftarrow \{0, 1\}^\lambda \setminus W''$.

Note that such W'', u' distribute identically to S, s . Hence, by an averaging argument, it is enough to show that for any fixed $s \in \{0, 1\}^\lambda$ and $S \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$ such that $s \in S$ and any fixed $w \in \{0, 1\}^\lambda \setminus S$, the following holds

$$\text{TD}(|S-\rangle^{\otimes t} \otimes |s\rangle, |W'-\rangle^{\otimes t} \otimes |s\rangle) = \text{TD}(|S-\rangle^{\otimes t}, |W'-\rangle^{\otimes t}),$$

where $W' = S \cup \{w\} \setminus \{s\}$. Since $|S-\rangle^{\otimes t}$ and $|W'-\rangle^{\otimes t}$ are pure states, the trace distance is simply

$$\begin{aligned} \sqrt{1 - \langle S - | W' - \rangle^{2t}} &= \sqrt{1 - \left(\frac{1}{2} + \frac{1}{2} \cdot \sqrt{1 - \frac{1}{2^{\frac{\lambda}{2}}}} \right)^{2t}} \\ &\leq \sqrt{1 - \left(\frac{1}{2} + \frac{1}{2} \cdot \left(1 - \frac{1}{2^{\frac{\lambda}{2}+1}} \right) \right)^{2t}} = \sqrt{1 - \left(1 - \frac{1}{2^{\frac{\lambda}{2}+2}} \right)^{2t}} \\ &\leq \sqrt{1 - \left(1 - \frac{2t}{2^{\frac{\lambda}{2}+2}} \right)} = \sqrt{\frac{t}{2^{\frac{\lambda}{2}+1}}}. \end{aligned} \quad (3)$$

The proof is complete by combining Equations (2) and (3). \square

Putting Proposition 3.3 and Lemma 4.5 together, we obtain the following proposition.

Proposition 4.6. *For any $(\mathcal{O}, \mathbf{QPSPACE})$ -aided quantum algorithm \mathcal{A} that makes at most $O(2^{\lambda/50})$ queries to \mathcal{O} , it holds that*

$$\mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] \right] \leq 2^{-2\lambda/25}.$$

Proof. Suppose not, and there exists an oracle-aided quantum algorithm \mathcal{A} making $r(\lambda) \in O(2^{\lambda/50})$ queries to the oracle U_{S_λ} ¹², such that

$$\mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] \right] > 2^{-\frac{2\lambda}{25}}.$$

Then, by Proposition 3.3 and letting $t(\lambda) = 2r(\lambda)^2 \cdot (4 \cdot 2^{2\lambda/25})^2$, there exists an algorithm \mathcal{B} that only queries $\mathbf{QPSPACE}$ and obtains the following distinguishing advantage

$$\begin{aligned} & \mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1] \right] \\ & > \frac{1}{2^{\frac{2\lambda}{25}}} - \frac{1}{4 \cdot 2^{\frac{2\lambda}{25}}} - \frac{1}{4 \cdot 2^{\frac{2\lambda}{25}}} = \frac{1}{2 \cdot 2^{\frac{2\lambda}{25}}} = 2^{-\frac{2\lambda}{25}-1} \in \Omega\left(2^{-\frac{2\lambda}{25}}\right). \end{aligned} \quad (4)$$

which is the same as saying for two uniformly random subsets of size $2^{\frac{\lambda}{2}}$ $S, W \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$, and uniformly random strings $s \leftarrow S$, and $u \leftarrow \{0, 1\}^\lambda$,

$$\mathbb{E}_{S,s} \left[\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S - \rangle^{\otimes t} \otimes |s\rangle) = 1] \right] - \mathbb{E}_{W,u} \left[\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |W - \rangle^{\otimes t} \otimes |u\rangle) = 1] \right] \in \Omega\left(2^{-\frac{2\lambda}{25}}\right). \quad (5)$$

Equation (5) follows from the observation that

$\mathbb{E}_{\mathcal{O}} [\Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1]]$ can be written as $\mathbb{E}_{S,s} [\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S - \rangle^{\otimes t} \otimes |s\rangle) = 1]]$ and $\mathbb{E}_{\mathcal{O}} [\Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1]]$ can be written as $\mathbb{E}_{W,u} [\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |W - \rangle^{\otimes t} \otimes |u\rangle) = 1]]$ where S, W, s, u are distributed as mentioned above, and hence by linearity of expectation,

$$\mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S_\lambda - \rangle^{\otimes t}, s) = 1] \right] \quad (6)$$

$$= \mathbb{E}_{S,s} \left[\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |S - \rangle^{\otimes t} \otimes |s\rangle) = 1] \right] - \mathbb{E}_{W,u} \left[\Pr_{\mathcal{B}} [\mathcal{B}^{\mathbf{QPSPACE}}(1^\lambda, |W - \rangle^{\otimes t} \otimes |u\rangle) = 1] \right]. \quad (7)$$

However, note that

$$t(\lambda) = 2r(\lambda)^2 (4 \cdot 2^{\frac{2\lambda}{25}})^2 \in O(2^{\frac{2\lambda}{50}} \cdot 2^{\frac{4\lambda}{25}}) = O(2^{\frac{9\lambda}{50}}). \quad (8)$$

¹²We can assume without loss of generality that \mathcal{A} only queries U_{S_λ} and not the oracles $U_{S_{\lambda'}}$ for $\lambda' \neq \lambda$ since the $U_{S_{\lambda'}}$ for $\lambda' \neq \lambda$ are sampled independently from the current computational indistinguishability game.

Combining Equation (8) with Lemma 4.5, we conclude that for some constant c , and λ large enough

$$\mathbb{E}_{S,s} \left[\Pr_{\mathcal{B}}[\mathcal{B}^{\text{QSPACE}}(1^\lambda, |S-\rangle^{\otimes t} \otimes |s\rangle) = 1] \right] - \mathbb{E}_{W,u} \left[\Pr_{\mathcal{B}}[\mathcal{B}^{\text{QSPACE}}(1^\lambda, |W-\rangle^{\otimes t} \otimes |u\rangle) = 1] \right] \quad (9)$$

$$\leq 2^{-\frac{\lambda}{2}+1} + \sqrt{c \frac{2^{\frac{9\lambda}{50}}}{2^{\frac{\lambda}{2}+1}}} = 2^{-\frac{\lambda}{2}+1} + \sqrt{c} 2^{-\frac{4\lambda}{25}} \leq (1 + \sqrt{c}) 2^{-\frac{4\lambda}{25}} \in o(2^{-\frac{2\lambda}{25}}), \quad (10)$$

which contradicts the bound in Equation (5), thereby completing the proof. \square

While it might seem like Proposition 4.6 implies computational indistinguishability of our QEFID construction for a random oracle $O \leftarrow \mathcal{O}$, this is not immediately the case since we have bounded the positive gap between the probability that the distinguisher's output is 1 at input from D_0 compared to an input from D_1 . Indistinguishability, however, requires bounding the absolute value between the two probabilities.

It is a well-known fact that such a ‘‘positive-gap’’ indistinguishability implies standard absolute-value indistinguishability [Yao82, BG11] (with some loss in advantage). This has been further extended to the case of a quantum oracle-aided distinguisher (where the gap is guaranteed in expectation over a random oracle, just as we require).

Lemma 4.7 (Absolute-Gap to Positive-Gap Distinguisher). *Let \mathcal{O} be a distribution over oracles. Let $D_0^{\mathcal{O}}$ and $D_1^{\mathcal{O}}$ be two \mathcal{O} -aided distributions over $\{0, 1\}^\lambda$ with corresponding sampling algorithms. Let \mathcal{A} be an \mathcal{O} -aided quantum algorithm such that*

$$\mathbb{E}_{O \leftarrow \mathcal{O}} \left[\left| \Pr_{x \leftarrow D_0} [\mathcal{A}(x) = 1] - \Pr_{x \leftarrow D_1} [\mathcal{A}(x) = 1] \right| \right] = \delta.$$

Then, there exists an \mathcal{O} -aided quantum algorithm \mathcal{B} such that

$$\mathbb{E}_{O \leftarrow \mathcal{O}} \left[\Pr_{x \leftarrow D_0} [\mathcal{B}(x) = 1] - \Pr_{x \leftarrow D_1} [\mathcal{B}(x) = 1] \right] \geq \delta^2.$$

The runtime and query complexity of \mathcal{B} is twice that of \mathcal{A} in addition to that of the sampling algorithms of D_0 and D_1 .

For completeness, we attach a proof to the lemma in Appendix B. As a corollary, we obtain the following.

Corollary 4.8. *For any $(\mathcal{O}, \text{QSPACE})$ -aided quantum algorithm \mathcal{A} that makes at most $O(2^{\lambda/50})$ queries to \mathcal{O} , it holds that*

$$\mathbb{E}_{O \leftarrow \mathcal{O}} \left[\left| \Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \text{QSPACE}}(1^\lambda, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \text{QSPACE}}(1^\lambda, s) = 1] \right| \right] \leq 2^{-\lambda/25}.$$

With Corollary 4.8 in hand, we proceed to complete the proof of Lemma 4.4.

Fix an $(\mathcal{O}, \text{QSPACE})$ -aided quantum adversary \mathcal{A} that makes at most $O(2^{\lambda/50})$ queries to \mathcal{O} . By Corollary 4.8 and Markov inequality, we conclude that

$$\Pr_{O \leftarrow \mathcal{O}} \left[\left| \Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \text{QSPACE}}(1^\lambda, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \text{QSPACE}}(1^\lambda, s) = 1] \right| \geq 2^{-\lambda/50} \right] \leq 2^{-\lambda/50}.$$

Since $\sum_{\lambda} 2^{-\lambda/50}$ converges, by Borel-Cantelli Lemma we have that, with probability 1 over the choice of $O \leftarrow \mathcal{O}$, it holds that

$$\left| \Pr_{s \leftarrow D_0^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] - \Pr_{s \leftarrow D_1^{\mathcal{O}}(1^\lambda)} [\mathcal{A}^{\mathcal{O}, \mathbf{QPSPACE}}(1^\lambda, s) = 1] \right| \leq 2^{-\lambda/50}, \quad (11)$$

except for finitely many $\lambda \in \mathbb{N}$. Since there are only countably many such quantum algorithms \mathcal{A} that make at most $O(2^{\lambda/50})$ queries to \mathcal{O} , we conclude that with probability 1 over the oracles $(\mathcal{O}, \mathbf{QPSPACE})$, it holds that: for every quantum algorithm \mathcal{A} making at most $O(2^{\lambda/50})$ queries to \mathcal{O} , Equation (11) holds, which completes the proof of Lemma 4.4.

5 Impossibility of OWSGs

In this section, we show that, for any fixed $O \in \mathcal{O}$, OWSGs (as defined in Definition 2.2) do not exist in the presence of the oracles $(O, \mathbf{QPSPACE})$. Together with Theorem 4.1, this completes the proof of our first main result from Theorem 1.1, namely, the oracle separation of OWSGs from QEFID pairs.

Theorem 5.1. *For any $O \in \mathcal{O}$, OWSGs do not exist relative to the oracles $(O, \mathbf{QPSPACE})$.*

5.1 Gentle Search for QPSPACE-aided POVMs

A key ingredient for breaking any oracle-relative OWSG candidate is the *gentle search* procedure [Aar19]. In our context, gentle search allows, given a quantum state and a collection of verification keys, to identify a key under which the state is accepted (by an a priori-fixed verification algorithm).

While various algorithms for standard gentle search exist in the literature [Aar19, WB24], we want to additionally allow the verification algorithm to have access to the $\mathbf{QPSPACE}$ oracle. To this end, via an observation made in [CCS24], we adapt the algorithm from [Aar19], in particular its *OR-tester component* [HLM17], to work also when the verification algorithm has polynomially-bounded access to $\mathbf{QPSPACE}$. Overall, we obtain the following general gentle search algorithm for $\mathbf{QPSPACE}$ -aided POVMs.

Lemma 5.2 (Gentle Search via QPSPACE Machine). *Let K be a finite set of strings and let $\{\Pi_k^{\mathbf{QPSPACE}}\}_{k \in K}$ be a family of $\mathbf{QPSPACE}$ -aided binary-valued POVMs, indexed by elements in K , each of which makes polynomially many queries to its oracle. Suppose $|\psi\rangle$ is a state such that there exist a real $c > 0$ and a key $k \in K$ for which $\text{Tr}(\Pi_k^{\mathbf{QPSPACE}} |\psi\rangle \langle \psi|) \geq c$. Let $\epsilon, \delta > 0$. Then, there exists a polynomial-time $\mathbf{QPSPACE}$ -aided quantum algorithm, which we denote by *GentleSearch*, that takes as input $|\psi\rangle^{\otimes t}$, for $t \in O\left((\log^4 |K| \log \log |K| + \log(1/\delta))/\epsilon^2\right)$, makes $\log(|K|)$ queries to its oracle and outputs a key $k' \in K$ such that, with probability at least $1 - \delta$,*

$$\text{Tr}(\Pi_{k'}^{\mathbf{QPSPACE}} |\psi\rangle \langle \psi|) \geq c - \epsilon.$$

Proof. Our algorithm applies the gentle search procedure devised by Aaronson [Aar19, Lemma 15]. Assuming, w.l.o.g., that $|K|$ is a power of 2, Aaronson's gentle search is a binary-search type of algorithm, that, at every iteration, divides the given collection of POVM elements into two halves and applies the quantum OR-tester, described below, on each to identify the half consisting of the target key, and proceeds recursively therein.

The OR-tester [HLM17, Corollary 3.1] is a test which, given a collection of binary-valued POVMs ζ_1, \dots, ζ_m and a state ρ (one copy thereof) distinguishes between the following two cases:

1. There exists $i \in [m]$ such that $\text{Tr}(\zeta_i \rho) \geq \gamma_1$, and
2. $\mathbb{E}_{i \leftarrow [m]}[\text{Tr}(\zeta_i \rho)] < \gamma_2$,

for $\gamma_1 > 1/2$ and $\gamma_1 > \gamma_2$, the choice of which affects the complexity of the test and its advantage.

Using such a tester, Aaronson [Aar19, Lemma 14] demonstrates given a collection of binary-valued POVMs $\tilde{\zeta}_1, \dots, \tilde{\zeta}_m$ and a state ρ , how to construct a tester that distinguishes between the following for any $0 \leq c \leq 1$ and $\epsilon > 0$,

- (i) There exists $i \in [m]$ such that $\text{Tr}(\tilde{\zeta}_i \rho) \geq c$, and
- (ii) $\forall i \in [m], \text{Tr}(\tilde{\zeta}_i \rho) < c - \epsilon$.

Specifically, he considers the amplified binary-valued POVMs $\{\zeta_i\}_i$ acting on $\rho^{\otimes \ell}$ for $\ell := O(\frac{\log(m)}{\epsilon^2})$ such that¹³ for the amplified POVMs and the state $\rho^{\otimes \ell}$, the OR-tester condition holds with $\gamma_1 = 1 - \frac{1}{m}$ in Item 1 and $\gamma_2 = \frac{1}{m}$ in Item 2. Then, he runs the OR-tester of [HLM17] for the POVMs ζ_1, \dots, ζ_m for the state $\rho^{\otimes \ell}$, and further amplifies it such that his test is able to distinguish Item (i) and Item (ii) with success probability $1 - \eta$ using a total of $O(\log(1/\eta) \log^3 m / \epsilon^2)$ copies of ρ . Evidently, such a test with $\eta = \delta \cdot \log |K|$ suffices to instantiate the binary search described above, to obtain overall success probability $1 - \delta$ with t copies of the input state.

Most importantly to us is the fact that the only place in Aaronson's gentle search where the POVMs are used is the OR-tester component. The OR-tester by Harrow et al. [HLM17] applies a variant of Marriott-Watrous gap-amplification procedure [MW05] to implement the POVM $\Pi = \sum_k \Pi_k$. Specifically, it consists of applying the projective measurements $\Lambda_0 = \sum_k \Pi_k \otimes (Q |k\rangle\langle k| Q^{-1})$ and $\Lambda_1 = I \otimes |0\rangle\langle 0|$ on the input state, in an alternating fashion, where Q denotes QFT over $\mathbb{Z}_{|K|}$.

As noted by Chen et al. [CCS24, Remark 5.3], when the POVM elements $\{\Pi_k\}_k$ are projectors implementable by polynomial-query **QSPACE**-aided quantum circuits in polynomial space, then also the POVM Λ_0 may be implemented by a polynomial-query **QSPACE**-aided quantum circuit. We can generalize the remark ([CCS24, Remark 5.3]) for arbitrary POVM elements $\{\Pi_k\}_k$ implementable by polynomial-query **QSPACE**-aided quantum circuits in polynomial space. This is because every polynomial space implementable oracle-aided binary-valued POVM and in particular Π_k for every k , can be implemented using a polynomial space implementable oracle-aided unitary U_k acting on some $n + r$ qubits where the last r qubits are additional ancillae initialized to zero (where $r \in \text{poly}$ is an upper bound on the polynomial space needed to implement $\{\Pi_k\}_k$), followed by measuring the last qubit in the computational basis, and then accept if the outcome is 1 else fail. Clearly, $\{U_k^\perp (I \otimes |1\rangle\langle 1|) U_k\}_k$ has the same acceptance statistics for a state $\rho \otimes |0^{\otimes r}\rangle\langle 0^{\otimes r}|$ as that of $\{\Pi_k\}_k$ on ρ , because

$$\text{Tr}[\Pi_k \rho] = \text{Tr}[(I \otimes |1\rangle\langle 1|) U_k (\rho \otimes |0^{\otimes r}\rangle\langle 0^{\otimes r}|) U_k^\perp] = \text{Tr}[U_k^\perp (I \otimes |1\rangle\langle 1|) U_k (\rho \otimes |0^{\otimes r}\rangle\langle 0^{\otimes r}|)], \quad (12)$$

and therefore by [CCS24, Remark 5.3], it is possible to implement the POVM Λ_0 by a polynomial-query **QSPACE**-aided quantum circuit.

To conclude, it is possible to perform the gentle search algorithm from [Aar19] in our setting, where the POVMs make a polynomially-bounded access to **QSPACE**. Its correctness follows by the analysis done in [HLM17, Aar19] as outlined above. For further details, we refer the reader to the aforementioned works. \square

¹³ ζ_i runs $\tilde{\zeta}_i$ on ρ ℓ times and accepts if ζ_i outputs accept at least $c - \frac{\epsilon}{2}$ times.

5.2 Proof of Theorem 5.1

Fix $O \in \mathcal{O}$ and let $(\mathcal{G}, \mathcal{V})$ be a OWSG candidate relative to (O, \mathbf{QSPACE}) (see Definition 2.2). We show that the fact that \mathcal{V} is a QPT (in particular, makes a polynomial number of queries to its oracles) necessarily implies an attack against the one-wayness of the candidate.

Our attack is based on the gentle search algorithm as implied by Lemma 5.2; Given (many copies of) a state ϕ_{k^*} corresponding to a key k^* and the collection of POVM elements $\{\Pi_k\}_k$, where Π_k corresponds to the event of \mathcal{V} accepting with key k , it is possible with access to the \mathbf{QSPACE} oracle to identify a key k such that (ϕ_{k^*}, k) is accepted by \mathcal{V} with good probability.

The main obstacle in this outline is that the algorithm \mathcal{V} is oracle-aided. While Lemma 5.2 allows the POVMs to be \mathbf{QSPACE} -aided, the oracle O remains an issue. Luckily, due to Proposition 3.3, we know that polynomially-many calls to O with queries of any length $\lambda \in \mathbb{N}$ can be emulated by polynomially-many copies of the state $|S_\lambda-\rangle$, where $S = \{S_\lambda\}$ are the random subsets underlying the oracle O . In particular, we can simulate \mathcal{V} using such copies for any λ w.r.t. which a query is made by \mathcal{V} , without accessing O .¹⁴

Let $q := q(\kappa)$ be the polynomial bound on the query complexity of \mathcal{V} and let $\bar{\lambda} := \bar{\lambda}(\kappa)$ be the polynomial bound on the length of these queries. That is, for any $\kappa \in \mathbb{N}$ and any $k \in \{0, 1\}^\kappa$, $\mathcal{V}(k, \cdot)$ always makes at most $q(\kappa)$ queries to O , where each is of length at most $\bar{\lambda}(\kappa)$. Then, by Proposition 3.3, there exists a polynomial-query \mathbf{QSPACE} -aided algorithm $\tilde{\mathcal{V}}$ such that, for any $k \in \{0, 1\}^\kappa$ and state ϕ ,

$$\left| \Pr[\mathcal{V}^{O, \mathbf{QSPACE}}(k, \phi) = 1] - \Pr[\tilde{\mathcal{V}}^{\mathbf{QSPACE}}(k, \phi \otimes \bigotimes_{\lambda=1}^{\bar{\lambda}} |S_\lambda-\rangle^{\otimes 2q^2\kappa^2}) = 1] \right| \leq 1/\kappa. \quad (13)$$

Given the above, to apply gentle search w.r.t. the POVMs defined by $\tilde{\mathcal{V}}$ we no longer need access to the oracle O . Instead, we must generate many copies of $|S_\lambda-\rangle$. To that end, we use the fact that O contains the controlled reflection about $|S_\lambda-\rangle$ (see Definition 3.2); due to Proposition 2.7, there exists a single-query algorithm which, on input $|0^\lambda\rangle$ and access to O , outputs $|S_\lambda-\rangle$ with probability $|\langle 0|S_\lambda-\rangle|^2 = 1/2$ and otherwise declares failure. By repeating such a procedure for κ iterations or until success, we get the following corollary.

Corollary 5.3. *For any $\kappa \in \mathbb{N}$, there exists an \mathcal{O} -aided quantum algorithm E that makes κ queries on the input 1^λ , such that, for any $O \in \mathcal{O}$ and $\lambda \in \mathbb{N}$,*

$$\Pr[E^O(1^\lambda) = |S_\lambda-\rangle] \geq 1 - 2^{-\kappa},$$

where $S_\lambda \subset \{0, 1\}^\lambda \setminus \{0^\lambda\}$ is the subset underlying the oracle O (see Definition 3.2).

We are now prepared to describe the attack on the OWSG candidate that requires $t(\kappa) \in O(\kappa^2 \log \kappa)$ copies of the challenge state.

$\mathcal{A}^{O, \mathbf{QSPACE}}(1^\kappa, \phi^{\otimes t})$:

1. For $\lambda = 1, \dots, \bar{\lambda}$, run $E^O(1^\lambda)$ (from Corollary 5.3) $2q^2\kappa^2t$ times to obtain $|S_\lambda-\rangle^{\otimes 2q^2\kappa^2t}$.
2. Run \mathbf{QSPACE} -aided gentle search, namely the algorithm *GentleSearch* from Lemma 5.2, with parameters $\epsilon = \delta = 1/2$, POVM elements $\{\Pi_k\}_{k \in \{0, 1\}^\kappa}$ (where Π_k corresponds the acceptance of $\tilde{\mathcal{V}}^{\mathbf{QSPACE}}(k, \cdot)$), and input $\phi \otimes \left(\bigotimes_{\lambda \in [\bar{\lambda}]} |S_\lambda-\rangle^{\otimes 2q^2\kappa^2t} \right)$ (of which we have t copies – this suffices due to Lemma 5.2), and output its outcome.

¹⁴Note that we use κ to denote the security parameter of the OWSG we aim to break. λ is used to denote the security parameter w.r.t. which our attack invokes its queries to the oracle O . While polynomially-related, κ and λ are not necessarily equal.

By Corollary 5.3, the first step succeeds with probability at least $1 - 2q^2\kappa^2\bar{\lambda}/2^\kappa$. Now, assuming ϕ was sampled by the OWSG under key $k^* \in \{0, 1\}^\kappa$, it holds by correctness that

$$\Pr[\mathcal{V}^{\mathcal{O}, \text{QPSpace}}(k^*, \phi) = 1] \geq 1 - \text{negl}(\kappa),$$

for some negligible function negl , and hence, by Equation (13),

$$\Pr[\tilde{\mathcal{V}}^{\text{QPSpace}}(k^*, \phi \otimes \bigotimes_{\lambda=1}^{\bar{\lambda}} |S_{\lambda-}\rangle^{\otimes 2q^2\kappa^2}) = 1] \geq 1 - 1/\kappa - \text{negl}(\kappa).$$

Given the above guarantee regarding ϕ , Lemma 5.2 implies that the second step finds a $k' \in \{0, 1\}^\kappa$ such that

$$\Pr[\tilde{\mathcal{V}}^{\text{QPSpace}}(k', \phi \otimes \bigotimes_{\lambda=1}^{\bar{\lambda}} |S_{\lambda-}\rangle^{\otimes 2q^2\kappa^2}) = 1] \geq 1/2 - 1/\kappa - \text{negl}(\kappa) > 1/3$$

and, consequently, $\Pr[\mathcal{V}^{\mathcal{O}, \text{QPSpace}}(k', \phi) = 1] \geq 1/3 - 1/\kappa$. This concludes the attack on the OWSG candidate successful and completes the proof of Theorem 5.1.

Remark 5.4. While we describe an attack against the OWSG that succeeds with probability $1/3 - 1/\kappa$ for security parameter κ , this is merely for simplicity of exposition. A more careful tuning of the parameters results in an attack that succeeds with probability $1 - \Theta(1/\kappa)$.

6 Further Implications of Theorem 1.1

Firstly, observe that EV-OWPuzzs (see Definition 2.3) are essentially a classical state version of OWSGs, up to a change of syntax for the sampling the state/puzzle. It was shown in [CGG24] that these syntaxes are equivalent, and hence EV-OWPuzzs imply OWSGs.

Theorem 6.1 (Implicit in [CGG24, Theorem 32, section 10]). *EV-OWPuzzs (see Definition 2.3) imply OWSGs (see Definition 2.2) in a black-box manner.*

Combining Theorem 6.1 with Theorems 4.1 and 5.1, we get the following corollary.

Corollary 6.2 (Restatement of Corollary 1.3). *There exists a unitary oracle relative to which QEFID pairs (see Definition 2.1) exist, but EV-OWPuzzs do not (see Definition 2.3).*

It was also shown in [CGG24] that QEFID pairs imply OWPuzzs.

Theorem 6.3 (Implicit in [CGG24, Lemma 8]). *QEFID pairs (see Definition 2.3) imply OWPuzzs (see Definition 2.2) in a black-box manner.*

Combining Theorem 6.3 with Theorems 4.1 and 5.1, we get the following corollary.

Corollary 6.4 (Restatement of Corollary 1.4). *There exists a unitary oracle relative to which OWPuzzs (see Definition 2.3) exist, but OWSGs (see Definition 2.3) do not.*

Next, we move deeper into Microcrypt, and start with the following simple observation.

Theorem 6.5. *Strongly unclonable state generators (strong-UCSGs) imply (efficiently verifiable) OWSGs.*

Proof. Let \mathcal{G} be the QPT generation algorithm of an UCSG with key space $\mathcal{K} = \{\mathcal{K}_\kappa\}_\kappa$ (see Definition 2.4). Let \mathcal{G}^κ denote the algorithm which, on input a key $k \in \mathcal{K}_\kappa$, for any $\kappa \in \mathbb{N}$, invokes $\mathcal{G}(k)$ κ times and outputs the tensor of all κ output states. Let \mathcal{V} be the algorithm that takes as input a key $k \in \mathcal{K}_\kappa$ and a state ψ , invokes $\mathcal{G}^\kappa(k)$ to obtain $|\phi_k\rangle^{\otimes \kappa}$, and applies a swap test over each of the copies of $|\phi_k\rangle$ and the corresponding register in ψ . \mathcal{V} outputs 1 if all tests succeed.

We argue that the pair $(\mathcal{G}^\kappa, \mathcal{V})$ constitutes an efficiently verifiable OWSG. First, correctness holds since, for any $k \in \mathcal{K}_\kappa$, $\Pr[\mathcal{V}(k, \mathcal{G}^\kappa(k)) = 1] = (\frac{1}{2}(1 + |\langle \phi_k | \phi_k \rangle|^2))^\kappa = 1$, where $|\phi_k\rangle = \mathcal{G}(k)$. For security, let \mathcal{A} be an adversary that breaks the one-wayness of $(\mathcal{G}^\kappa, \mathcal{V})$. That is, there exists a polynomial t and a non-negligible function η such that, for a random key $k \leftarrow \mathcal{K}_\kappa$ and the corresponding state $|\phi_k\rangle = \mathcal{G}(k)$, it holds that $\mathcal{A}(1^\kappa, |\phi_k\rangle^{\otimes \kappa \cdot t})$ outputs a k' such that $\mathcal{V}(k', |\phi_k\rangle^{\otimes \kappa}) = 1$ with probability $\eta := \eta(\kappa)$. The latter necessarily implies that, for infinitely many κ , $|\langle \phi_{k'} | \phi_k \rangle|^2 > 0.8$ with probability at least $\eta/2$ since, otherwise,

$$\Pr[\mathcal{V}(k', |\phi_k\rangle^{\otimes \kappa}) = 1] < \eta/2 + (1 - \eta/2) \left(\frac{1}{2}(1 + |\langle \phi_{k'} | \phi_k \rangle|^2) \right)^\kappa < \eta/2 + 0.9^\kappa(1 - \eta/2) < \eta$$

for all but finitely many κ .

Given the above, we derive an attack against the presumed unclonable state family. The attack, which we recall has access to the controlled reflection oracle $R_{\phi_k}^c$, takes as input security parameter 1^κ and $\kappa \cdot t$ copies of the state $|\phi_k\rangle$ and acts as follows:

1. Run $\mathcal{A}(1^\kappa, |\phi_k\rangle^{\otimes \kappa \cdot t})$ to obtain a key k' .
2. Run the following for at most $\kappa^2 \cdot t$ iterations or until $\kappa \cdot t + 1$ iterations succeed:
 - 2.1. Generate the state $|\phi_{k'}\rangle \leftarrow \mathcal{G}(k')$.
 - 2.2. Invoke the algorithm from Proposition 2.7 with input $|\phi_{k'}\rangle$ using the oracle $R_{\phi_k}^c$. If the algorithm fails, abort. Otherwise, the output state is $|\phi_k\rangle$.
3. Output the $\kappa \cdot t + 1$ copies of $|\phi_k\rangle$ generated by the loop in the previous step.

We now analyze the success probability of our proposed attack. First, recall that for infinitely many κ , it holds $|\langle \phi_{k'} | \phi_k \rangle|^2 > 0.8$ with probability at least $\eta/2$. Assuming this event occurs, due to Proposition 2.7, every iteration of step 2.2. in the attack is successful with probability at least 0.8. Hence, by a standard Chernoff bound, the probability that $\kappa \cdot t + 1$ attempts among the $\kappa^2 \cdot t$ iterations succeed is at least $1 - 2^{-\kappa}$. This concludes the proof. □

Note that in the proof of Theorem 6.5, it was crucial that the adversary \mathcal{A} has access to the controlled-reflection about the state, i.e., the state family is a *strongly* unclonable family. Next, we use our emulation result for arbitrary reflections to show that this additional access to the controlled-reflection does not add any strength.

Theorem 6.6. *Any UCSG is also strongly unclonable.*

Proof. Let $\Phi = \{|\phi_k\rangle\}_{k \in \mathcal{K}}$ be an unclonable state generator. Assume towards contradiction that it is not strongly unclonable. Then, there exist an oracle-aided QPT algorithm \mathcal{A} and polynomial $m := m(\kappa)$, such that

$$\mathbb{E}_{k \leftarrow \mathcal{K}_\kappa} \left[\text{Tr} \left(|\phi_k\rangle \langle \phi_k |^{\otimes m+1} \mathcal{A}_{\phi_k}^{R^c}(1^\kappa, |\phi_k\rangle^{\otimes m}) \right) \right] > \eta(\kappa) \quad (14)$$

for infinitely many κ , where $\eta := \eta(\kappa)$ is a non-negligible function and $R_{\phi_k}^c$ is the controlled reflection oracle about ϕ_k . Let $q : \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial bound on the query complexity of \mathcal{A} , i.e. \mathcal{A} makes at most $q := q(\kappa)$ queries to its oracle on input 1^κ . Then, due to Theorem 2.8, we may simulate \mathcal{A} by a QPT algorithm \mathcal{B} that does not have access to $R_{\phi_k}^c$ but requires additional copies of $|\phi_k\rangle$.¹⁵ More concretely, \mathcal{B} satisfies the following

$$\text{TD} \left(\mathcal{A}^{R_{\phi_k}^c} (1^\kappa, |\phi_k\rangle^{\otimes m}) \otimes |\phi_k\rangle^{\otimes 8q^2/\eta^2}, \mathcal{B} (1^\kappa, |\phi_k\rangle^{\otimes m} \otimes |\phi_k\rangle^{\otimes 8q^2/\eta^2}) \right) \leq \eta/2. \quad (15)$$

By combining Equations (14) and (15), \mathcal{B} constitutes a successful attack against the (standard) unclonability of Φ : On input $m + 8q^2/\eta^2$ copies of $|\phi_k\rangle$, it produces $m + 1 + 8q^2/\eta^2$ with advantage at least $\eta/2$. \square

Theorem 6.7 ([JLS18, Theorem 2]). *PRSGs imply UCSGs.*

Remark 6.8. In [JLS18, Theorem 5], it was shown that PRSGs imply strong UCSGs. However, they did not show that UCSGs imply strong UCSGs, which we observe in Theorem 6.6.

Corollary 6.9. *For any fixed $O \in \mathcal{O}$, strong UCSGs, UCSGs and PRSGs do not exist relative to O , QPSPACE. Hence, we conclude that strong UCSGs, UCSGs and PRSGs are separated from QEFID pairs (see Definition 2.1).*

The proof is immediate by combining Theorems 6.5 to 6.7 with Theorem 5.1.

Remark 6.10. An incomparable but related result is [NZ24] where the authors separated UCSGs (or unclonable states as referred to in [NZ24]) and non-telegraphable states. Since non-telegraphy seems to imply one-wayness, which implies EFI pairs [KT24], intuitively, it seems that [NZ24, KT24] together imply a separation between UCSGs and EFI pairs. However, this is not true because: 1) the definition of non-telegraphable and unclonable states considered in [NZ24], only allows the adversary to get a single copy of the state in the respective security games, due to which the resulting OWSGs from non-telegraphable states are also single-copy secure, but all existing constructions [KT24, BJ24] of EFI pairs from OWSGs require multi-copy security, 2) we consider a multi-copy definition of UCSGs where the cloner in the cloning game can get any polynomial number of copies of the state, hence ruling out single-copy secure UCSGs as shown in [NZ24] does not rule out all UCSGs as per our definition.

7 Separating Private-key Quantum Money from QEFID

In this last section, we prove our second main separation result, namely that between private-key quantum money schemes and QEFID pairs. Here, in contrast to our oracle separation of OWSGs from QEFID pairs that is laid down in the previous sections, we are able to establish a weaker notion of separation, specifically a *fully black-box separation*.

Fully black-box separation means the impossibility of fully black-box constructions. These include any realization of a cryptographic primitive based on another (referred to as the base primitive), where the construction uses the algorithms underlying the base primitive in a black-box way (i.e. independently of their implementation) and, additionally, the security reduction breaks any implementation of the base primitive given a black-box access to an adversary that breaks the construction based on that implementation. For a formal definition, we refer the reader to Appendix A.

¹⁵In fact, by Theorem 2.8, \mathcal{B} requires additional copies of $|\phi'_k\rangle = |1\rangle \otimes |\phi_k\rangle$, since $R_{\phi_k}^c = R_{\phi'_k}$. However, \mathcal{B} may easily transform $|\phi_k\rangle$ into $|\phi'_k\rangle$.

To rule out a fully black-box construction of private-key quantum money schemes from QEFID pairs it is sufficient, then, to devise an oracle world where (i) there exists a QEFID construction that is secure against polynomial-query adversaries, even if unbounded in runtime, and, on the other hand, (ii) any private-key quantum money scheme may be broken by a polynomial-query attack that is possibly unbounded otherwise. This is a relaxation compared to an oracle separation where we consider polynomially bounded adversaries, both in queries and runtime (where the “unboundedness” may be “pushed” entirely to the oracles). For an explanation on why we are able to achieve only a fully black-box separation in the private quantum money case, and the challenges in obtaining an oracle separation, we refer the reader to the technical overview in Section 1.2.

The oracle world through which we separate between private-key quantum money schemes and QEFID pairs consists of an oracle sampled from the distribution \mathcal{O} that we defined in Definition 3.2. Recall, the oracle \mathcal{O} has been useful to establish our previous separation. Specifically, it provided us with the source of hardness needed to build a QEFID pair that is secure against any \mathcal{O} -aided adversary that makes a bounded number of queries to its oracle, but is otherwise unbounded; see Lemma 4.4. Hence, we already have Item (i) in hand (namely the existence of statistically-secure QEFID pairs) and it remains to show how to break any private-key quantum money scheme relative to \mathcal{O} .

To that end, we will use *shadow tomography* [Aar19], which is a powerful tool that, roughly speaking, enables estimating the outcome of exponentially many POVMs on a given quantum state given only polynomially many copies of that states. In particular, shadow tomography can be seen as a strengthening of the gentle search procedure which we used to break any OWSG in Theorem 5.1.

We will use the following result from [Aar19].

Theorem 7.1 ([Aar19, Theorem 2, Problem 1]). *Let $\epsilon, \delta \in \mathbb{R}$ and $D, M \in \mathbb{N}$. There exists $k \in \tilde{O}(\log D \cdot \log^4 M \cdot \log(1/\delta)/\epsilon^4)$, where \tilde{O} notation hides a $\text{poly}(\log \log(M), \log \log(D), \log(1/\epsilon))$ factor, and a quantum algorithm that takes as input k copies of a (possibly mixed) D -dimensional quantum state ρ , as well as a list of binary-values POVMs Π_1, \dots, Π_M , and outputs real numbers $b_1, \dots, b_M \in [0, 1]$ such that, with probability at least $1 - \delta$, for every $i \in [M]$,*

$$|b_i - \text{Tr}(\Pi_i \rho)| \leq \epsilon.$$

Equipped with Theorem 7.1, we are able to demonstrate a statistical attack that breaks any private-key quantum money scheme relative to any oracle $O \in \mathcal{O}$.

Theorem 7.2. *For any $O \in \mathcal{O}$ and any private-key quantum money scheme \mathbf{M} relative to O that satisfies correctness (as defined in Definition 2.6), there exists a polynomial-query oracle-aided adversary A that breaks the unforgeability of \mathbf{M} .*

Interestingly, we stress that our statistical attack does not make any queries to the verification oracle. Theorem 7.2 in combination with Theorem 4.1, gives us the following corollary: in any fully black-box construction of private quantum money schemes from QEFID pairs, the reduction either has complexity greater than $2^{\lambda/51}$ or success probability smaller than $2^{-\lambda/51}$ (see Appendix A for formal definitions).

Corollary 7.3 (Formal Restatement of Theorem 1.5). *Private-key quantum money schemes are $2^{\lambda/51}$ -separated from QEFID pairs.*

Proof. Suppose there exists a (q, ϵ) -fully black-box construction of a private-key quantum money scheme $\mathbf{M} = (\text{KeyGen}, \text{Mint}, \text{Verify})$ from a QEFID pair with a corresponding reduction \mathcal{R} .

Let \mathcal{O} be the distribution over oracles defined in Definition 3.2 and consider a random $O \leftarrow \mathcal{O}$. Let (D_0^O, D_1^O) be the oracle-aided QEFID pair from Construction 4.2. By construction, (D_0^O, D_1^O) is efficiently sampleable and, by Proposition 4.3, it satisfies statistical fairness as required by Definition 2.1.

By Theorem 7.2, there exists a polynomial-query quantum adversary \mathcal{A} that breaks \mathbf{M} relative to any $O \in \mathcal{O}$ and, in particular, relative to a random $O \leftarrow \mathcal{O}$ with probability 1. Therefore, by the correctness of \mathcal{R} , it must hold that

$$\left| \Pr_{x \leftarrow D_0^O(1^\lambda)} [\mathcal{R}^{\mathcal{A}, O}(1^\lambda, x) = 1] - \Pr_{x \leftarrow D_1^O(1^\lambda)} [\mathcal{R}^{\mathcal{A}, O}(1^\lambda, x) = 1] \right| \geq \epsilon(\lambda)$$

for infinitely many $\lambda \in \mathbb{N}$.

Let $p := p(\kappa)$ be the polynomial bound on the query complexity of \mathcal{A} and $p' := p'(\kappa)$ be the polynomial bound on the query complexity of \mathcal{V} erify. We can replace \mathcal{R} by an algorithm $\tilde{\mathcal{R}}$ that makes $p \cdot p' \cdot q$ queries to O , while not accessing \mathcal{A} at all, by simulating the queries to \mathcal{A} and \mathcal{A} 's queries to \mathcal{V} erify in by itself (in fact, as we already mention, our attack \mathcal{A} queries O only and does not query the verification oracle, hence, in our case there are no “nested” queries to \mathcal{V} erify that must be simulated).

By Lemma 4.4, if $\epsilon(\lambda) \in \omega(2^{-\lambda/51}) \subset \omega(2^{-\lambda/50})$, such an algorithm $\tilde{\mathcal{R}}$ must perform $\omega(2^{\lambda/50})$ queries to O . Hence $p \cdot p' \cdot q \in \omega(2^{\lambda/50})$ and, since p and p' are polynomials, $q \in \omega(2^{\lambda/51})$. \square

The remaining of this section is dedicated to the proof of Theorem 7.2.

7.1 Proof of Theorem 7.2

Fix $O \in \mathcal{O}$. Let $\mathbf{M} = (\text{KeyGen}, \text{Mint}, \text{Verify})$ be a private-key quantum money scheme relative to O . Let κ is the security parameter for \mathbf{M} . Let $n := n(\kappa)$ denote the size (in qubits) of a money state output by Mint on input key of length κ . Let $\mu := \mu(\kappa) \in [0, 1]$ be an inverse polynomial function representing the correctness of \mathbf{M} . Since Mint and Verify are QPT algorithms, they only make polynomially many queries to O . Let $q := q(\kappa)$ be a polynomial bound on the query complexity of both Mint and Verify . Let $\bar{\lambda} := \bar{\lambda}(\kappa)$ be the polynomial bound on the length of the queries made by Mint and Verify . That is, for any $\kappa \in \mathbb{N}$ and any $k \in \{0, 1\}^\kappa$, $\text{Mint}(1^\kappa)$ and $\text{Verify}(k, \cdot)$ make at most $q(\kappa)$ queries to O , where each query is of length at most $\bar{\lambda}(\kappa)$. Let

$$\tau := \bigotimes_{\lambda=1}^{\bar{\lambda}} |S_\lambda\rangle \langle S_\lambda|^{2q^2\kappa^2/\mu^2}. \quad (16)$$

We invoke reflection emulation again to “de-oracle” the algorithms Mint and Verify : By Proposition 3.3, since O consists of reflection oracles and all queries made to O are of length at most $\bar{\lambda}$, there exist QPT algorithms $\widetilde{\text{Mint}}$ and $\widetilde{\text{Verify}}$ such that for every $\kappa \in \mathbb{N}$, key $k \in \{0, 1\}^\kappa$ and state ρ ,¹⁶

$$\text{TD}[\text{Mint}^O(k) \otimes \tau, \widetilde{\text{Mint}}(k, \tau)] \leq \frac{\mu}{\kappa}, \quad (17)$$

$$\text{and } \left| \Pr[\text{Verify}^O(k, \rho) = 1] - \Pr[\widetilde{\text{Verify}}(k, \rho, \tau) = 1] \right| \leq \frac{\mu}{\kappa}. \quad (18)$$

Let \mathcal{VM} denote the algorithm that takes as input two keys $k, k' \in \{0, 1\}^\kappa$, for some $\kappa \in \mathbb{N}$, and a $2\bar{\lambda}q^2\kappa^2/\mu^2$ -register state τ (which will be of the form of Equation (16)), and behaves as follows:

¹⁶Note we define $\widetilde{\text{Mint}}$ to simulate Mint while outputting the state τ that it receives as input. By Proposition 3.3, the state τ does not get disturbed by too much. For the simulation of Verify , however, we do not care about preserving the state τ , and therefore, we omit it from the output.

$\mathcal{VM}(k, k', \tau)$: 1. Compute $(\rho, \tau') \leftarrow \widetilde{\text{Mint}}(k', \tau)$. 2. Output $\widetilde{\text{Verify}}(k, \rho, \tau')$.¹⁷

Our attack against **M** consists of two invocation of shadow tomography, as formalized in Theorem 7.1. Let $m := m(\kappa) = \kappa^{10} \cdot n(\kappa) / \mu(\kappa)^4$, which is polynomial since μ is inverse polynomial in κ .

First, we consider the collection of POVM elements $\{\Pi_k\}$ where, for any $k \in \{0, 1\}^\kappa$, Π_k corresponds to that $\widetilde{\text{Verify}}(k, \cdot)$ accepts, i.e., outputs 1. By Theorem 7.1, there exists an algorithm \mathcal{B}_V that, for any $\kappa \in \mathbb{N}$ and any $n(\kappa)$ -qubit state ρ , takes as input $(\rho \otimes \tau)^{\otimes m(\kappa)}$ and outputs estimates $\{b_V(k, \rho)\}_{k \in \{0, 1\}^\kappa}$ such that with probability $1 - \frac{1}{2^\kappa}$, for every $k \in \{0, 1\}^\kappa$,

$$|b_V(k, \rho) - \Pr[\widetilde{\text{Verify}}(k, \rho, \tau)] = 1| \leq \frac{\mu(\kappa)}{\kappa}. \quad (19)$$

Second, we consider the collection of POVM elements $\{\Lambda_{k, k'}\}$ where, for any $(k, k') \in \{\{0, 1\}^\kappa \times \{0, 1\}^\kappa\}_\kappa$, $\Lambda_{k, k'}$ corresponds to that $\mathcal{VM}(k, k', \cdot)$ outputs 1. By Theorem 7.1 again, there exists an algorithm \mathcal{B}_{VM} that, on input $\tau^{\otimes m}$ outputs estimates $\{b_{VM}(k, k')\}_{k, k' \in \{0, 1\}^\kappa}$, such that with probability $1 - \frac{1}{2^\kappa}$, for every $k, k' \in \{0, 1\}^\kappa$,

$$|b_{VM}(k, k') - \Pr[\mathcal{VM}(k, k', \tau)] = 1| \leq \frac{\mu(\kappa)}{\kappa}. \quad (20)$$

Using \mathcal{B}_V and \mathcal{B}_{VM} , we propose the following query-efficient forger \mathcal{A} against **M**. For any $\kappa \in \mathbb{N}$, on input $m(\kappa)$ copies of an $n(\lambda)$ -qubit money state \mathbb{S}^* , \mathcal{A} behaves as follows:

1. For $\lambda = 1, \dots, \bar{\lambda}$, \mathcal{A} runs $E^O(1^\lambda)$ (from Corollary 5.3) $6q^2\kappa^2m(\kappa)/\mu(\kappa)^2$ times and obtains $3m$ copies of $|S_\lambda\rangle\langle S_\lambda|^{-\otimes 2q^2\kappa^2/\mu^2}$. Overall, this results in the state $\tau^{\otimes 3m}$, where τ is as defined in Equation (16).
2. \mathcal{A} runs \mathcal{B}_V on $(\mathbb{S}^* \otimes \tau)^{\otimes m}$ and gets back estimates $\{b_V(k, \mathbb{S}^*)\}_{k \in \{0, 1\}^\kappa}$.
3. \mathcal{A} runs \mathcal{B}_{VM} on $\tau^{\otimes 2m}$ and gets back estimates $\{b_{VM}(k, k')\}_{k, k' \in \{0, 1\}^\kappa}$.
4. Find the lexicographically first k' such that

$$|b_V(k, \mathbb{S}^*) - b_{VM}(k, k')| \leq 5\mu(\kappa)/\kappa \quad (21)$$

for all $k \in \{0, 1\}^\kappa$. If such a k' does not exist, abort.

5. Let $m' = \lceil 2\kappa m / (\mu(1 - 10/\kappa)) \rceil$ ¹⁸. Run $\text{Mint}^O(k')$ for m' times and output all resulting states.

Theorem 7.2 follows by the following lemma, which we fully prove in Section 7.2.

Lemma 7.4. *The algorithm \mathcal{A} described above breaks the unforgeability of the arbitrary private-key quantum money scheme **M**.*

¹⁷Note that the states ρ and τ' can be entangled.

¹⁸Here $\lceil \cdot \rceil$ denotes the ceiling function on real numbers.

7.2 Proof of Lemma 7.4

Clearly, $m'(\kappa) \in \text{poly}(\kappa)$ since $\mu(\kappa)$ and $m(\kappa)$ are inverse polynomial function and polynomial function of κ respectively. Therefore, \mathcal{A} is polynomial-query.

The rest of the proof is dedicated to analyze the success probability of \mathcal{A} in the forging game against \mathbf{M} . For any $\kappa \in \mathbb{N}$, $k \in \{0, 1\}^\kappa$ and $n(\kappa)$ -qubit state ρ , define

$$\begin{aligned} a_{\mathcal{V}}(k, \rho) &:= \Pr[\widetilde{\text{Verify}}(k, \rho, \tau) = 1], & a_{\mathcal{VM}}(k, k') &:= \Pr[\mathcal{VM}(k, k', \tau) = 1], \\ \$(k) &:= \text{Mint}^O(k), & \text{and } (\widetilde{\$}(k), \tau') &:= \widetilde{\text{Mint}}(k, \tau), \end{aligned}$$

where $\widetilde{\$}(k)$ and τ' can be potentially entangled.

Proposition 7.5. *For any $\kappa \in \mathbb{N}$ and $k, k' \in \{0, 1\}^\kappa$, it holds that*

$$|a_{\mathcal{VM}}(k, k') - a_{\mathcal{V}}(k, \$(k'))| \leq \mu/\kappa.$$

Proof. Let $(\widetilde{\$}(k'), \tau') \leftarrow \widetilde{\text{Mint}}(k', \tau)$. By Equation (17), $\text{TD}[(\widetilde{\$}(k'), \tau'), \$(k') \otimes \tau] \leq \frac{\mu}{\kappa}$. Hence, $|a_{\mathcal{VM}}(k, k') - a_{\mathcal{V}}(k, \$(k'))|$ is the same as

$$|\Pr[\widetilde{\text{Verify}}(k, \widetilde{\$}(k'), \tau') = 1] - \Pr[\widetilde{\text{Verify}}(k, \$(k'), \tau) = 1]| \leq \frac{\mu}{\kappa}.$$

□

In the following claim, we establish a connection between the acceptance probabilities (in fact, the differences therein) by the verification algorithm Verify and its emulated version, namely $\widetilde{\text{Verify}}$.

Proposition 7.6. *For any $\kappa \in \mathbb{N}$, $k, k' \in \{0, 1\}^\kappa$ and $n(\kappa)$ -qubit state ρ ,*

$$\left| |a_{\mathcal{V}}(k, \rho) - a_{\mathcal{VM}}(k, k')| - \left| \Pr[\text{Verify}^O(k, \rho) = 1] - \Pr[\text{Verify}^O(k, \$(k')) = 1] \right| \right| \leq 3\mu/\kappa.$$

Proof. By Equation (18),

$$\left| \Pr[\text{Verify}^O(k, \rho) = 1] - a_{\mathcal{V}}(k, \rho) \right| = \left| \Pr[\text{Verify}^O(k, \rho) = 1] - \Pr[\widetilde{\text{Verify}}(k, \rho, \tau) = 1] \right| \leq \frac{\mu}{\kappa}. \quad (22)$$

Further, using Equation (22) for $\rho = \$(k')$, along with Proposition 7.5 and triangle inequality, we have

$$\begin{aligned} & \left| \Pr[\text{Verify}^O(k, \$(k')) = 1] - a_{\mathcal{VM}}(k, k') \right| \\ & \leq \left| \Pr[\text{Verify}^O(k, \$(k')) = 1] - a_{\mathcal{V}}(k, \$(k')) \right| + \mu/\kappa && \text{By Proposition 7.5.} \\ & \leq \mu/\kappa + \mu/\kappa && \text{By Equation (22).} \\ & = 2\mu/\kappa. \end{aligned} \quad (23)$$

Combining Equations (22) and (23) and by triangle inequality, we get

$$\begin{aligned} |a_{\mathcal{V}}(k, \rho) - a_{\mathcal{VM}}(k, k')| &\leq |\Pr[\text{Verify}^O(k, \rho) = 1] - \Pr[\text{Verify}^O(k, \$(k')) = 1]| + 3\mu/\kappa \\ &|\Pr[\text{Verify}^O(k, \rho) = 1] - \Pr[\text{Verify}^O(k, \$(k')) = 1]| \leq |a_{\mathcal{V}}(k, \rho) - a_{\mathcal{VM}}(k, k')| + 3\mu/\kappa. \end{aligned}$$

The proposition follows by combining the last two equations.

□

We fix $\kappa \in \mathbb{N}$ and let $k^* \leftarrow \text{KeyGen}(1^\lambda)$ be the key sampled by the challenger in the unforgeability game and $\$^* \leftarrow \text{Mint}(k^*)$ be the corresponding input to \mathcal{A} . Let E denote the event that the estimators \mathcal{B}_ν and $\mathcal{B}_{\nu\mathcal{M}}$ run by \mathcal{A} were successful. That is, E occurs if and only if Equations (19) and (20) hold for the values $\{b_\nu(k, \$^*)\}$ and $\{b_{\nu\mathcal{M}}(k, k')\}$ that are output in Steps 2 and 3 of \mathcal{A} , respectively. By the guarantee of Theorem 7.1,

$$\Pr[E] \geq 1 - \frac{2}{2^\kappa}. \quad (24)$$

In the next claim, we rely on the guarantee from shadow tomography to show that, if a state minted under a key k' behaves like the input state $\* w.r.t. verification under any arbitrary key k , then the key k' exhibits an estimate $b_\nu(k, \$^*)$ that is similar to the estimate $b_{\nu\mathcal{M}}(k, k')$.

Proposition 7.7. *Conditioned on E , for any $\kappa \in \mathbb{N}$ and $k, k' \in \{0, 1\}^\kappa$*

$$\left| |b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| - |\Pr[\text{Verify}^O(k, \$^*)] - \Pr[\text{Verify}^O(k, \$ (k'))]| \right| \leq 5\mu/\kappa. \quad (25)$$

Proof. It is enough to show that conditioned on E ,

$$\left| |b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| - |a_\nu(k, \$^*) - a_{\nu\mathcal{M}}(k, k')| \right| \leq 2\mu/\kappa, \quad (26)$$

because then by Proposition 7.6 for $\rho = \* , and triangle inequality,

$$\left| |b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| - \left| \Pr[\text{Verify}^O(k, \$^*)] - \Pr[\text{Verify}^O(k, \$ (k'))] \right| \right| \quad (27)$$

$$\leq \left| |b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| - |a_\nu(k, \$^*) - a_{\nu\mathcal{M}}(k, k')| \right| + 3\mu/\kappa \quad (28)$$

$$\leq 2\mu/\kappa + 3\mu/\kappa = 5\mu/\kappa, \quad (29)$$

which concludes the proof of the proposition. Therefore, we prove Equation (26).

Note that conditioned on E , by its definition, it holds that

$$|b_\nu(k, \$^*) - a_\nu(k, \$^*)| \leq \frac{\mu}{\kappa} \quad \text{and} \quad |b_{\nu\mathcal{M}}(k, k') - a_{\nu\mathcal{M}}(k, k')| \leq \frac{\mu}{\kappa}, \quad (30)$$

and, consequently, by triangle inequality on Equation (30), we get

$$|b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| \leq |a_\nu(k, \$^*) - a_{\nu\mathcal{M}}(k, k')| + \frac{2\mu}{\kappa}. \quad (31)$$

$$|a_\nu(k, \$^*) - a_{\nu\mathcal{M}}(k, k')| \leq |b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k')| + \frac{2\mu}{\kappa}. \quad (32)$$

The last two equations together imply Equation (26). □

In particular, applying Proposition 7.7 with $k' = k^*$, it follows that, conditioned on E

$$|b_\nu(k, \$^*) - b_{\nu\mathcal{M}}(k, k^*)| \leq 5\mu/\kappa. \quad (33)$$

for all $k \in \{0, 1\}^\kappa$. Hence, conditioned on E , there exists k' such that Equation (21) is satisfied, in which case \mathcal{A} does not abort. From this point on, we assume E occurs. By Equation (24), this incurs only a negligible loss in the success probability.

Let k'' be the k' that \mathcal{A} finds in step 4. In the following claim, we show that a quantum money state minted under k'' , namely $\$(k'')$, verifies under k^* with good probability.

Claim 7.8. Conditioned on E ,

$$\Pr[\mathcal{V}erify^O(k^*, \$(k'')) = 1] \geq \mu(1 - 10/\kappa). \quad (34)$$

Proof. By our construction of \mathcal{A} (step 4), it holds that

$$|b_{\mathcal{V}}(k^*, \$^*) - b_{\mathcal{V}\mathcal{M}}(k^*, k'')| \leq 5\mu/\kappa.$$

By the above and Proposition 7.7, we conclude that

$$\begin{aligned} & |\Pr[\mathcal{V}erify^O(k^*, \$^*) = 1] - \Pr[\mathcal{V}erify^O(k^*, \$(k'')) = 1]| \\ & \leq |b_{\mathcal{V}}(k^*, \$^*) - b_{\mathcal{V}\mathcal{M}}(k^*, k'')| + 5\mu/\kappa \leq 10\mu/\kappa, \end{aligned}$$

which implies the inequality in the claim by the μ -correctness of the scheme. \square

Having shown that a single money state $\$(k'') \leftarrow \mathcal{M}int^O(k'')$ passes verification under k^* , all that separates us from showing that \mathcal{A} is successful is a simple concentration bound, using which we argue that $m' = \lceil 2\kappa m / (\mu(1 - 10/\kappa)) \rceil$ such states contain at least $m + 1$ states that pass verification, with good enough probability.

Let r denote the number of money states among the m' states that \mathcal{A} outputs that pass verification under k^* . By Equation (34), $\mathbb{E}[r] \geq (\mu(1 - 10/\kappa)) \cdot m' \geq 2\kappa m$. Thus, by Hoeffding's inequality,

$$\begin{aligned} \Pr[r \leq m] & \leq \Pr[|r - \mathbb{E}[r]| \geq \kappa m] \leq 2e^{-2\kappa^2 m^2 / m'} \leq 2e^{-2\kappa m \cdot (\mu(1-10/\kappa) \cdot \frac{m}{2m + \mu(1-10/\kappa)})} \\ & \leq 2e^{-2\kappa m \cdot (\mu(1-10/\kappa)/4)} = 2e^{-\kappa \frac{m\kappa}{2\mu^4} \cdot (\mu(1-10/\kappa))} = 2e^{-\frac{\kappa^{11}}{2\mu^3} (1-10/\kappa)}, \end{aligned}$$

which is negligible. This, together with Equation (24), completes the proof of Lemma 7.4.

Acknowledgments. We want to thank Or Sattath for the helpful discussion on the specification of the QPSPACE machine oracle.

Amit Behera was funded by the Israel Science Foundation (grant No. 2527/24) and the European Union (ERC-2022-COG, ACQUA, 101087742). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.



Giulio Malavolta is supported by the European Research Council through an ERC Starting Grant (Grant agreement No. 101077455, ObfusQation).

Tomoyuki Morimae is supported by JST CREST JPMJCR23I3, JST Moonshot R&D JPMJMS2061-5-1-1, JST FOREST, MEXT QLEAP, the Grant-in Aid for Transformative Research Areas (A) 21H05183, and the Grant-in-Aid for Scientific Research (A) No.22H00522.

Tamer Mour is supported by European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (Grant agreement No. 101019547)

References

- [Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016. (Cited on page 9.)

- [Aar19] Scott Aaronson. Shadow tomography of quantum states. *SIAM J. Comput.*, 49(5):STOC18–368, 2019. (Cited on page 7, 8, 9, 24, 25, 30.)
- [AGL24] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common haar state model: Feasibility results and separations. Cryptology ePrint Archive, Paper 2024/1043, 2024. (Cited on page 10, 11.)
- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Cham, August 2022. (Cited on page 3.)
- [BBD⁺97] Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, October 1997. (Cited on page 41.)
- [BCN24] John Bostanci, Boyang Chen, and Barak Nehoran. Oracle separation between quantum commitments and quantum one-wayness, 2024. (Cited on page 11, 12.)
- [BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *ITCS 2023*, volume 251, pages 24:1–24:21. LIPIcs, January 2023. (Cited on page 3.)
- [BEM⁺23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the uhlmann transformation problem. *arXiv preprint arXiv:2306.13073*, 2023. (Cited on page 15.)
- [BG11] Zvika Brakerski and Oded Goldreich. From absolute distinguishability to positive distinguishability. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, volume 6650 of *Lecture Notes in Computer Science*, pages 141–155. Springer, 2011. (Cited on page 23, 40.)
- [BJ24] Rishabh Batra and Rahul Jain. Commitments are equivalent to one-way state generators, 2024. (Cited on page 4, 5, 13, 29.)
- [CCS24] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: constructing and separating quantum pseudorandomness. *arXiv preprint arXiv:2404.03295*, 2024. (Cited on page 8, 10, 11, 12, 18, 19, 24, 25.)
- [CGG⁺23] Bruno Cavalari, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. On the computational hardness of quantum one-wayness. *arXiv preprint arXiv:2312.08363*, 2023. (Cited on page 7, 8.)
- [CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 215–248. Springer, Cham, August 2024. (Cited on page 4, 5, 6, 11, 12, 14, 27.)

- [CM24] Andrea Coladangelo and Saachi Mutreja. On black-box separations of quantum digital signatures from pseudorandom states, 2024. (Cited on page 10, 11.)
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986. (Cited on page 3.)
- [GL89] Oded Goldreich and Leonid A Levin. A hard-core predicate for all one-way functions. In *STOC*, pages 25–32. ACM, 1989. (Cited on page 3.)
- [Go190] Oded Goldreich. A note on computational indistinguishability. *Information Processing Letters*, 34(6):277–281, 1990. (Cited on page 3, 6.)
- [GTB23] Tudor Giurgica-Tiron and Adam Bouland. Pseudorandomness from subset states, 2023. (Cited on page 12.)
- [Har13] Aram W. Harrow. The church of the symmetric subspace, 2013. (Cited on page 43.)
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. (Cited on page 3.)
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 2020. (Cited on page 7.)
- [HLM17] Aram W. Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. *Proc. SODA 2017*, pp. 1598-1611, 2017. (Cited on page 24, 25.)
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th FOCS*, pages 230–235. IEEE Computer Society Press, October / November 1989. (Cited on page 3.)
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989. (Cited on page 3.)
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Cham, August 2018. (Cited on page 3, 4, 6, 7, 9, 11, 15, 17, 19, 29, 41.)
- [JMW23] Fernando Granha Jeronimo, Nir Magrafta, and Pei Wu. Pseudorandom and pseudoentangled states from subset states, 2023. (Cited on page 12.)
- [KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1589–1602. ACM Press, June 2023. (Cited on page 3, 40.)
- [Kre21] W. Kretschmer. Quantum pseudorandomness and classical complexity. *TQC 2021*, 2021. (Cited on page 3, 4, 10, 11.)

- [KT24] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *56th ACM STOC*, pages 968–978. ACM Press, June 2024. (Cited on page 4, 5, 7, 12, 14, 29.)
- [LMW24] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *56th ACM STOC*, pages 979–990. ACM Press, June 2024. (Cited on page 3.)
- [LR86] Michael Luby and Charles Rackoff. Pseudo-random permutation generators and cryptographic composition. In *18th ACM STOC*, pages 356–363. ACM Press, May 1986. (Cited on page 3.)
- [MMWY24] Giulio Malavolta, Tomoyuki Morimae, Michael Walter, and Takashi Yamakawa. Exponential quantum one-wayness and efi pairs, 2024. (Cited on page 4.)
- [MW05] Chris Marriott and John Watrous. Quantum arthur—merlin games. *Comput. Complex.*, 14(2):122–152, June 2005. (Cited on page 25.)
- [MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Cham, August 2022. (Cited on page 3, 13, 14.)
- [MY24] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In Frédéric Magniez and Alex Bredariol Grilo, editors, *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024, September 9-13, 2024, Okinawa, Japan*, volume 310 of *LIPICs*, pages 4:1–4:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. (Cited on page 3, 5, 9, 14.)
- [Nao90] Moni Naor. Bit commitment using pseudo-randomness. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 128–136. Springer, New York, August 1990. (Cited on page 3.)
- [NZ24] Barak Nehoran and Mark Zhandry. A computational separation between quantum no-cloning and no-telegraphing. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, January 30 to February 2, 2024, Berkeley, CA, USA*, volume 287 of *LIPICs*, pages 82:1–82:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. (Cited on page 29.)
- [Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394. ACM Press, May 1990. (Cited on page 3.)
- [Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, April 2018. (Cited on page 10.)
- [WB24] Adam Bene Watts and John Bostanci. Quantum event learning and gentle random measurements, 2024. (Cited on page 24.)
- [Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983. (Cited on page 4.)
- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 80–91. IEEE Computer Society, 1982. (Cited on page 23, 40.)

A Defining Fully Black-Box Separation of Private-key Quantum Money Schemes from QEFID pairs

We hereby define a fully black-box construction of private-key quantum money schemes from QEFID pairs. Roughly speaking, such a construction is called fully black-box if it makes only a black-box used of the underlying QEFID pair and, additionally, the security reduction breaks the underlying QEFID pair while making only black-box access to the presumed adversary against the quantum money scheme.

Definition A.1 (Fully Black-box Construction of Private-key Quantum Money Schemes from QEFID Pairs). *Let $q : \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon : \mathbb{N} \rightarrow [0, 1]$. A (q, ϵ) -fully black-box construction of private-key quantum money schemes from QEFID pairs, with key space $\mathcal{K} = \{\mathcal{K}_\kappa\}_{\kappa \in \mathbb{N}}$, is a triple of polynomial-query oracle-aided algorithms $(\text{KeyGen}, \text{Mint}, \text{Verify})$ that satisfy the syntax of a private-key quantum money scheme (see Definition 2.6) and an oracle-aided reduction \mathcal{R} satisfying the following properties:*

- **Construction Correctness:** *For any unitary oracle O , $(\text{KeyGen}, \text{Mint}, \text{Verify})$ satisfy correctness relative to O . That is, there exists an inverse-polynomial function $\mu : \mathbb{N} \rightarrow [0, 1]$ such that, for any $\kappa \in \mathbb{N}$,*

$$\Pr[\text{Verify}^O(k, \$k); k \leftarrow \text{KeyGen}^O(1^\kappa), \$k \leftarrow \text{Mint}^O(k)] \geq \mu(\kappa).$$

- **Black-box Security Reduction:** *For any unitary oracle O , any O -aided QEFID pair (D_0^O, D_1^O) that is efficiently sampleable and statistically far as required by Definition 2.1, and any non-uniform oracle-aided quantum adversary \mathcal{A} , if there exist polynomials $m, m' : \mathbb{N} \rightarrow \mathbb{N}$ where $m' > m$ such that, for infinitely many $\kappa \in \mathbb{N}$,*

$$\Pr[\exists S \subseteq [m'(\kappa)], |S| > m, \text{Verify}^O(k, \$i) = 1 \forall i \in S;$$

$$k \leftarrow \text{KeyGen}^O(1^\kappa), \$k \leftarrow \text{Mint}^O(k), \$_{1, \dots, m'(\kappa)} \leftarrow \mathcal{A}^{\text{Verify}^O(k, \cdot)}(1^\kappa, \$k^{\otimes m(\kappa)})] > \frac{1}{2},$$

(where $\$_{1, \dots, m'(\kappa)}$ is a state on $m'(\kappa)$ registers and $\$_i$ denotes its i^{th} register), then \mathcal{R} breaks the computational indistinguishability of (D_0^O, D_1^O) with advantage ϵ . That is, for infinitely many $\lambda \in \mathbb{N}$,

$$\left| \Pr_{x \leftarrow D_0^O} [\mathcal{R}^{\mathcal{A}, O}(x) = 1] - \Pr_{x \leftarrow D_1^O} [\mathcal{R}^{\mathcal{A}, O}(x) = 1] \right| \geq \epsilon(\lambda).$$

- **Reduction Efficiency:** *For any $\lambda \in \mathbb{N}$ and $y \in \{0, 1\}^\lambda$, $\mathcal{R}^{\mathcal{A}, O}(y)$ makes at most $q(\lambda)$ queries to the oracles O and \mathcal{A} .*

We define a fully black-box α -separation to embody the impossibility of any fully black-box construction abiding a trade-off (parameterized by α) between the complexity of the underlying reduction and its success probability. A larger value of α gives stronger separation and, in particular, superpolynomial α indicates the impossibility of a reduction that is both polynomial time and has non-negligible advantage, as typically required in the traditional cryptographic setting.

Definition A.2 (Black-box Separation of Private-key Quantum Money Schemes From QEFID Pairs). *We say that private-key quantum money schemes are α -separated from QEFID pairs, for $\alpha(\lambda) > 1$, if for any (q, ϵ) -fully black-box construction of private-key quantum money schemes from QEFID pairs, it holds that either*

1. $q(\lambda) > O(\alpha(\lambda))$, or
2. $\epsilon(\lambda) \leq O(1/\alpha(\lambda))$.

B From Absolute-Gap Distinguisher to Positive-Gap Distinguisher

Lemma 4.7 (Absolute-Gap to Positive-Gap Distinguisher). *Let \mathcal{O} be a distribution over oracles. Let $D_0^{\mathcal{O}}$ and $D_1^{\mathcal{O}}$ be two \mathcal{O} -aided distributions over $\{0, 1\}^\lambda$ with corresponding sampling algorithms. Let \mathcal{A} be an \mathcal{O} -aided quantum algorithm such that*

$$\mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\left| \Pr_{x \leftarrow D_0} [\mathcal{A}(x) = 1] - \Pr_{x \leftarrow D_1} [\mathcal{A}(x) = 1] \right| \right] = \delta.$$

Then, there exists an \mathcal{O} -aided quantum algorithm \mathcal{B} such that

$$\mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{x \leftarrow D_0} [\mathcal{B}(x) = 1] - \Pr_{x \leftarrow D_1} [\mathcal{B}(x) = 1] \right] \geq \delta^2.$$

The runtime and query complexity of \mathcal{B} is twice that of \mathcal{A} in addition to that of the sampling algorithms of D_0 and D_1 .

The proof follows by the same reduction as in [KQST23, Corollary 32] based on Yao's distinguishing/predictor lemma [Yao82] (also see [BG11]). We repeat the proof below.

Proof. We define the distinguisher $\mathcal{B}^{\mathcal{O}}$ to behave as follows upon receiving an input x :

1. Sample $c \leftarrow \{0, 1\}$.
2. Sample $x' \leftarrow D_c^{\mathcal{O}}$ and run $\mathcal{A}^{\mathcal{O}}(x')$ to get an outcome d .
3. Run $\mathcal{A}^{\mathcal{O}}(x)$ to get an outcome e .
4. Output $c \oplus d \oplus e$.

For any $O \in \mathcal{O}$, we define

$$a(O) := \Pr_{x \leftarrow D_0^{\mathcal{O}}} [\mathcal{A}^{\mathcal{O}}(x) = 1] \quad \text{and} \quad b(O) := \Pr_{x \leftarrow D_1^{\mathcal{O}}} [\mathcal{A}^{\mathcal{O}}(x) = 1].$$

It is easy to see that, for any fixed $O \in \mathcal{O}$, $\Pr[c \oplus d = 0] = (1 + a(O) - b(O))/2$ and $\Pr[c \oplus d = 1] = (1 + b(O) - a(O))/2$. Hence we get,

$$\begin{aligned} & \Pr_{x \leftarrow D_0^{\mathcal{O}}} [\mathcal{B}^{\mathcal{O}}(x) = 1] - \Pr_{x \leftarrow D_1^{\mathcal{O}}} [\mathcal{B}^{\mathcal{O}}(x) = 1] \\ &= \Pr[c \oplus d = 0] \cdot (a(O) - b(O)) + \Pr[c \oplus d = 1] \cdot (b(O) - a(O)) \\ &= \frac{1 + a(O) - b(O)}{2} \cdot (a(O) - b(O)) + \frac{1 + b(O) - a(O)}{2} \cdot (b(O) - a(O)) \\ &= (a(O) - b(O))^2 \\ &= |a(O) - b(O)|^2 \end{aligned}$$

For a last step, we apply the Cauchy-Schwarz inequality as follows

$$\begin{aligned} \mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} \left[\Pr_{x \leftarrow D_0} [\mathcal{B}(x) = 1] - \Pr_{x \leftarrow D_1} [\mathcal{B}(x) = 1] \right] &= \mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} [|a(O) - b(O)|^2] \\ &\geq \mathbb{E}_{\mathcal{O} \leftarrow \mathcal{O}} [|a(O) - b(O)|^2] \\ &= \delta^2. \end{aligned}$$

□

C Proof of Generalized Reflection Emulation

Theorem 2.8 (Reflection Emulation [JLS18, Theorem 4]). *Let Q be a quantum oracle. Let $|\psi\rangle$ be a quantum state and let $R_\psi = I - 2|\psi\rangle\langle\psi|$ be the corresponding reflection unitary. Let $|\phi\rangle$ be a state not necessarily independent of $|\psi\rangle$. Let \mathcal{A} be a (Q, R_ψ) -aided quantum circuit that makes q queries to the oracle R_ψ . Then, there exists a Q -aided quantum circuit \mathcal{B} such that, for any $\ell \in \mathbb{N}$,*

$$\text{TD}\left(\mathcal{A}^{Q, R_\psi}(|\phi\rangle) \otimes |\psi\rangle^{\otimes \ell}, \mathcal{B}^Q(|\phi\rangle \otimes |\psi\rangle^{\otimes \ell})\right) \leq \frac{2q}{\sqrt{\ell+1}}.$$

Further, if \mathcal{A} is of polynomial size then so is \mathcal{B} .

Proof. Assume without loss of generality that \mathcal{A}^{Q, R_ψ} is just a unitary followed by the discarding of some registers, and that the input state is pure.¹⁹ We define \mathcal{B} as follows. Let T denote the input register where the copies of $|\psi\rangle$ reside and let $|\theta\rangle_T := |\psi\rangle^{\otimes \ell}$. \mathcal{B} behaves the same as \mathcal{A} does, except that every query to the oracle R_ψ on some register D is replaced by applying the reflection about the symmetric subspace, i.e., $R_{\sqrt{\ell+1}\mathbb{C}^N}$, on the registers D and T . Here, \mathbb{C}^N represents the Hilbert space in which the state $|\psi\rangle$ resides and $\sqrt{\ell+1}\mathbb{C}^N$ is the symmetric subspace over the $\ell+1$ registers (i.e., the ℓ registers of T and the input register D) with respect to \mathbb{C}^N .

Clearly, the runtime of \mathcal{B} is polynomial in ℓ, q , and the runtime of \mathcal{A} , because implementing the reflection about the symmetric subspace can be done in time polynomial in the number of the number of registers, and the number of qubits for each register [JLS18, Page 17, 18] and [BBD⁺97].

We will show that the intermediate state of \mathcal{A} and \mathcal{B} , right after the first R_ψ -query made by \mathcal{A} or (resp.) the simulation thereof is made by \mathcal{B} , are statistically close. A bound on the distance between the final outcome of the two algorithms is then implied by inducting the same argument over all oracle queries.

Let $|\phi\rangle$ be the intermediate state just before the first oracle query to R_ψ or the simulated query under the algorithms \mathcal{A} or, respectively, \mathcal{B} (up to this point the algorithms are identical and so is their intermediate state). Let D be the query register on which \mathcal{A} queries R_ψ and W denote the rest of the registers. We can write $|\phi\rangle$ as $\sum_s c_s |s\rangle_W \otimes |\phi^s\rangle_D$ for some pure states $|\phi^s\rangle$ and amplitudes c_s .

Define

$$\begin{aligned} |\Psi_{\mathcal{A}}^s\rangle &:= R_\psi(|\phi^s\rangle_D) \otimes |\theta\rangle_T, \text{ and} \\ |\Psi_{\mathcal{B}}^s\rangle &:= R_{\sqrt{\ell+1}\mathbb{C}^N}(|\phi^s\rangle_D \otimes |\theta\rangle_T). \end{aligned}$$

Clearly, the state of all the registers, including the T register under the algorithms \mathcal{A} and \mathcal{B} are, respectively,

$$|\tilde{\Psi}_{\mathcal{A}}\rangle = \left(\sum_s c_s |s\rangle_W \otimes |\Psi_{\mathcal{A}}^s\rangle_{D,T} \right) \quad \text{and} \quad |\tilde{\Psi}_{\mathcal{B}}\rangle = \left(\sum_s c_s |s\rangle_W \otimes |\Psi_{\mathcal{B}}^s\rangle_{D,T} \right).$$

Note that,

$$|\langle \tilde{\Psi}_{\mathcal{A}} | \tilde{\Psi}_{\mathcal{B}} \rangle| = \left| \sum_s |c_s|^2 \langle \Psi_{\mathcal{A}}^s | \Psi_{\mathcal{B}}^s \rangle \right|. \quad (35)$$

In the proof of [JLS18, Theorem 4], the following was shown:

¹⁹If the input state is a mixed state, then we can instead consider a purification of the input state, and instead consider that \mathcal{A} acts on the entire purified state while acting as identity on the purification registers.

Claim C.1. For any pre-query state $|\phi\rangle_D$, the corresponding output states $|\Psi_A\rangle_{D,T}$ and $|\Psi_B\rangle_{D,T}$ defined as

$$|\Psi_A\rangle := R_\psi(|\phi\rangle) \otimes |\theta\rangle_T \quad \text{and} \quad |\Psi_B\rangle := R_{\sqrt{\ell+1}\mathbb{C}^N}(|\phi\rangle \otimes |\theta\rangle_T),$$

satisfy, $\langle \Psi_A | \Psi_B \rangle \in \mathbb{R}$ and

$$\langle \Psi_A | \Psi_B \rangle \geq 1 - \frac{2}{\ell+1},$$

which is non-negative since $\ell \geq 4$.

We defer the proof of the above claim to the sequel. Combining Claim C.1 with Equation (35) we conclude that

$$|\langle \tilde{\Psi}_A | \tilde{\Psi}_B \rangle| = \left| \sum_s |c_s|^2 \langle \Psi_A^s | \Psi_B^s \rangle \right| = \sum_s |c_s|^2 |\langle \Psi_A^s | \Psi_B^s \rangle| \geq \min_s \langle \Psi_A^s | \Psi_B^s \rangle = \langle \Psi_A^{s_{\min}} | \Psi_B^{s_{\min}} \rangle, \quad (36)$$

where $s_{\min} = \arg \min_s |\langle \Psi_A^s | \Psi_B^s \rangle|$.

Hence,

$$\begin{aligned} \text{TD} [|\tilde{\Psi}_A\rangle, |\tilde{\Psi}_B\rangle] &= \sqrt{1 - |\langle \tilde{\Psi}_A | \tilde{\Psi}_B \rangle|^2} \\ &\leq \sqrt{1 - |\langle \Psi_A^{s_{\min}} | \Psi_B^{s_{\min}} \rangle|^2} && \text{By Equation (36).} \\ &\leq \sqrt{1 - \left(1 - \frac{2}{\ell+1}\right)^2} && \text{By Claim C.1} \\ &\leq \sqrt{1 - \left(1 - 2 \cdot \frac{2}{\ell+1}\right)} = \sqrt{\frac{4}{\ell+1}} = \frac{2}{\sqrt{\ell+1}}. \end{aligned}$$

Let $|\tilde{\Psi}_A^q\rangle$ and $|\tilde{\Psi}_B^q\rangle$ denote the final states of all registers (including the T registers) of algorithms \mathcal{A} and \mathcal{B} before discarding any registers. Then, by inducting on the number of oracle queries, we conclude that if $|\tilde{\Psi}_A^q\rangle$ and $|\tilde{\Psi}_B^q\rangle$ denote the state of all registers including the T registers after q queries and just before discarding any registers, under algorithms \mathcal{A} and \mathcal{B} respectively, then,

$$\text{TD} [|\tilde{\Psi}_A^q\rangle, |\tilde{\Psi}_B^q\rangle] \leq \frac{2q}{\sqrt{\ell+1}}.$$

Since discarding registers can only reduce the trace distance between two pure states, we conclude that

$$\text{TD} \left[\mathcal{A}^{R_\psi, Q}(|\phi\rangle) \otimes |\psi\rangle^{\otimes \ell}, \mathcal{B}^Q(|\phi\rangle \otimes |\psi\rangle^{\otimes \ell}) \right] \leq \frac{2q}{\sqrt{\ell+1}}. \quad (37)$$

We now recall the proof of Claim C.1 for completeness. First, note that $R_{\sqrt{\ell+1}\mathbb{C}^N}$ can be written as

$$\begin{aligned} R_{\sqrt{\ell+1}\mathbb{C}^N} &= I - 2 \cdot \text{Proj}_{\sqrt{\ell+1}\mathbb{C}^N} \\ &= I - 2 \cdot \frac{1}{(\ell+1)!} \sum_{\pi \in S_{\ell+1}} W_\pi, \end{aligned} \quad (38)$$

where $S_{\ell+1}$ is the symmetric group over $\ell+1$ elements,

and $W_\pi := \sum_{x_1, \dots, x_{\ell+1} \in \{0, 1, \dots, N-1\}} |x_{\pi^{-1}(1)}, \dots, x_{\pi^{-1}(\ell+1)}\rangle \langle x_1, \dots, x_{\ell+1}|$. For a proof of the last equality, see [Har13, Proposition 6]. Next, for any computational basis state $|x\rangle|y\rangle$, note that

$$\begin{aligned}
& \langle x| \otimes \langle \theta| R_{\sqrt{\ell+1}\mathbb{C}^N} |y\rangle \otimes |\theta\rangle \\
&= \langle x| \otimes \langle \theta| \left(I - 2 \cdot \frac{1}{(\ell+1)!} \sum_{\pi \in \mathcal{S}_{\ell+1}} W_\pi \right) |y\rangle \otimes |\theta\rangle \\
&= \langle x|y\rangle - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1)=1} \langle x| \otimes \langle \theta| W_\pi |y\rangle \otimes |\theta\rangle - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1) \neq 1} \langle x| \otimes \langle \theta| W_\pi |y\rangle \otimes |\theta\rangle \\
&= \langle x|y\rangle - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1)=1} \langle x| \otimes \langle \psi|^{\otimes \ell} W_\pi |y\rangle \otimes |\psi\rangle^{\otimes \ell} - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1) \neq 1} \langle x| \otimes \langle \psi|^{\otimes \ell} W_\pi |y\rangle \otimes |\psi\rangle^{\otimes \ell} \\
&= \langle x|y\rangle - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1)=1} \langle x|y\rangle - \frac{2}{(\ell+1)!} \sum_{\pi: \pi(1) \neq 1} \langle \psi|y\rangle \langle x|\psi\rangle \\
&= \langle x|y\rangle - \frac{2\ell}{(\ell+1)!} \langle x|y\rangle - \frac{2((\ell+1)! - \ell!)}{(\ell+1)!} \langle \psi|y\rangle \langle x|\psi\rangle \\
&= \frac{\ell-1}{\ell+1} \langle x|y\rangle - \frac{2\ell}{\ell+1} \langle \psi|y\rangle \langle x|\psi\rangle. \tag{39}
\end{aligned}$$

Therefore, we conclude that

$$\begin{aligned}
(I \otimes \langle \theta|) R_{\sqrt{\ell+1}\mathbb{C}^N} (I \otimes |\theta\rangle) &= \left(\sum_x |x\rangle \langle x| \otimes \langle \theta| \right) R_{\sqrt{\ell+1}\mathbb{C}^N} \left(\sum_y |y\rangle \langle y| \otimes |\theta\rangle \right) \\
&= \sum_{x,y} |x\rangle (\langle x| \otimes \langle \theta| R_{\sqrt{\ell+1}\mathbb{C}^N} |y\rangle \otimes |\theta\rangle) \langle y| \\
&= \sum_{x,y} \left(\frac{\ell-1}{\ell+1} \langle x|y\rangle - \frac{2\ell}{\ell+1} \langle \psi|y\rangle \langle x|\psi\rangle \right) |x\rangle \langle y| \quad \text{By Equation (39)} \\
&= \frac{\ell-1}{\ell+1} \sum_x |x\rangle \langle x| - \frac{2\ell}{\ell+1} \sum_{x,y} \langle \psi|y\rangle \langle x|\psi\rangle |x\rangle \langle y| \\
&= \frac{\ell-1}{\ell+1} I - \frac{2\ell}{\ell+1} |\psi\rangle \langle \psi|. \tag{40}
\end{aligned}$$

Hence,

$$\begin{aligned}
\langle \Psi_{\mathcal{A}} | \Psi_{\mathcal{B}} \rangle &= \text{Tr}[(|\phi\rangle \otimes |\theta\rangle)(\langle \phi| \otimes \langle \theta|) ((R_\psi \otimes I) R_{\sqrt{\ell+1}\mathbb{C}^N})] \\
&= \text{Tr}[|\phi\rangle \langle \phi| R_\psi (I \otimes \langle \theta|) R_{\sqrt{\ell+1}\mathbb{C}^N} (I \otimes |\theta\rangle)] \quad \text{By Cyclicity of trace,} \\
&= \text{Tr}[|\phi\rangle \langle \phi| (I - 2|\psi\rangle \langle \psi|) \left(\frac{\ell-1}{\ell+1} I - \frac{2\ell}{\ell+1} |\psi\rangle \langle \psi| \right)] \quad \text{By Equation (40)} \\
&= \text{Tr}[|\phi\rangle \langle \phi| \left(\frac{\ell-1}{\ell+1} I - \frac{2\ell}{\ell+1} |\psi\rangle \langle \psi| - \frac{2(\ell-1)}{\ell+1} |\psi\rangle \langle \psi| + \frac{4\ell}{\ell+1} |\psi\rangle \langle \psi| \right)] \\
&= \text{Tr}[|\phi\rangle \langle \phi| \left(\frac{\ell-1}{\ell+1} I + \frac{2}{\ell+1} |\psi\rangle \langle \psi| \right)] \\
&= \frac{\ell-1}{\ell+1} + \frac{2}{\ell+1} |\langle \psi|\phi\rangle|^2 \\
&\geq \frac{\ell-1}{\ell+1} = 1 - \frac{2}{\ell+1}. \tag{41}
\end{aligned}$$

This concludes the proof of Claim C.1.

□