IACR Communications in Cryptology
ISSN 3006-5496, Vol. 1, No. 4, 26 pages.

https://doi.org/10.62056/a6ksr-10k

A Note on the Minimality of One-Way Functions in Post-Quantum Cryptography

Sam Buxbaum[1] and Mohammad Mahmoody[2]

[1] Boston University, USA
[2] University of Virginia, USA

Abstract. In classical cryptography, one-way functions (OWFs) play a central role as the minimal primitive that (almost) all primitives imply. The situation is more complicated in *quantum* cryptography, in which honest parties and adversaries can use quantum computation and communication, and it is known that analogues of OWFs in the quantum setting might not be minimal.

In this work we ask whether OWFs are minimal for the intermediate setting of *post-quantum* cryptography, in which the protocols are classical while they shall resist quantum adversaries. We show that for a wide range of natural settings, if a primitive $\mathcal{Q}$ implies OWFs, then so does its (uniformly or non-uniformly secure) post-quantum analogue. In particular, we show that if a primitive $\mathcal{Q}$ implies any other primitive $\mathcal{P}$ that has a 2-message security game (e.g., OWFs) through a black-box classical security reduction $\mathcal{R}$, then one can always (efficiently) turn any polynomial-size quantum adversary breaking $\mathcal{P}$ into a polynomial-size quantum adversary breaking $\mathcal{Q}$. Note that this result holds even if the implementation of $\mathcal{P}$ using that of $\mathcal{Q}$ is arbitrarily non-black-box.

We also prove extensions of this result for when the reduction $\mathcal{R}$ anticipates its oracle adversary to be deterministic, whenever either of the following conditions hold: (1) the adversary needs to win the security game of $\mathcal{Q}$ only with non-negligible probability (e.g., $\mathcal{Q}$ is collision-resistant hashing) or (2) that either of $\mathcal{P}$ and $\mathcal{Q}$ have "falsifiable" security games (this is the case when $\mathcal{P}$ is OWFs). Our work leaves open answering our main question when $\mathcal{Q}$ implies OWFs through a non-black-box security reduction, or when $\mathcal{P}$ uses a more complicated security game than a two-message one.

Keywords: one-way functions · black-box reductions · post-quantum cryptography

Contents

E-mail: sambux@bu.edu (Sam Buxbaum), mohammad@virginia.edu (Mohammad Mahmoody)

Received: 2024-10-09 Accepted: 2024-12-03

1 Introduction

One-way functions, or functions that are easy to compute but (typically) hard to invert, are one of the most foundational concepts in cryptography. The seminal work of Impagliazzo and Luby [IL89] showed that one-way functions are necessary for many basic cryptographic primitives. Specifically, they showed that the existence of secure protocols for identification, symmetric-key encryption, commitments, and coin flipping imply the existence of one-way functions. Further works have extended these results to show the *equivalence* of several other primitives with one-way functions as well, including pseudorandom generators [ILL89, HILL99] and digital signatures [Rom90, NY89].

We now know of an entire hierarchy of computational assumptions that are necessary for various cryptographic primitives, with one-way functions at the base of the hierarchy. Simply put, classical cryptography cannot exist without one-way functions. Their role has earned them the name of the "minimal assumption" as their existence is necessary for a big bulk of cryptography.

The rise of quantum computing has shuffled this landscape. Quantum computing introduces both new risks, such as the ability to break many popular public key encryption schemes [Sho94], and new opportunities, such as the ability to do cryptography from different, potentially weaker assumptions than are necessary for classical cryptography. One of the first indications of a different quantum landscape was the quantum key distribution protocol of Bennett and Brassard [BB84]. They proposed a protocol for key agreement which only required a quantum channel as well as a classical authenticated channel, whereas key agreement in the classical world requires computationally bounded adversaries[1] and has been black-box separated from one-way functions [IR90]. We briefly note that despite not all constructions and (security) reductions being black-box (e.g., [FFS87, Bar01]), still many techniques in cryptography are black-box in the sense that they use the primitive and adversary (in the security reduction) as a black box, and hence they hold relative to any oracle as well [IL89, RTV04]. In fact, non-black-box security reductions à la [Bar01] are less common in cryptography.

Minimal Assumptions in Quantum Cryptography. More recently, many foundational primitives have been constructed from pseudorandom quantum states, a loose quantum analog to classical pseudorandom generators [JLS18, AQY22, MY22]. Although pseudorandom states can be constructed from quantum one-way functions, the reverse direction is not known. In contrast, recent results have shown *oracle separations* between pseudorandom states and one-way functions. First, Kretschmer constructed a (quantum) oracle relative to which pseudorandom states exist but $\mathbf{BQP} = \mathbf{QMA}$ [Kre21], and then later Kretschmer et al. constructed a classical oracle relative to which (single-copy) pseudorandom states exist but $\mathbf{P} = \mathbf{NP}$ [KQST23]. Additional works have explored the notion of EFI pairs—pairs of quantum states that are efficiently samplable, statistically far, and computationally indistinguishable—as a candidate to represent the minimal computational hardness required for quantum cryptography [BCQ23]. In short, quantum cryptography might have a different assumption hierarchy than classical cryptography, and the quantum analogues of one-way functions might not be at the center.

Minimal Assumptions in Post-Quantum Cryptography. Relatively little work has explored the computational assumptions required for post-quantum cryptography, which is the hybrid of the two worlds of classical and quantum cryptography. In post-quantum cryptography, the protocols are still implemented classically, but they must be secure against quantum attackers. Post-quantum cryptography models the most accessible intermediate future where cryptographically relevant quantum computers exist but have

[1] Specifically, key agreement cannot be secure information theoretically, and it implies OWFs.

not yet reached the widespread adoption necessary to implement quantum cryptographic protocols. Given the upheaval of the assumption hierarchy brought about by quantum computation, it is reasonable to ask what that hierarchy looks like in the post-quantum world. With that, we state the primary question of this work.

Are one-way functions minimal for post-quantum cryptography?

The question above might seem to have a straightforward "yes" answer, but the non-triviality of this question stems from the fact that a reduction might rewind the adversary, potentially messing up the internal state of the adversary if it is quantum.

The minimality of one-way functions in the classical world derives from the fact that they are implied by most basic cryptographic primitives. That is, there are black-box reductions from one-way functions to most other primitives. With this in mind, we can rephrase the above question as a more concrete technical question.

If a classically secure cryptographic primitive $\mathcal{Q}$ implies one-way functions, does the post-quantum-secure variant of $\mathcal{P}$ imply (post-quantum) one-way functions as well?

1.1 Our Results

We answer this question (mostly) in the positive by showing that in a large variety of natural settings, one-way functions are minimal for post-quantum cryptography as well.

2-Message Games. We study the problem more generally by using a primitive $\mathcal{P}$ rather than OWFs, but we still restrict our focus to primitives whose own security is defined using a 2-message game, which is the case for OWFs. We sometimes refer to such 2-message games as *non-interactive* "assumptions."[2] That is, the security game of the primitive $\mathcal{P}$ models a non-interactive assumption, and this is crucial for our results. While this may seem restrictive, we observe that the security of many (if not most) foundational cryptographic primitives can be modeled in terms of 2-message games, including one-way functions, trapdoor one-way functions, semantically secure PKE, key agreement, etc.

Specifically, we show that if there exists a classical black-box reduction from some primitive $\mathcal{P}$ defined by a 2-message security game (e.g., OWFs) to some other primitive $\mathcal{Q}$, then there also exists a post-quantum (black-box) reduction from $\mathcal{P}$ to $\mathcal{Q}$. That is, if the existence of $\mathcal{Q}$ implies the existence of $\mathcal{P}$ in the classical world, then the same is true in the post-quantum world.

There is an important corollary of this result. Since one-way functions are defined by a 2-message game, then any primitive which classically implies one-way functions (through the black-box method above) also implies one-way functions in the post-quantum setting. In other words, the existence of post-quantum one-way functions lies at the base of the post-quantum computational assumption hierarchy, just like in the classical world. A more complete description is found in Section 4.

Since 2-message adversaries carry minimal state, it is perhaps not surprising that classical reductions to non-interactive assumptions lift to the post-quantum setting. In some sense, our results can be seen as a verification of this intuition. However, proving this to be the case turns out to be more nuanced and non-trivial than expected, and we believe the final result deserves a full proof. In fact, even though the adversary used in the reduction plays in a non-interactive game, it can still be used interactively by the reduction through both rewinding (while keeping the randomness) and resetting (to a fresh randomness). Additionally, in light of how drastically interaction with a quantum machine can change the computational assumption landscape, it should not be taken for granted that a reduction that interactively uses its adversary carries over unscathed from the classical setting to the post-quantum setting.

[2] Here the term assumption refers to the belief that no efficient adversary can win in this game.

Extension: Deterministic-Adversary Reductions. We additionally present a strengthening of this result, showing that even reductions that expect a *deterministic* adversary, which we call *deterministic-adversary reductions*, can be lifted to the post-quantum setting by imposing additional restrictions on the games involved in the reduction. The extension covers the following cases:

- First, we show that if $\mathcal{Q}$ has success threshold zero (i.e., the adversary only has to win with non-negligible probability), then our main result holds, and a deterministic-adversary reduction can be lifted to the quantum-adversary setting. This result requires the security reduction to call its adversary on a single security parameter (which is a mild restriction, as almost all reductions follow this restriction; see the work of [CLMP12] for discussions on why this is usually the case).

- Second, we show that even if $\mathcal{Q}$ has an arbitrary security threshold (e.g., threshold $\frac{1}{2}$ for indistinguishability games), then as long as either of the security games of $\mathcal{P}$ or $\mathcal{Q}$ is falsifiable [Nao03] (i.e., has a polynomial-time challenger), then a deterministic-adversary reduction can be lifted to the post-quantum setting. For the setting that $\mathcal{Q}$ is falsifiable, we again limit the reduction to call its adversary on a single security parameter, but when $\mathcal{P}$ is falsifiable (e.g., OWFs), such a limitation is not needed for the result to go through.

The motivation for the case of deterministic-adversary security reductions stems from the fact that classical reductions can benefit from such assumptions at the cost of requiring stronger (non-uniform) assumptions in the end. Namely, reductions can pretend to use the "best possible randomness" of their adversaries (e.g., one that is provided as non-uniform advice) and then benefit from this determinism. For example, the work of [IMS12] achieved their black-box sublinear malicious-verifier ZK arguments using this technique. As for previous work on such lower bounds, the work of [CLMP12] also achieved lower bounds against deterministic-verifier reductions for the different problem of ruling out provability of 2-Message Zero Knowledge.

Extension: Polynomial-Size (Non-Uniform) Quantum Adversaries. Our results also extend to cover the non-uniform quantum hardness regime. In particular, if we start from the implication that primitive $\mathcal{Q}$ implies $\mathcal{P}$ through a black-box game-based security reduction while $\mathcal{P}$ has a 2-message game, then we can lift this implication to cover the post-quantum variants of $\mathcal{P}$ and $\mathcal{Q}$. Indeed, we can do this lifting while we cover the *non-uniformly*-hard post-quantum variants in which the given adversary of $\mathcal{P}$ has non-uniform *quantum* advice. However, to obtain this implication, we pay the price of using *polynomially more* quantum advice during our (post-quantum) reduction.

We clarify two settings that our results *do not* cover. The first is the case where the quantum auxiliary input given to the reduction is an *arbitrary quantum function* of the oracle adversary and the oracle primitive, similar to classical non-uniform security reductions [CLMP12]. We emphasize that this setting is irrelevant to our results, because we start from a classical reduction with a classical interface, and hence the reduction itself cannot take quantum advice. The second case is where we are given only a single copy of the adversary (and of the quantum auxiliary input), known as constructive post-quantum reductions [BBK22]. In this case, the no-cloning theorem prohibits us from creating polynomially many copies of the arbitrary quantum auxiliary input $|s\rangle$.

1.2 Technical Overview

Our primary technical contribution is to prove the following. If there exists a classical reduction from winning in a game $\mathcal{Q}$ to winning in a 2-message game $\mathcal{P}$, then for any efficient quantum adversary $\mathcal{A}_Q$ who wins in $\mathcal{P}$, we can construct a new efficient quantum adversary $\mathcal{A}'_Q$ who wins in $\mathcal{Q}$.

We now describe slightly informally the techniques which allow us to lift reductions to 2-message games to the post-quantum setting. This represents the bulk of the technical contribution and is the crux in proving our results. Consider a reduction $\mathcal{R}$ from a game $\mathcal{Q}$ to a 2-message game $\mathcal{P}$, and let $\mathcal{A}_Q$ be an efficient quantum adversary that wins in $\mathcal{P}$. For now suppose $\mathcal{A}_Q$ is uniform and starts from the $|0\rangle$ state.[3] Even though $\mathcal{A}_Q$ plays in a non-interactive game, we cannot claim that $\mathcal{R}^{\mathcal{A}_Q}$ wins in $\mathcal{Q}$. Since $\mathcal{R}$ expects a classical adversary, by interacting with $\mathcal{A}_Q$, it may potentially be able to distinguish it from a classical adversary using techniques such as rewinding. In fact, it is not quite clear what it formally means to classically access a quantum algorithm $\mathcal{A}_Q$ under rewinding. What we want instead is a successful quantum adversary that is indistinguishable from *some* successful classical adversary that can be rewound, even though this classical adversary might be *inefficient*. Therefore, since $\mathcal{R}$ must be able to win in $\mathcal{Q}$ given any such successful classical adversary (even if it is inefficient), if it cannot tell that it was given a quantum adversary, then it must win in $\mathcal{Q}$ even with the quantum adversary. We remark that the running time of the adversary cannot be used to distinguish an inefficient classical adversary from an efficient quantum adversary, because both are provided as oracles and respond to queries instantaneously.[4]

This high-level plan is reminiscent of the so-called “two-oracle” technique [GW11, Pas11, Wic13]. To achieve the above goal, we propose modifications both to $\mathcal{A}_Q$ and a classical algorithm $(\mathcal{A}_Q)_C$, which is any inefficient classical algorithm that behaves (perfectly) identically to $\mathcal{A}_Q$ on a single query. Our modifications will essentially make both the classical and quantum adversaries useless to rewind, allowing us to argue that they are still indistinguishable from each other from the perspective of both the reduction *and* the 2-message game.

First, we construct the modified classical adversary $(\tilde{\mathcal{A}}_Q)_C$ by fixing an exponentially long random tape with an entry in the tape for each possible query and evaluating $(\mathcal{A}_Q)_C$ on each query, which makes it so the output of $(\tilde{\mathcal{A}}_Q)_C$ on any query is determined in advance (based on its long randomness) and identical for repeated queries. Next, we construct the modified quantum adversary $\mathcal{A}'_Q$ by lazy-evaluating $\mathcal{A}_Q$ and also memorizing the previously asked queries by storing the responses to previous queries. This modification makes $\mathcal{A}'_Q$ no longer an actual oracle (due to its memory), but that is fine for our purpose of successfully running a reduction, as the reduction does not know that its oracle is stateful. From here, we can argue that $(\tilde{\mathcal{A}}_Q)_C$ and $\mathcal{A}'_Q$ are also indistinguishable from the perspective of the reduction, even if the reduction behaves arbitrarily.

Deterministic-Adversary Reductions. We first briefly discuss the challenges that arise when using deterministic-adversary reductions. At first, it might seem that we can simply use the same adversary $\mathcal{A}'_Q$ to emulate the reduction $\mathcal{R}$. After all, doing so is identical to “picking and fixing a (very long) randomness” for the equivalent adversary $(\tilde{\mathcal{A}}_Q)_C$. For simplicity, let us now suppose that the reduction calls its oracle adversary only on a single security parameter $n_{\mathcal{P}}$.

Now, $\mathcal{R}$ treats $(\tilde{\mathcal{A}}_Q)_C$ as a deterministic adversary, and we want to analyze their interaction using some fixed randomness for the adversary. Let ε be the adversary's advantage as a randomized adversary. By fixing its randomness, we use an averaging argument to claim that with probability at least $\frac{\varepsilon}{2}$, the adversary has advantage at least $\frac{\varepsilon}{2}$.

While this is a non-negligible advantage, the $1 - \frac{\varepsilon}{2}$ probability of choosing a “bad” randomness prohibits us from making claims about the reduction's success in a game $\mathcal{Q}$

[3] We extend this argument to the case of non-uniform adversaries by using more advice in the reduction.

[4] Reductions that are almost black-box but can depend on the running time of the adversary are called “class reductions” [Sha20].

with an arbitrary "success threshold"; the success threshold τ is the trivial probability of winning in the game that the adversary wants to beat with non-negligible advantage. $\mathcal{R}$ wins in $\mathcal{Q}$ with probability $\frac{\varepsilon}{2} \cdot \delta(\frac{\varepsilon}{2}) + (1 - \frac{\varepsilon}{2}) \cdot \texttt{unknown}$, where $\delta(\gamma)$ is the corresponding non-negligible advantage of winning in $\mathcal{Q}$ if the oracle adversary wins in $\mathcal{P}$ with advantage γ. In a game $\mathcal{Q}$ with an arbitrary success threshold, the unknown component of $\mathcal{R}$'s advantage may be negative, so we cannot claim that $\mathcal{R}$'s advantage in total is non-negligible.

We now sketch the key ideas that we use to obtain this extension.

1. **$\mathcal{Q}$ is threshold-zero.** For the first case, the depiction of the challenge above already shows that if the security threshold of $\mathcal{Q}$ is zero and $\mathcal{R}$ only calls $\mathcal{A}$ on a single security parameter $n_{\mathcal{P}}$, then the advantage of emulating $\mathcal{R}$ using $\mathcal{A}'_Q$ will be at least $\frac{\varepsilon}{2} \cdot \delta(\frac{\varepsilon}{2})$, which is also non-negligible. This is the easiest case.

2. **$\mathcal{P}$ is falsifiable.** For this case, we cannot simply rely on the picked "randomness" of $\mathcal{A}'_Q$ (simulating a random choice of $(\tilde{\mathcal{A}}_Q)_C$) being good with probability $\frac{\varepsilon}{2}$. However, we know that *if* we get lucky, we can "test" that we have become lucky due to the falsifiability of $\mathcal{P}$. So, we try our luck many times until we succeed in finding a "good" instantiation of $(\tilde{\mathcal{A}}_Q)_C$ by resetting the internal table of $\mathcal{A}'_Q$ until hitting a good table. We do this for *every* time that $\mathcal{R}$ calls its adversary oracle on a new security parameter. Once we find a good table for each security parameter, we use the "good" adversaries in the interaction in the game $\mathcal{Q}$.

3. **$\mathcal{Q}$ is falsifiable.** For this case, we will use a similar idea to the one above, but instead of testing whether a table M for $\mathcal{A}'_Q$ is good by running $\mathcal{A}'_Q$ in $\mathcal{P}$, we must infer whether it wins in $\mathcal{P}$ with sufficient probability based on its success probability when the reduction is run in $\mathcal{Q}$ with oracle access to $\mathcal{A}'_Q$.

Polynomial-Size (Non-Uniform) Quantum Adversaries. We now explain the simple idea that allows us to obtain the extension to the case of non-uniform quantum hardness against adversaries that come with non-uniform advice. To explain the idea, we need an observation about the big picture of our baseline proofs for the uniform case.

At the core of the baseline proofs, we plug a uniform quantum adversary into a classical reduction to win a security game by running the adversary polynomially many times from scratch with quantum auxiliary input $|0\rangle$. However, in doing so, once an adversary starts an execution, we do not rewind it. Now, suppose instead that the quantum adversary has hard-coded quantum auxiliary input $|s\rangle$ to begin with. Since we allow arbitrary advice so long as it has polynomial length, we can simply begin with polynomially many copies of $|s\rangle$. Each time the proofs reset an adversary, we use a *fresh* (unentangled) copy of $|s\rangle$ (as opposed to a fresh copy of $|0\rangle$).

1.3 Further Discussion and Related Work

Constructive and Universal Reductions. Bitansky, Brakerski, and Kalai studied the question of which classical reductions lift to the post-quantum setting in the presence of *single-copy* adversaries, showing both feasibility and impossibility results [BBK22]. Namely, their focus is to understand which reductions carry over to the setting where we cannot make any copy of the adversary (or the quantum advice that it comes with). Additionally, Chan, Freitag, and Pass studied the notion of *universal reductions*, or reductions that hold for a wide variety of computational models while making minimal assumptions about which models are "physically realizable" [CFP22]. Again, they focus on settings where a fully stateful environment carries state across its interactions and hence sometimes cannot be reset. Similarly, they show feasibility and impossibility results in this setting. Much of the technical contribution of both [BBK22] and [CFP22] is dedicated to remedying the difficulties imposed by one-shot adversaries which we avoid due to our more relaxed model. Hence, the results of both works are incomparable to ours, as we achieve stronger implications in a more relaxed model.

Rewinding. As mentioned, a primary difficulty against boosting a classical reduction to the post-quantum setting is that the reduction may rewind its oracle. Quantum rewinding is challenging in the general case due to the no-cloning theorem [WZ82], although there have been numerous quantum rewinding algorithms that are proven to be effective in certain restricted scenarios [Wat06, Unr12, ARU14]. (See the tutorial [LM22] for an exposition of such subtleties.) Due to these difficulties, many existing works show positive results only for classical reductions which are *straight-line*, meaning they do not rewind their oracle [Son14, CFP22], or *non-adaptive*, meaning they cannot change their strategy based on the responses they get from their oracle [BBK22]. In contrast, we show results even for classical reductions which behave arbitrarily, but we impose restrictions on the games involved in the reduction.

Beyond 2-Message Games? The focus of this work is on understanding when one-way functions remain minimal in post-quantum cryptography. We study the more general setting when primitives with a 2-message security game remain minimal, as this covers OWFs. A reasonable question would be whether our results also extend to cases where $\mathcal{P}$ is defined by a security game with more than two messages. Namely, can we still conclude that "$\mathcal{Q}$ classically implies $\mathcal{P}$" can be lifted to the post-quantum setting?

Here we observe that our results demonstrably do not hold if we extend $\mathcal{P}$ to cover primitives with 4-message security games. In particular, consider the 4-message qubit certification protocol of [BCM$^+$21] used in their construction of a proof of quantumness. That protocol presents a challenge that is secure against classical polynomial-time provers if the LWE assumption holds. Furthermore, their security proof is done through a black-box reduction to the learning with errors (LWE) assumption. However, that classical reduction *necessarily* rewinds the adversary in a way that inherently invalidates the reduction in the post-quantum setting. Indeed, an efficient quantum prover *can* provably win that challenge, while LWE is conjectured to hold in the post-quantum setting as well. This means that such a reduction is unlikely to exist when the given adversary is quantum.

To connect this to our results, we can construct a contrived primitive $\mathcal{P}$ with a 4-message security game (based on the qubit certification test of [BCM$^+$21]) such that: (1) $\mathcal{P}$ is implied by LWE in the classical setting through a black-box security reduction, but (2) $\mathcal{P}$ cannot be obtained from LWE in the post-quantum setting through a black-box reduction, unless LWE can be broken in quantum polynomial time. The description of the primitive $\mathcal{P}$ is as follows. We simply consider the identity function as the only legitimate "implementation" for $\mathcal{P}$, while its security depends on whether an "efficient" adversary can win the 4-message security game.[5]

Stronger Ad-Hoc Post-Quantum Security Notions. In this work, we consider the post-quantum variant of classical primitives, where this post-quantum variant is derived automatically from the *same* (classical) security game of the classical primitives. One can imagine ad-hoc (perhaps *stronger*) post-quantum variants in which the quantum adversary can do *more* by asking superposition queries [BZ13]. However, as long as such notions are *stronger*, such primitives also imply the weaker version as defined by us and hence imply OWFs as well whenever the weaker notion does.

We finally comment that the work of Yamakawa and Zhandry [YZ20] also studies the question of "what reductions from the classical world carry to the post-quantum world" in the context of the random oracle model. However, their proof allows the new algorithm itself to remain quantum even after making their oracle queries classical.

[5] This primitive plausibly exists in the classical world but not in the post-quantum world.

2 Preliminaries

In this section we define the terms and ideas that we will use throughout the paper. A *negligible function* $f(n)$ is a function that decreases faster than any inverse polynomial function $1/poly(n)$ for sufficiently large n. We say that a probability is negligible if it is a negligible function of its parameter (which is usually the security parameter), and we say it is non-negligible otherwise. We write $x \leftarrow D$ to represent sampling x according to the distribution D. We begin by defining efficient classical and quantum algorithms. Both definitions are standard. We consider a classical efficient algorithm to be a probabilistic polynomial-time Turing machine and a quantum algorithm to be a family of efficiently generated quantum circuits. When we deal with computationally unbounded algorithms, we do *not* limit them to be implemented using Turing machines and similarly to [BBF13] model them with arbitrary distributions.

Definition 1 (Classical Algorithms)**.** A classical (perhaps inefficient) algorithm A is a collection of (potentially correlated) random variables, such that for any given input $x \in \{0,1\}^\star$, it outputs y from a specific definition, and we denote this process as $y = A(x; r_{|x|})$, where x is the input and $r_{|x|}$ is the (perhaps exponentially long or unbounded) source of randomness for input length n. Sometimes the randomness is indexed with "security parameter" n rather than the input length. An *efficient* classical algorithm is one that can be implemented by a probabilistic Turing machine that runs in time $poly(|x|)$ (and so it will also necessarily use $poly(|x|)$ bits of randomness). We sometimes refer to efficient algorithms as PPT machines for short. More generally, a PPT machine A with access to an oracle O is denoted by A^O and shall run in probabilistic polynomial-time for every O if one counts each query to the oracle as a single step.

Interactive Algorithms and Rewinding. Here we clarify how Definition 1 covers the notion of interactive algorithms and what it means to rewind interactive or even non-interactive algorithms provided as oracles. Interactive algorithms can use non-interactive algorithms to define their "next-message" function. Namely, the interactive randomized algorithm A first picks its randomness r_n for security parameter n. Then, given inputs $x_1, \dots, x_i$, it outputs the i answers $y_1, \dots, y_i$, or equivalently it only outputs the i^{th} answer y_i. Then, one can interact with A by iteratively obtaining the answers. To rewind A means to "go back in time" and ask a different branch of queries (e.g., (x_1, x'_2) after previously having asked (x_1, x_2)). Even a non-interactive algorithm can be rewound by asking it a different query x'_1 for the *same randomness* r_n. We also allow the randomness of an (even unbounded) algorithm to be reset to r'_n. Since the randomness can be of super-polynomial length, we allow the oracle algorithm A that uses oracle O to have a handler for each "randomness session." Namely, A can ask O to respond with respect to randomness handler h_1, and then change the handler to h_2, which means the randomness of O needs to be reset. Therefore, $O(x; h), O(x'; h)$ for $|x| = |x'| = n$ will be answered using the same randomness r_n, while $O(x, h), O(x, h')$ will be answered using independent randomness r_n, r'_n. When dealing with *efficiently computable* oracles, we can always allow the algorithm A to pick the randomness, but for inefficient algorithms we can provide the access indirectly as stated above, which is first formalized and used in [BJY97].

Loosely speaking, an efficient quantum algorithm is a family of efficiently generated quantum circuits with one circuit for each input size. A circuit for input size n has three sets of qubit registers for input, output, and work.

Definition 2 (Quantum Algorithm)**.** An efficient quantum algorithm Q is a family of quantum circuits $\{Q_n\}_{n\in\mathbb{N}}$, where the description of Q_n is computed by a classical PPT algorithm A given input 1^n. Q_n has 3 registers: the first n qubits receive the input, the second $q(n)$-qubit register is reserved for storing the output, and the third set of $q(n)$

qubits are reserved for the work during the computation, while $q(\cdot)$ is some polynomial function. On input $x \in \{0,1\}^n$, the initial quantum state is $|x\rangle\,|0\rangle\,|0\rangle$. At the end of computation, the second (output) register is measured in the computational basis to yield a classical output. The number of gates in the circuit for input length n is bounded as $|Q_n| \leq p(n)$, where $p(\cdot)$ is some polynomial function. A *non-uniform* efficient quantum algorithm Q is defined similarly, but the work registers are initialized with $|\phi\rangle_n$, where $|\phi\rangle_n$ is the quantum advice for input length n.

Since we are concerned with post-quantum cryptography, where the parties implementing the protocols are classical, and their "default" security game involves classical challengers and adversaries, we constrain quantum algorithms to work with only classical input and output. For this reason, we can define a *classical interface* of a (non-interactive) quantum algorithm Q as any (perhaps inefficient) randomized algorithm which has the same input-output distribution as that of Q. We denote a classical interface of Q by $(Q)_C$.

Definition 3 (Classical Interface of a Quantum Algorithm)**.** Let Q be a quantum algorithm with classical inputs and outputs. A classical interface $(Q)_C$ for Q is any (perhaps inefficient) classical randomized algorithm[6] such that for each possible input x, the distributions $Q(x)$ and $(Q)_C(x)$ are identical.

Note that in the definition above $(Q)_C$ is not necessarily unique for a fixed Q, as there might be different ways to sample $Q(x)$. Moreover, this definition does not state how the answers $(Q)_C(x)$ and $(Q)_C(x')$ are correlated. One algorithm might use smaller randomness that leads to correlated answers, while another algorithm uses more randomness and leads to independent answers. In fact, this very point (that $(Q)_C$ is not unique) points to the subtle point that rewinding a classical interface of a quantum algorithm might lead to different behavior.

We now define the notion of an interactive game between a challenger and an adversary.

Definition 4 (Security Games)**.** A security game is an interactive game between two machines, a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$, in which k messages are sent between $\mathcal{C}$ and $\mathcal{A}$, n is the common input referred to as the security parameter, and the last message is always sent by the adversary. $\mathcal{C}$ and $\mathcal{A}$ each have their own independent randomness, $r_{\mathcal{C}}$ and $r_{\mathcal{A}}$, respectively. The game is merely described by the description of the challenger $\mathcal{C}$. Suppose ℓ is the total size of its inputs given to the challenger by an adversary. We call the game *falsifiable* if $\mathcal{C}$ runs in polynomial time over $(n\ell)$ and simply round efficient if $k \leq \mathit{poly}(n\ell)$.[7] At the end, $\mathcal{C}$ outputs 1 if $\mathcal{A}$ has won this particular execution of the game and outputs 0 otherwise. We assign a "success threshold" $t(n) \in [0,1)$ to any such game, and we say that adversary $\mathcal{A}$ has advantage $\varepsilon(n)$ in (winning) the game if it wins with probability $t(n) + \varepsilon(n)$. In this case, we call the game a threshold-$t(n)$ game.

In the following, we assume that games come with security thresholds, unless stated otherwise. Intuitively, a game is falsifiable if it is possible to efficiently simulate the role of the challenger by sampling a challenge and verifying a response. This means that one can efficiently verify if the adversary could win and "break" the computational assumption behind the security game. This was first proposed by Naor [Nao03].

Usually, security games of cryptographic primitives are either threshold-zero (like OWFs or any other search problem) or threshold-($\frac{1}{2}$) (like PRGs or any other distinguishing game), and in all of these settings the adversary's goal is to win with non-negligible advantage.

A *2-message game* is a special case of a k-message game where $k = 2$ and where the only messages are the challenger's challenge followed by the adversary's response. We might refer to such games as non-interactive games as well.

[6] Recall that we use the term algorithm generically, without requiring the algorithm to be implemented using any specific computing model such as Turing Machines.

[7] Note that a special case is to have $k \leq \mathit{poly}(n)$, but for security games like that of digital signatures, the number of rounds might be chosen by the polynomial time adversary.

In this work, in both the classical and post-quantum settings, the challenger is a classical algorithm. In the classical setting, the adversary is a classical algorithm, and in the post-quantum setting, the adversary is a quantum algorithm with a classical interface. We comment that post-quantum security does *not* have to have such a limitation in general, and the security games of post-quantum-secure schemes could have their own quantum-tailored definitions, but such primitives will not have a direct classical counterpart primitive, as it is not automatically clear what is the security game of the corresponding classical primitive. However, by merely allowing a quantum adversary to play in the classical security game of the original classical primitive $\mathcal{Q}$, we can talk about the post-quantum variant of $\mathcal{Q}$.

Definition 5 (Black-Box Reduction Between Games)**.** An (ε, δ)-reduction $\mathcal{R}$ from (winning in) the game G_2 to (winning in) the game G_1 is a PPT oracle algorithm that, given access to any oracle adversary $\mathcal{A}$ which wins in G_1 with advantage $\varepsilon(n_1)$ (for all security parameters n_1), $\mathcal{R}^{\mathcal{A}}$ wins in G_2 with advantage $\delta(n_2)$, where n_2 is the security parameter of G_2. $\mathcal{R}$ can call $\mathcal{A}$ on multiple security parameters, rewind it, or reset its randomness, but if it only calls it on security parameter n_1 (that is determined by the security parameter of G_2, n_2), we call this a single-security-parameter reduction. We also give $1^{1/\varepsilon}$ as an input parameter to $\mathcal{R}$.

Giving $1^{1/\varepsilon}$ as an input parameter to $\mathcal{R}$ is to make the definition non-vacuous by allowing the running time of $\mathcal{R}$ to depend on its adversary.[8]

Syntax of Oracle Queries. In Definition 5 the reduction is allowed to call its oracle adversary on different security parameters, rewind it, or reset the adversary's randomness. This can be syntactically enabled by letting $\mathcal{R}$ call $\mathcal{A}$ on $(x, n; h)$ where x is the "input", n is the security parameter, and h is the randomness handler. x itself could be a vector $(x_1, \ldots, x_i)$ for interactive adversaries. Therefore, $\mathcal{R}$ can rewind $\mathcal{A}$ by keeping the same randomness handler h and asking a different branch of queries, like $(x_1, \ldots, x_{i-1}, x'_i)$, and it can reset the randomness of $\mathcal{A}$ by changing the randomness handler h.

In general, as mentioned in the definition, the oracle adversary $\mathcal{A}$ that wins in G_1 can model a randomized algorithm that allows fixing or resetting its randomness. If the reduction assumes that the oracle is deterministic, we call it a *deterministic-adversary reduction.* Note that such an assumption might be beneficial for the reduction, as it limits the way the oracle adversary can perform.

Non-Negligible-Advantage Security Reductions. In general, δ could be a function of both ε and the security parameter n_2. By default, security reductions can use any non-negligible $\varepsilon \geq 1/poly(n_1)$ and will produce $\delta_{\varepsilon,n_2} \geq 1/poly(n_2)$ that is non-negligible. From this point forward, we always assume that we work with a non-negligible-advantage security reduction, and for simplicity we will refer to the advantage in G_1 as ε and the advantage in G_2 as δ_ε, leaving the dependence on security parameters implicit.

Next, we will define a cryptographic primitive, following [RTV04] but also with some refinements from [BBF13] that model implementations as random variables, allowing inefficient implementations with long randomness. We consider only primitives implemented by classical algorithms, but they can be both classically secure and quantum-secure. In other words, we consider classical primitives and post-quantum primitives (those implemented classically with classical or quantum security, respectively) but not quantum primitives (those implemented by quantum algorithms).

[8] Otherwise, if ε is too small, the reduction might not even get one "useful" execution out of its adversary. See [BBF13] for more discussions.

Definition 6 (Classical Primitives with Classical Security)**.** A cryptographic primitive $\mathcal{P}$ is a pair $(F_{\mathcal{P}}, R_{\mathcal{P}})$. $F_{\mathcal{P}}$ is a set of potentially inefficient randomized algorithms (as defined in Definition 1) implementing $\mathcal{P}$. An implementation f is considered (classically) *efficient* if it can be computed by a classical PPT algorithm. $R_{\mathcal{P}}$ is a relation with pairs of the form (f, A) in which both f, A are potentially inefficient randomized algorithms. For an implementation $f \in F_{\mathcal{P}}$, an adversary A is said to $\mathcal{P}$-break f if $(f, A) \in R_{\mathcal{P}}$. $f \in \mathcal{P}$ is a secure implementation if no PPT algorithm $\mathcal{P}$-breaks f. The primitive $\mathcal{P}$ *exists* if there exists an efficient and secure implementation of $\mathcal{P}$.

Post-Quantum Primitives. Definition 6 above can be extended to classical primitives with quantum security (i.e., the post-quantum setting). To do so, one shall allow adversaries A to be quantum algorithms with a *potentially quantum interface*. However, as mentioned before, such post-quantum primitives might not have a direct classical variant when it comes to their security definition. One can further limit the adversaries to have classical interfaces. Indeed, even though a family of primitives, like public-key encryption with various security definitions such as CPA and CCA, can include post-quantum variants in which the security definition is fully tailored to quantum adversaries (e.g., by allowing superposition queries to be asked by the adversary [BZ13]), here we aim to define what it means to take a classical primitive and systematically derive its post-quantum variant. We now define the post-quantum variant of a classical primitive that falls into this category.

Definition 7 (The Post-Quantum-Secure Variant of a Game-Based Classical Primitive)**.** Suppose $\mathcal{P}$ is a classical primitive. Suppose further that the security relation of $\mathcal{P}$ is defined through a security game with challenger $\mathcal{C}$ that can depend on the implementation f of $\mathcal{P}$; namely, $(f, A) \in R_{\mathcal{P}}$ iff A wins in $\mathcal{C}_f$ with non-negligible advantage (above its success threshold). Then, we define the post-quantum variant of $\mathcal{P}$ to be the same primitive, but now the security definition allows efficient quantum adversaries to participate in the same security game (using classical communication).

Black-box security reductions between games are a special way of designing black-box security proofs as depicted in [RTV04]. The reason is that the notion of a primitive as defined above (based on [RTV04]) allows the security definition to be arbitrary, while we focus on security definitions that are in the form of interactive security games. This means that our results do not depend on how the *implementation reduction* is performed, as formalized in [RTV04], meaning that we can handle non-black-box constructions, so long as the security reduction is a black-box reduction between games.

The definition below first formalizes what it means for a primitive's existence to imply that of another without emphasizing any of the components being black-box. This formalization is also presented for the post-quantum setting (in the second bullet). However, the definition formalizes along the way what it means to have a free (non-constructive) implementation reduction while the security reduction is still black-box. Although we could also focus on defining the security reductions for the post-quantum setting in a constructive way, we use the simpler definition as our goal is merely showing that classical implications can be lifted to the post-quantum setting.

Definition 8 (Potentially Non-constructive Reductions)**.** A (potentially non-constructive) reduction exists from a classical primitive $\mathcal{P} = (F_{\mathcal{P}}, R_{\mathcal{P}})$ to a classical primitive $\mathcal{Q} = (F_{\mathcal{Q}}, R_{\mathcal{Q}})$ if the following two conditions hold:

1) Implementation. For every PPT implementation $f \in F_{\mathcal{Q}}$, there is a PPT implementation $g_f \in F_{\mathcal{P}}$.

2) Security. For every PPT implementation $f \in F_{\mathcal{Q}}$, if there is a PPT adversary $\mathcal{A}_{\mathcal{P}}$ that $\mathcal{P}$-breaks g_f, then there is a PPT adversary $\mathcal{A}_{\mathcal{Q}}$ that $\mathcal{Q}$-breaks f.

Furthermore, if $\mathcal{P}$ and $\mathcal{Q}$ have security definitions that are based on security games with challengers $C_{\mathcal{P}}, C_{\mathcal{Q}}$, we define the following two cases:

- **Game-Based Black-Box Security Reduction.** We say that the reduction from $\mathcal{P}$ to $\mathcal{Q}$ has a black-box game-based reduction if the existence of $\mathcal{A}_{\mathcal{Q}}$ above (based on $\mathcal{A}_{\mathcal{P}}$) is proved through a black-box reduction from winning in $C_{\mathcal{Q},f}$ with non-negligible advantage to winning in $C_{\mathcal{P},g_f}$ with non-negligible advantage. Note that such a reduction shall work even if the adversary is *not efficient*.
- **Security Reduction for the Post-Quantum Setting.** Let $\mathcal{P}'$ and $\mathcal{Q}'$ be the post-quantum variants of $\mathcal{P}$ and $\mathcal{Q}$, respectively. Then, we say that there is a reduction from $\mathcal{P}'$ to $\mathcal{Q}'$ if (in addition to the implementation reduction above) it holds that whenever there is an efficient *quantum* adversary $\mathcal{A}_{\mathcal{P}}$ that $\mathcal{P}$-breaks g_f, then there is an efficient quantum adversary $\mathcal{A}_{\mathcal{Q}}$ that $\mathcal{Q}$-breaks f. The notion of "efficient quantum algorithm" here could be either uniform or non-uniform, leading to two different definitions.

3 Lifting Reductions to 2-Message Games

In this section, we prove our primary technical results. At a high level, we show that if there is a classical reduction from winning in some game $\mathcal{Q}$ to winning in a 2-message game $\mathcal{P}$, then given any efficient quantum adversary that wins in the 2-message game $\mathcal{P}$, we can use this quantum adversary to win the second game in quantum polynomial time. This holds both for uniform and non-uniform notions of quantum efficiency.

3.1 Useful Tools

Before turning to the main results, it will be helpful to discuss ways of modifying adversaries to be more amenable to analysis in a reduction. Specifically, we wish to show that quantum adversaries can be modified in such a way that classical reductions cannot distinguish them from an equally successful classical adversary. Given a quantum adversary that wins in some 2-message game but which may not be easily usable by a classical reduction, we construct both a modified quantum adversary and a corresponding classical adversary such that both have equal advantage in the game and are indistinguishable from each other, but a classical reduction can effectively use the classical adversary. From this, we conclude that the adversary can effectively use the modified quantum adversary as well.

We will talk at several points about inefficient classical adversaries. While it is important for our quantum adversaries to be efficient, it is not an issue if the constructed classical adversaries are inefficient, so long as the reduction is supposed to work with such adversaries as well. The only purpose of the classical constructions is to show that a certain quantum adversary is indistinguishable from *some* useful classical adversary. That is, the inefficient classical adversaries are no more than a proof device.

To begin, we observe that the classical interface of a quantum adversary performs equally well to the quantum adversary. When we later aim to substitute such an adversary into a *reduction*, it becomes crucial that the original game (that the reduction is made to) is 2-message, as otherwise, the reduction's rewinding might make a distinction between the two adversaries.

Lemma 1. *Let $\mathcal{A}_Q$ be a quantum adversary that wins a 2-message security game $\mathcal{G}$ with probability ρ. Then any classical interface $(\mathcal{A}_Q)_C$ of $\mathcal{A}_Q$ also wins in the same game $\mathcal{G}$ with the same exact probability ρ.*

Proof. The proof is immediate considering that in the eyes of the challenger there is no difference between the two adversaries as long as $(\mathcal{A}_Q)_C$ is not rewound. □

A classical interface $(\mathcal{A}_Q)_C$ is not particularly useful in making claims about quantum behavior when we want to combine this adversary with a reduction. A classical reduction may be able to effectively use $(\mathcal{A}_Q)_C$, but this provides no immediate guarantee on its effectiveness in using the original $\mathcal{A}_Q$, since there is no guarantee about the similarity of the behavior of $\mathcal{A}_Q$ and $(\mathcal{A}_Q)_C$ under rewinding.

Always Resetting Algorithms. We now describe a method of constructing a new classical inefficient adversary that is useful for winning a 2-message game $\mathcal{G}$, starting from any classical adversary in $\mathcal{G}$ with the same advantage. The new adversary, however, uses *fresh independent* randomness for each independent query. Roughly speaking, this adversary is "reset all the time" for each query asked. In particular, let $\mathcal{A}_C$ be any classical adversary which wins in $\mathcal{G}$ with some advantage ε. $\mathcal{A}_C$ is a potentially randomized algorithm with some arbitrary correlation between its answers to queries under the same security parameter. (For example, $\mathcal{A}_C$ could be a classical interface for a quantum algorithm $\mathcal{A}_Q$.) On challenge message α from the challenger, we denote the distribution of responses by $\mathcal{A}_C(\alpha)$. We write $\mathcal{A}_C(\alpha; r)$ to denote $\mathcal{A}_C$'s response on input α where we fix its randomness to r.

We now define the new inefficient classical adversary $\tilde{\mathcal{A}}_C$ as follows. $\tilde{\mathcal{A}}_C$ uses an *exponentially long* random tape T, with an entry $T[\alpha]$ in the random tape for each possible query α. For every query α, $\tilde{\mathcal{A}}_C$ computes $\beta = \mathcal{A}_C(\alpha, T[\alpha])$ and returns β. $\tilde{\mathcal{A}}_C$ is described in Algorithm 1 below. As we observe later, $\tilde{\mathcal{A}}_C$ has the same advantage in $\mathcal{G}$ as $\mathcal{A}_C$ does.

Algorithm 1: Construction of $\tilde{\mathcal{A}}_C$

```
Input: $\mathcal{A}_C$, with randomness space $\Lambda_n$ for security parameter $n$
Initialize empty $T_n = [\,]$ for each $n$, and sample it as follows:
for each possible query $(\alpha, n)$ for security parameter $n$ do
    $T_n[\alpha] \leftarrow \Lambda_n$
end
Answer queries as follows:
for each query $(\alpha, n)$ for security parameter $n$ do
    $\beta \leftarrow \mathcal{A}_C(\alpha; T_n[\alpha])$
    Return $\beta$
end
```

One interesting point about $\tilde{\mathcal{A}}_C$ is that it is an inefficient algorithm, even if the original $\mathcal{A}_C$ is efficient. Another point is that Λ_n could be an infinite space if the original algorithm $\mathcal{A}_C$ is inefficient (e.g., an inefficient classical interface of an efficient quantum algorithm).

Remark 1 (Resetting Randomness of $\tilde{\mathcal{A}}_C$). Although $\tilde{\mathcal{A}}_C$ already uses independent randomness for each of the queries, a reduction that uses $\tilde{\mathcal{A}}_C$ as an oracle is still allowed to reset its (very long) randomness to a fresh value. This is doable exactly in the way that *any* algorithm can be reset. Namely, we we can provide an extra parameter h along with each query (α, n) that is the "randomness handler," and changing h to a new value means resetting randomness. In that case, each table of algorithm $\tilde{\mathcal{A}}_C$ needs to be parameterized with h as well (in addition to n). For simplicity of the presentation, we skip this parameter in Algorithm 1, but it shall be added whenever a reduction uses this algorithm.

Lemma 2. *If a classical adversary $\mathcal{A}_C$ wins a 2-message game $\mathcal{G}$ with probability ρ, then the classical adversary $\tilde{\mathcal{A}}_C$ constructed from $\mathcal{A}_C$ wins with the same exact probability.*

Proof. The proof is immediate, considering that in a 2-message game, $\tilde{\mathcal{A}}_C$ and $\mathcal{A}_C$ are asked only one query, and for a single query their answers are statistically identical. □

Lazy Evaluation of Quantum Oracle Algorithms. Next, we describe a method of modifying a quantum adversary to return identical outputs to repeated inputs to simulate the specific algorithm $(\tilde{\mathcal{A}}_Q)_C$ rather than $(\mathcal{A}_Q)_C$, where $(\mathcal{A}_Q)_C$ is an arbitrary classical interface of $\mathcal{A}_Q$. $(\tilde{\mathcal{A}}_Q)_C$ can be obtained from any classical interface using the modification of Algorithm 1.

Let $\mathcal{A}_Q$ be any quantum adversary. We construct a new quantum adversary $\mathcal{A}'_Q$ that lazy-evaluates $\mathcal{A}_Q$ on each query and stores responses to previous queries. At the beginning of execution, $\mathcal{A}'_Q$ defines an empty mapping M from queries to responses. Upon each query α,[9] $\mathcal{A}'_Q$ checks whether α exists in M. If $M[\alpha]$ is defined, $\mathcal{A}'_Q$ returns $M[\alpha]$. If $M[\alpha]$ is not defined, $\mathcal{A}'_Q$ computes $\beta \leftarrow \mathcal{A}_Q(\alpha)$, stores $M[\alpha] = \beta$, and returns β. $\mathcal{A}'_Q$ is described in Algorithm 2 below. We note that the construction of $\mathcal{A}'_Q$ makes only black-box use of $\mathcal{A}_Q$. After describing the uniform case, we describe how $\mathcal{A}'_Q$ shall be modified when $\mathcal{A}_Q$ comes with a hardcoded quantum auxiliary input $|\phi_n\rangle$ for security parameter n.

Algorithm 2: Construction of $\mathcal{A}'_Q$

```
Input : Uniform quantum algorithm $\mathcal{A}_Q$ with classical inputs and outputs
Initialize empty $M_n = [\,]$ for each $n$
for each classical query $(\alpha, n)$ for security parameter $n$ do
    if $M_n[\alpha]$ is already defined then
        Return $M_n[\alpha]$
    else
        $\beta \leftarrow \mathcal{A}_Q(\alpha)$
        $M[\alpha] = \beta$
        Return $\beta$
    end
end
```

Resetting Randomness of $\mathcal{A}'_Q$. Similarly to Remark 1, once $\mathcal{A}'_Q$ is used by a reduction, we can allow "resetting" its randomness by allowing an extra parameter h as the randomness handler, so whenever h changes, the table M_n needs to be filled from scratch, hence we add the parameter h next to n.

When $\mathcal{A}_Q$ Comes with Quantum Advice. Whenever the algorithm $\mathcal{A}_Q$ comes with quantum advice $|\phi_n\rangle$, we shall enhance $\mathcal{A}'_Q$ with *polynomially many copies* of this advice. Namely, we will start with $|\phi_n\rangle^m = |\phi_n^1\rangle \otimes \cdots \otimes |\phi_n^m\rangle$, where each $|\phi_n^i\rangle$ is a copy of $|\phi_n\rangle$ and m is an upper bound on the number of queries asked to $\mathcal{A}'_Q$. Then, whenever $\mathcal{A}'_Q$ wants to run $\mathcal{A}_Q$ on a *new* input (α, n, h), it will use the next unused copy $|\phi_n^i\rangle$. This of course means that $\mathcal{A}'_Q$ uses polynomially larger quantum advice, however this is not an issue for our final conclusions. Namely, we can still conclude that if there is a non-uniform algorithm winning in a 2-message challenge of $\mathcal{P}$, we can use it (and more advice) and win in another challenge of $\mathcal{Q}$.

Below, we show that for a given quantum adversary $\mathcal{A}_Q$, the corresponding adversaries $(\tilde{\mathcal{A}}_Q)_C$ and $\mathcal{A}'_Q$ are statistically indistinguishable. More importantly, this holds *even under rewinding.* The reason is that since both the classical modified adversary and the quantum modified adversary will always return identical answers to repeated queries, it is useless to rewind either of them under the same randomness, and even if their randomness is reset, their behavior continues identically. In effect, we use the structure of 2-message games to eliminate the main technical hurdle caused by the quantum-ness of the original $\mathcal{A}_Q$.

[9] Every query comes with a security parameter and option for resetting randomness; for now we ignore these extra inputs.

Lemma 3. *Let $\mathcal{A}_Q$ be any quantum algorithm with a classical interface $(\mathcal{A}_Q)_C$. Let $(\tilde{\mathcal{A}}_Q)_C$ be constructed from $(\mathcal{A}_Q)_C$ as described in Algorithm 1, and let $\mathcal{A}'_Q$ be constructed from $\mathcal{A}_Q$ as described in Algorithm 2. Then:*

1. *$(\tilde{\mathcal{A}}_Q)_C$ and $\mathcal{A}'_Q$ win in their 2-message game with the same probability.*
2. *In the eyes of a reduction $\mathcal{R}$ who can rewind and reset its oracle (which is an adversary for a 2-message game), $\mathcal{A}'_Q$ is identical to $(\tilde{\mathcal{A}}_Q)_C$.*

Proof. The first part follows from the second part, as $\mathcal{R}$ could run the challenger $\mathcal{C}$ of the 2-message game and ask a single query. So, we only prove the second part.

To see why Part 2 holds, it is enough to note that $\mathcal{A}'_Q$ is simply doing a lazy evaluation, so if we run $\mathcal{A}'_Q$ on all possible queries once and record all their answers once in a huge table M, its behavior against the reduction $\mathcal{R}$ would remain the same. Finally, this variant of $\mathcal{A}'_Q$ is doing exactly what $(\tilde{\mathcal{A}}_Q)_C$ does. □

3.2 Main Result

We now present the main result of the paper. Given a classical reduction from winning some game to winning another 2-message game, and given some efficient quantum adversary who wins in the 2-message game with non-negligible advantage, we show how to construct a new efficient quantum adversary that the classical reduction can successfully use.

Theorem 1 (Main Result)**.** *If there exists a classical reduction $\mathcal{R}$ from winning in security game $\mathcal{Q}$ to winning in 2-message game $\mathcal{P}$, and if there exists an efficient uniform (resp. non-uniform) quantum adversary $\mathcal{A}_Q$ who wins in $\mathcal{P}$ with non-negligible advantage, then there exists another efficient uniform (resp. non-uniform) quantum adversary that wins in $\mathcal{Q}$ with non-negligible advantage.*

Intuition. Before presenting the formal proof, we first describe the intuition and chain of thoughts that lead to the proof of Theorem 1.

We begin with a classical reduction $\mathcal{R}$ which is given oracle access to a classical adversary $\mathcal{A}_C$ that has advantage $\varepsilon \geq 1/poly(n_{\mathcal{P}})$ in the 2-message game $\mathcal{P}$. Suppose $\delta_\varepsilon \geq 1/poly(n_{\mathcal{Q}})$ is the probability that $\mathcal{R}$ succeeds in winning in game $\mathcal{Q}$ whenever it is given a classical ε-successful adversary in $\mathcal{P}$.[10]

Our goal is to show that we can construct some ε'-advantage quantum adversary $\mathcal{A}'_Q$ for $\varepsilon' \approx \varepsilon$ such that we can effectively run $\mathcal{R}^{\mathcal{A}'_Q}$ and show that it can still win in $\mathcal{Q}$ with advantage δ_ε. Note that in order to win against $\mathcal{Q}$ it is permissible to use any stateful adversary who might internally run efficient quantum computation.

We observe that since the game $\mathcal{P}$ is a 2-message game, any classical interface of the quantum adversary, denoted by $(\mathcal{A}_Q)_C$, wins in $\mathcal{P}$ with the same advantage ε as well (by Lemma 1). Therefore, a first attempt might be to try to show that for any quantum adversary $\mathcal{A}_Q$ with advantage ε in $\mathcal{P}$, the reduction $\mathcal{R}^{(\mathcal{A}_Q)_C}$ still wins in $\mathcal{Q}$ with advantage δ_ε, just like it does when given a classical adversary. Such an approach, however, suffers when the reduction $\mathcal{R}$ might want to “rewind” $(\mathcal{A}_Q)_C$ and ask it to answer a new query using the same randomness used to answer the previous query.

Instead, we use the modified version of the classical interface such that the new classical adversary retains an ε advantage in $\mathcal{P}$, while it is clear how to rewind and reset it during the reduction. This means we can still run an efficient quantum adversary that is indistinguishable from the new classical adversary, allowing $\mathcal{R}$ to use it successfully. We modify the quantum (oracle) adversary to make it look like a deterministic algorithm by obtaining $\mathcal{A}'_Q$ as described in Algorithm 2. However, we are still left with the problem

[10] As mentioned before, δ can also depend on the security parameter, but for simplicity of notation we only denote its dependence on ε explicitly.

of proving that $\mathcal{R}$ can successfully use this modified quantum adversary, which can be resolved by showing that $\mathcal{A}'_Q$ is indistinguishable from a successful classical adversary. For this classical adversary, we use the modified classical interface $(\tilde{\mathcal{A}}_Q)_C$ which is a classical adversary that is both useful in winning the 2-message game and is indistinguishable from $\mathcal{A}'_Q$ even under rewinding (see Lemma 2 and Lemma 3). We now give a direct proof.

Proof of Theorem 1. Suppose $\mathcal{R}$ is a reduction from winning in $\mathcal{Q}$ to winning in $\mathcal{P}$ that uses any ε-advantage classical adversary in $\mathcal{P}$ as an oracle and wins in $\mathcal{Q}$ with advantage δ_ε. Let $\mathcal{A}_Q$ be any quantum algorithm who wins in $\mathcal{P}$ with advantage ε.

Emulating $\mathcal{R}$. We now describe an algorithm $\mathcal{B}$ that uses $\mathcal{R}$ and $\mathcal{A}_Q$ to win against $\mathcal{Q}$. The algorithm $\mathcal{B}$ emulates $\mathcal{R}$, and whenever it asks a query from its oracle, $\mathcal{B}$ uses the *stateful* algorithm $\mathcal{A}'_Q$ to answer the query. We now elaborate on a few points.

- $\mathcal{R}$ is allowed to call $\mathcal{A}'_Q$ on any security parameter that it wants. However, recall that whenever $\mathcal{A}'_Q$ is launched on a new security parameter n, we initiate the stateful mapping M of $\mathcal{A}'_Q$ as empty.

- Similarly, $\mathcal{R}$ can "reset" the randomness of its oracle (even for the same security parameter n that it has previously queried). This can be done using randomness handlers that are given as an extra parameter along with the security parameter n. Again, in this case, we shall initiate $\mathcal{A}'_Q$ using a fresh table.

- Therefore, the reduction $\mathcal{R}$ can indeed switch its queries between security parameters and different reset instantiations of its oracle for the same security parameter, while using parameters n, h that accompany a query α.

- Finally, if the algorithm $\mathcal{A}_Q$ comes with quantum advice $|\phi_n\rangle$, then the emulation $\mathcal{B}$ uses a fresh copy of $|\phi_n\rangle$ for each instance of $\mathcal{A}'_Q$ and each reset instance of $\mathcal{A}'_Q$.

We now claim that this emulation of $\mathcal{R}$ (using quantum computation) is supposed to win its challenge with advantage δ_ε. Roughly speaking, we will argue that $\mathcal{R}$ can pretend that it is using $(\tilde{\mathcal{A}}_Q)_C$ and that $(\tilde{\mathcal{A}}_Q)_C$ shall be helpful to let $\mathcal{R}^{(\tilde{\mathcal{A}}_Q)_C}$ win its challenge. For simplicity, let's assume that $\mathcal{R}$ does not change the randomness of its oracle to a new value and that it only works with a single randomness and a single security parameter. The argument will be the same when these restrictions do not exist. In this simpler setting, by Lemma 3, $\mathcal{R}$'s behavior will be identical when we switch between giving it oracle access to $(\tilde{\mathcal{A}}_Q)_C$ such as $\mathcal{R}^{(\tilde{\mathcal{A}}_Q)_C}$ or emulating it as above using the algorithm $\mathcal{B}$. Moreover, by Lemma 2 and Lemma 1, $(\tilde{\mathcal{A}}_Q)_C$ indeed has advantage ε in winning its challenge, and hence it is a useful adversary that leads to $\mathcal{R}^{(\tilde{\mathcal{A}}_Q)_C}$ having advantage δ_ε. □

3.3 Extension to Deterministic-Adversary Reductions

We now prove an extension to the main result of the previous subsection.

Assuming that a given adversary is deterministic is a useful technique in proving security and works as follows. Suppose you have a reduction $\mathcal{R}$ that uses an ε-advantage adversary, where the advantage is also over the randomness of the adversary. Then, one can *always* "fix" the randomness of the adversary to "its best" leading to an adversary that is deterministic and still has advantage ε. Sometimes such a determinism assumption comes with benefits during the security proof (e.g., see [IMS12] for an example). Therefore, in light of such techniques, one can wonder if our results cover such reductions as well.

Implications of Deterministic-Adversary Reductions. Before answering the question above, we shall clarify: how can we find such "best" randomness for the adversary efficiently? After all, the adversary's randomness could even be not polynomial size. The answer is that: (1) we can fix the adversary's randomness to "its best" theoretically and provide oracle access to the reduction, even if it is not polynomial size, and (2) when the adversary *is* a polynomial time algorithm, it will inevitably have a polynomial-size random seed as well, in which case we can use "best randomness" provided as non-uniform advice. This means that we end up with a non-uniform proof of security [CLMP13]. Non-uniform proofs of security are useful for proving non-uniform implications. Therefore, any reduction that assumes its adversary to be deterministic still yields that the non-uniform variant of its computational assumption (i.e., that breaking $\mathcal{Q}$ non-uniformly is hard) implies that breaking $\mathcal{P}$ is (also non-uniformly) hard as well.

Definition 9 (Black-Box Deterministic-Adversary Reductions)**.** An (ε, δ)-reduction $\mathcal{R}$ from winning game G_2 to winning game G_1 is *deterministic-adversary* if it is defined similarly to Definition 5 with the difference that the reduction $\mathcal{R}$ is only required to succeed for deterministic adversaries. Namely, it will win its game G_2 with advantage δ whenever the oracle adversary $\mathcal{A}$ who wins in G_1 with advantage ε is also deterministic (i.e., resetting its randomness will not change its answers). In particular, if $\mathcal{A}$ is a *randomized* adversary who wins with advantage ε over its own randomness, $\mathcal{R}^{\mathcal{A}}$ comes with no immediate guarantees. As usual, also give $1^{1/\varepsilon}$ as an input parameter to $\mathcal{R}$.

We now state the extension to our main result.

Theorem 2 (Extending Theorem 1 to Deterministic-Adversary Reductions)**.** *Suppose there is a classical reduction $\mathcal{R}$ from winning in the game $\mathcal{Q}$ with non-negligible advantage to* deterministically *winning in the 2-message game $\mathcal{P}$ with non-negligible advantage. Additionally suppose either of the following three cases hold:*

1. *$\mathcal{Q}$ is threshold-zero, and $\mathcal{R}$ only calls $\mathcal{A}$ on a single security parameter.*

2. *$\mathcal{P}$ is falsifiable. (There is no restriction on calling $\mathcal{A}$ on multiple security parameters.)*

3. *$\mathcal{Q}$ is falsifiable, and $\mathcal{R}$ only calls $\mathcal{A}$ on a single security parameter.*

Then, in all three cases above, the existence of any efficient uniform (resp. non-uniform) quantum adversary $\mathcal{A}_Q$ who wins in $\mathcal{P}$ with a non-negligible advantage implies the existence of an efficient uniform (resp. non-uniform) quantum adversary that wins in $\mathcal{Q}$ with a non-negligible advantage.

Note that in the theorem above we did not mention whether or not $\mathcal{R}$ resets its adversary for any security parameter. The reason is that the adversary is assumed to be deterministic, and resetting their randomness does not change their behavior.

See Section 1.2 for an overview of the challenges and ideas for proving Theorem 2.

Proof of Theorem 2. Suppose $\mathcal{R}$ is a non-negligible-advantage reduction from winning in $\mathcal{Q}$ to winning in $\mathcal{P}$ that uses any non-negligible ε-advantage deterministic classical adversary in $\mathcal{P}$ to win in $\mathcal{Q}$ with a corresponding non-negligible advantage δ_ε. Let $\mathcal{A}_Q$ be any quantum algorithm who wins in $\mathcal{P}$ with advantage ε. Looking ahead, we will show how to emulate the reduction using $\mathcal{A}_Q$ and win in the game $\mathcal{Q}$ with advantage $\delta_{\Theta(\varepsilon)}$.

Case 1: $\mathcal{Q}$ is Threshold-Zero. When $\mathcal{Q}$ is threshold-zero and the reduction $\mathcal{R}$ only calls its adversary on a single security parameter $n_{\mathcal{P}}$, then roughly speaking, we can simply run the reduction normally without allowing its emulated adversary to be reset, and we can pretend its adversary oracle is deterministic. The reason is as follows. If $\mathcal{A}$ is an adversary who has advantage ε in $\mathcal{P}$ for the fixed security parameter $n_{\mathcal{P}}$, then by an averaging argument, it will have advantage at least $\frac{\varepsilon}{2}$ with probability at least $\frac{\varepsilon}{2}$ over its own randomness (when we sample and fix it). We can now construct an algorithm $\mathcal{B}$ as in Theorem 1 to run $\mathcal{R}$ and respond to its oracle queries using $\mathcal{A}'_Q$. $\mathcal{R}^{\mathcal{A}'_Q}$ then wins with probability at least $\frac{\varepsilon}{2} \cdot \delta_{\frac{\varepsilon}{2}}$, which is still non-negligible if ε is non-negligible. Since $\mathcal{Q}$ is a threshold-zero game, then $\mathcal{R}$'s advantage is equal to its success probability, so $\mathcal{R}$ has non-negligible advantage in $\mathcal{Q}$.

Case 2: $\mathcal{P}$ is Falsifiable. If $\mathcal{P}$ is falsifiable, then we can run the challenger of $\mathcal{P}$ efficiently. This means that we can efficiently determine whether a given fixed randomness of the adversary leads to a "good" adversary or not, and we can continue to guess a new fixed randomness until we arrive at a good one. In order to do so, it will be crucial here that we know ε (given to $\mathcal{R}$ as input) as stated in Definitions 5 and 9.

How to Find and Fix an Exponentially Large Random Tape? Here, we would like to fix the randomness of a potentially inefficient or quantum adversary, and we would like to do it *efficiently.* This seems like an oxymoron. First, recall that here we would like to gradually answer oracle queries while keeping them in a table that grows over time. For each of these partially generated tables, we have to "test" them (to conclude that they reflect a "good randomness to be fixed") and then use them against a reduction that anticipates a good and deterministic adversary. In doing so, we will not change the quantum nature of the emulation algorithm. We merely change the way the reduction uses the adversary in a way that the reduction can pretend its adversary is a classical successful deterministic adversary. We first describe the high level idea in the classical context of a random oracle, and then we will use the same idea while a quantum algorithm is emulating access to an inefficient randomized classical algorithm (just like the ROM).

Suppose we want to emulate a random oracle honestly, but we want to do it under a condition C about the ROM that can be probabilistically tested using an efficient tester. Also, suppose we eventually would like to run the test C in a final execution, but we want to first find (and completely fix the randomness of) a good random oracle that has a good chance of passing the test C in a final execution. In doing so, we are lazy-evaluating the random oracle, but we have to stick to the answers that we simulate before the final test C. If we could truly sample a complete random oracle and test it, this would be easy, as we could try several random oracles and pick the one that is good against C. However, even if we are doing a lazy evaluation efficiently, we can still stick to the partially generated oracle that we have sampled and extend it every time that we run the next test (even during the final test C). The key idea is that after k tests against C are passed, the probability of failing the next (final) test C (when we again continue to extend our partial oracle using lazy evaluation) goes down to $O(1/k)$. This trick works even though we are *not* actually fixing the whole randomness of the oracle, yet using this trick is *equivalent* to actually finding and fixing a randomness for the ROM.

Now, going back to running a quantum adversary, we have a similar challenge – one can pretend that the quantum adversary is also an oracle with exponentially long randomness (just by modeling the input-output behavior of the algorithm in a probabilistic way). However, we can use the same "lazy evaluation" trick by making sure that the same queries are answered similarly if asked twice, and also testing the partially generated oracle (table) many times before continuing to the actual final real test/reduction/game C. In both cases of ROM and quantum adversaries, we end up skewing the distribution of the sampled oracle/adversary by picking the good one, but it is fine as long as the adversary helps us succeed in what we want to do (i.e., winning the reduction's challenge).

We now use the above idea and describe an algorithm $\mathcal{B}'$ that uses $\mathcal{R}$ and $\mathcal{A}_Q$ to win against $\mathcal{Q}$ as described in Algorithm 3, where $k = \mathit{poly}(n_{\mathcal{Q}})$ for security parameter $n_{\mathcal{Q}}$, and $t_{\mathcal{P}}$ is the success threshold of $\mathcal{P}$. Intuitively, for every security parameter $n_{\mathcal{P}}$ for which there is an oracle query asked, we will first create polynomially many instantiations of $\mathcal{A}'_Q$, run them each against $\mathcal{P}$ polynomially many times (which we can do because $\mathcal{P}$ is falsifiable), and select an instantiation of $\mathcal{A}'_Q$ with good "fixed randomness".

We denote by $\mathcal{P}_n$ the game $\mathcal{P}$ with security parameter n, and let $\mathcal{N}_\mathcal{P}$ be the set of security parameters of $\mathcal{P}$ that the reduction $\mathcal{R}$ calls upon its own security parameter $n_\mathcal{Q}$.

Algorithm 3: Construction of $\mathcal{B}'$

```
Input: $k$, $\varepsilon$, $\mathcal{N}_\mathcal{P}$
for $n \in \mathcal{N}_\mathcal{P}$ do
    for $i \leftarrow 1$ to $k$ do
        Create a new $\mathcal{A}'^{(i,n)}_Q$ with an empty mapping for security parameter $n$
        Run $\mathcal{A}'^{(i,n)}_Q$ against $\mathcal{P}_n$ $k$ times, and let $\rho_i = j/k$ be the fraction of its wins
        (In case $\mathcal{A}_Q$ had non-uniform advice $|\phi_n\rangle$, use fresh copies of this advice in
          each of the $k$ invocations of the adversary $\mathcal{A}'^{(i,n)}_Q$ above or when it is reset.)
        if $\rho_i > t_\mathcal{P} + \frac{\varepsilon}{4}$ then
            Use the partial table of $\mathcal{A}'^{(n_\mathcal{P})}_Q = \mathcal{A}'^{(i,n)}_Q$ for security parameter $n$
            Break the loop and continue to the next security parameter
        end
    end
end
If no such adversary was found for any of $n \in \mathcal{N}_\mathcal{P}$, exit without success.
Run $\mathcal{R}$, and for security parameter $n_\mathcal{P}$ use the partially fixed oracle $\mathcal{A}'^{(n_\mathcal{P})}_Q$, and
  continue running it on new inputs by expanding its table.
```

For simplicity of the presentation, in what follows we will drop the explicit reference to the security parameter $n_\mathcal{P}$ in the instantiation of $\mathcal{A}'^{(n_\mathcal{P})}_Q$ and its statistically similar classical algorithm $(\tilde{\mathcal{A}}_Q)_C$. One can think of $(\tilde{\mathcal{A}}_Q)_C$ as a randomized adversary with a huge random tape (like the ROM). We say the randomness used by $(\tilde{\mathcal{A}}_Q)_C$ is "good" if by fixing that randomness, $(\tilde{\mathcal{A}}_Q)_C$ has advantage at least $\frac{\varepsilon}{2}$ to win against $\mathcal{P}$. By an averaging argument, $(\tilde{\mathcal{A}}_Q)_C$ picks a good randomness with probability $\geq \frac{\varepsilon}{2}$. By Lemma 3, $\mathcal{R}$'s behavior will be identical when we switch between giving it oracle access to $(\tilde{\mathcal{A}}_Q)_C$ such as $\mathcal{R}^{(\tilde{\mathcal{A}}_Q)_C}$ or emulating it as above using algorithm $\mathcal{B}'$.

In a similar vein, we can pretend that the simulated adversary $\mathcal{A}'_Q$ also uses (infinite) randomness r_i that describes the *full* (infinite) table of answers in its ith execution in Algorithm 3. We can also similarly define a randomness to be good if by fixing it, $\mathcal{A}'_Q$ becomes a deterministic adversary that wins against $\mathcal{P}$ with advantage at least $\frac{\varepsilon}{2}$.

We will first show that over k independent samplings of random seeds $r_1, \ldots, r_k$, at least one randomness will be good with high probability. We use ρ_i to denote the fraction of wins of $\mathcal{A}'^{(i)}_Q$ in $\mathcal{P}$ over k attempts. We next show that for any individual randomness r_i, if r_i is good, then $\rho_i > t_\mathcal{P} + \frac{\varepsilon}{4}$ with high probability, where $\frac{\varepsilon}{4}$ is chosen arbitrarily as a non-negligible value less than $\frac{\varepsilon}{2}$.

Claim. *At least one of the random seeds $r_1, \ldots, r_k$ (implicitly) used in Algorithm 3 to generate its k tables will be good with probability* $1 - \exp(-k\varepsilon^{O(1)})$.

Proof. The probability of any individual randomness being good is at least $\frac{\varepsilon}{2}$. Thus, the probability of any individual randomness being bad is upper bounded by $(1 - \frac{\varepsilon}{2})$. The probability of the randomness being bad all k times is upper bounded by $(1 - \frac{\varepsilon}{2})^k$. □

Claim. *If a randomness r_i is good, then $\rho_i > t_\mathcal{P} + \frac{\varepsilon}{4}$ holds with probability* $1 - \exp(-k\varepsilon^{O(1)})$.

Proof. If r_i is good, then $\mathcal{A}_Q'^{(i)}$ wins in $\mathcal{P}$ with advantage at least $\frac{\varepsilon}{2}$, so the expected value of ρ_i is $E[\rho_i] \geq t_{\mathcal{P}} + \frac{\varepsilon}{2}$, where $t_{\mathcal{P}}$ is the threshold of $\mathcal{P}$. Hoeffding's inequality states that, for k trials, it holds that

$$\Pr\left[\rho_i \leq t_{\mathcal{P}} + \frac{\varepsilon}{4}\right] \leq e^{-2k(\frac{\varepsilon}{4})^2}.$$

Therefore, if r_i is good, the probability that the condition $\rho_i > t_{\mathcal{P}} + \frac{\varepsilon}{4}$ is not satisfied is upper bounded by $e^{-2k(\frac{\varepsilon}{4})^2}$. □

From these claims, the probability that no good randomness is found is upper bounded by $1 - 2\exp(-k\varepsilon^{O(1)}) = 1 - \exp(-k\varepsilon^{O(1)})$. Therefore, with probability $1 - \exp(-k\varepsilon^{O(1)})$, it holds that $\mathcal{B}'$ finds an adversary $\mathcal{A}_Q'^{(i)}$ with "empirical" advantage at least $\frac{\varepsilon}{4}$ in its k executions against $\mathcal{P}$.

Another use of Hoeffding shows that the probability that an adversary has empirical advantage $\frac{\varepsilon}{4}$ in k executions, while the true advantage is less than $\frac{\varepsilon}{8}$, is at most $1 - \exp(-k\varepsilon^{O(1)})$. This bound holds for every security parameter $n_{\mathcal{P}}$ for which the oracle is asked a query. We now note that the maximum security parameter n_m for which the reduction $\mathcal{R}$ can ask a query is at most $n_m \leq \mathit{poly}(n_{\mathcal{Q}})$, since it is supposed to run in time $\mathit{poly}(n_{\mathcal{Q}})$, as $n_{\mathcal{Q}}$ is its own security parameter. It is also true that $\varepsilon(n_m) \geq 1/\mathit{poly}(n_m) \geq 1/\mathit{poly}(n_{\mathcal{Q}})$. Therefore, if we pick $k = \mathit{poly}(n_{\mathcal{Q}})$ large enough, with probability $1 - \exp(-k\varepsilon^{O(1)}) \leq 1 - \exp(-n_{\mathcal{Q}})$, it holds that $\mathcal{B}'$ finds and uses an adversary $\mathcal{A}_Q'^{(i)}$ with (actual) advantage at least $\frac{\varepsilon}{8}$ in $\mathcal{P}$ for *all* the security parameters that it queries. In that case, the reduction $\mathcal{R}$ will succeed against $\mathcal{Q}$ with probability at least $\delta_{(\frac{\varepsilon}{8})} - \exp(-n_{\mathcal{Q}})$, which is non-negligible whenever ε is non-negligible. This finishes the proof for the second case.

Note that the above algorithm and its analysis show that we can find "good" adversaries whenever the reduction $\mathcal{R}$ wants them, and this can be done for different security parameters, even if the reduction "resets" its adversary for a single security parameter multiple times. We now present a variant of this argument for the third case.

Case 3: $\mathcal{Q}$ is Falsifiable. If $\mathcal{Q}$ is falsifiable and $\mathcal{R}$ queries its adversary only on a single security parameter $n_{\mathcal{P}}$, we change (in fact simplify) our strategy of testing if an adversary is good in the challenge game $\mathcal{P}$, and instead we will merely test if the result of using it by the reduction leads to good enough success advantage in challenge $\mathcal{Q}$.

Algorithm 4: Construction of $\mathcal{B}''$

```
Input: $k$, $\varepsilon$, $\delta$
for $i \leftarrow 1$ to $k$ do
    Create a new $\mathcal{A}_Q'^{(i)}$ with an empty mapping
    Run $\mathcal{R}^{\mathcal{A}_Q'^{(i)}}$ against $\mathcal{Q}$ $k$ times, and let $\rho_i = j/k$ be the fraction of its wins
    (In case $\mathcal{A}_Q$ had non-uniform advice $|\phi_n\rangle$, use fresh copies of this advice in each
      of the $k$ invocations of the adversary $\mathcal{A}_Q'^{(i)}$ above or when it is reset.)
    if $\rho_i > t_{\mathcal{Q}} + (\delta_{\varepsilon/2})/2$ then
        Use $\mathcal{R}^{\mathcal{A}_Q'^{(i)}}$ for the final execution against $\mathcal{Q}$ by continuing to extend its table
        Terminate execution
    end
end
```

More formally, we construct a slightly different algorithm $\mathcal{B}''$ as described in Algorithm 4, where again $k = \mathit{poly}(n_{\mathcal{Q}})$, and $t_{\mathcal{Q}}$ is the threshold of $\mathcal{Q}$. For simplicity, we drop the fixed security parameters $n_{\mathcal{Q}}, n_{\mathcal{P}}$ from the description. Now, rather than observing whether any given $\mathcal{A}'^{(i)}_Q$ wins in $\mathcal{P}$ directly, we infer the effect of using it by the success of the reduction $\mathcal{R}^{\mathcal{A}'^{(i)}_Q}$ in the challenge $\mathcal{Q}$.

Just as in Case 1, the probability that no good randomness is chosen over k trials is at most $(1-\frac{\varepsilon}{2})^k$. Now, for any good randomness r_i, by definition, we will have to have $\mathbf{E}[\rho_i] \geq t_{\mathcal{Q}} + \delta_{\varepsilon/2}$. Therefore, as in Case 1, again by Hoeffding, when r_i is good, the probability that $\rho_i < t_{\mathcal{Q}} + \frac{\delta_{(\varepsilon/2)}}{2}$ in Algorithm 4 is upper bounded by $\exp(-k\gamma)$ for $\gamma = (\delta_{(\varepsilon/2)})^{O(1)}$. Finally, when Algorithm 4 runs its ith execution, the probability of obtaining $\rho_i \geq t_{\mathcal{Q}} + \frac{\delta_{(\varepsilon/2)}}{2}$, while the actual advantage is less than $\frac{\delta_{(\varepsilon/2)}}{4}$, is at most $\exp(-k\gamma)$ for $\gamma = (\delta_{(\varepsilon/2)})^{O(1)}$.

Since ε and consequently $\delta_{(\varepsilon/2)}$ are both non-negligible in $n_{\mathcal{Q}}$, all of the error probabilities above can be made at most $\exp(-n_{\mathcal{Q}})$ by picking $k = poly(n_{\mathcal{Q}})$ large enough. Therefore, the above algorithm $\mathcal{B}''$ wins its challenge with probability $(\delta_{(\varepsilon/2)}/4) - \exp(-n_{\mathcal{Q}}) \geq 1/poly(n_{\mathcal{Q}})$, and this finishes the proof. □

4 Deriving Minimal Post-Quantum Primitives

We now present the main conclusions of the paper. We use the machinery of the previous section to lift classical reductions between cryptographic primitives to the post-quantum setting. In Theorem 3, we show that if there is a classical reduction from a primitive $\mathcal{P}$ defined by a 2-message security game to another primitive $\mathcal{Q}$, then there is also a post-quantum reduction from $\mathcal{P}$ to $\mathcal{Q}$.

Theorem 3. *If there exists a classical reduction from a primitive $\mathcal{P}$ with a 2-message security game to a primitive $\mathcal{Q}$ defined using a multi-message security game, where the security reduction is black-box between their games, then the existence of the uniform (resp. non-uniform) post-quantum variant of $\mathcal{Q}$ (as defined in Definition 7) implies the existence of the uniform (resp. non-uniform) post-quantum variant of $\mathcal{P}$ (again as defined in Definition 7).*

Proof. Let $\mathcal{P}'$ and $\mathcal{Q}'$ be the post-quantum variants of $\mathcal{P}$ and $\mathcal{Q}$, respectively. The implementation of a post-quantum variant of a primitive is identical to the classical implementation of the primitive. Therefore, the implementation reduction property from Definition 8 holds, as switching from the classical setting to the post-quantum setting changes nothing.

To show that the security reduction property of Definition 8 holds, we must show that for any efficient uniform (resp. non-uniform) quantum algorithm $\mathcal{A}_{\mathcal{P}}$ that breaks the security game of $\mathcal{P}'$, then there is an efficient uniform (resp. non-uniform) quantum algorithm $\mathcal{A}_{\mathcal{Q}}$ that breaks the security game of $\mathcal{Q}'$. This property holds by Theorem 1. □

We stated the theorem above only for the main case, but for the case that the classical reduction is a deterministic-adversary reduction, we can use Theorem 2 to derive similar extensions as long as either of the three cases of Theorem 2 hold. In particular, we can use the special case that $\mathcal{P}$ is falsifiable to derive a stronger result for the minimality of OWFs.

Implications for One-Way Functions. We now connect Theorem 3 with the existence of one-way functions in post-quantum cryptography. We observe that with appropriate definitions, post-quantum one-way functions are minimal for post-quantum cryptography, just as classical one-way functions are minimal for classical cryptography.

Below we recall the definition of one-way functions for completeness with explicit mention of the type of the attacker (either classical or quantum). We call one-way functions with classical or quantum security simply classical one-way functions and post-quantum one-way functions, respectively.

Definition 10 (One-Way Function)**.** A function $f : \{0,1\}^* \to \{0,1\}^*$ is considered *one-way* with classical (resp. post-quantum) security if the following two conditions hold:

1. **Easy to Compute.** There exists a deterministic classical polynomial time algorithm A that computes f. That is, for all inputs x, $A(x) = f(x)$.
2. **Hard to Invert.** For all efficient classical (resp. uniform or non-uniform quantum polynomial-time) algorithms A', there is a negligible function λ, such that

$$\Pr_{x \leftarrow \{0,1\}^n, x' \leftarrow A'(f(x),1^n)}[f(x') = f(x)] \leq \lambda(n).$$

One-way functions trivially have a 2-message security game in which the challenger sends $y = f(x)$ for a (secretly) randomly chosen x to the adversary, and the adversary wins if its message x' maps to the same output $f(x)$. The challenger is also polynomial-time, so OWFs are falsifiable.

Finally, with this definition in hand, we can conclude that if the existence of some cryptographic primitive implies the existence of one-way functions classically through a black-box reduction between their security games, then the implication also holds in the post-quantum setting.

Corollary 1. *Let $\mathcal{Q}$ be a cryptographic primitive with a game-based security definition such that the classical existence of $\mathcal{Q}$ implies the existence of classical one-way functions through a black-box game-based reduction (but the implementation could be arbitrarily non-black-box). Then the existence of the post-quantum uniformly (resp. non-uniformly) secure variant of $\mathcal{Q}$ (as defined in Definition 7) implies the existence of uniformly (resp. non-uniformly) secure post-quantum one-way functions. Furthermore, this holds even if the reduction assumes its adversaries to be deterministic.*

To derive Corollary 1 from Theorem 3 we merely observe that one-way functions have a 2-message security game and that their security game is falsifiable. Therefore, Corollary 1 holds without restrictions on how the reduction from (post-quantum) OWFs to (post-quantum) $\mathcal{Q}$ calls its adversary.

One can use Theorem 3 to derive other similar results about the post-quantum minimality of other foundational primitives in classical cryptography, and this goes beyond one-way functions. For example, we could model the security of trapdoor one-way functions, key agreement, and semantically secure public-key encryption with 2-message games. Among these, trapdoor (one-way) functions (TDFs) play a central role in the Cryptomania world [Imp95], and our results demonstrate their minimality for post-quantum Cryptomania for every primitive that implies TDFs in a black-box way.

Many primitives, however, have security games with more than 2 messages (e.g., signatures, CCA-secure encryption, etc.). This raises the following question: do implications such as $\mathcal{Q} \to \mathcal{P}$ carry to the post-quantum setting when $\mathcal{P} = \mathcal{S}$ is, for example, the primitive of signatures? The particular question about signatures can be answered using our results indirectly, due to the following reasons: (1) signatures and OWFs are equivalent in the classical setting through a fully black-box reduction, and (2) post-quantum OWFs imply post-quantum signatures as well. So, $\mathcal{Q} \to \mathcal{S}$ implies $\mathcal{Q} \to$ OWFs, and therefore the post-quantum $\mathcal{Q}$ implies post-quantum OWFs, which in turn implies post-quantum signatures. For other complicated primitives, it remains open to study their minimality in the post-quantum world (for the appropriate class of primitives).

References

[AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009. doi:10.1017/CBO9780511804090.

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 208–236, Cham, 2022. Springer Nature Switzerland. doi:10.48550/arXiv.2112.10020.

[ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS '14, page 474–483, USA, 2014. IEEE Computer Society. doi:10.1109/FOCS.2014.57.

[Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 106–115. IEEE, 2001. doi:10.1109/SFCS.2001.959885.

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, page 175–179, December 1984. doi:10.48550/arXiv.2003.06557.

[BBF13] Paul Baecher, Christina Brzuska, and Marc Fischlin. Notions of black-box reductions, revisited. In *Advances in Cryptology-ASIACRYPT 2013: 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I 19*, pages 296–315. Springer, 2013. doi:10.1007/978-3-642-42033-7_16.

[BBK22] Nir Bitansky, Zvika Brakerski, and Yael Tauman Kalai. Constructive post-quantum reductions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 654–683, Cham, 2022. Springer Nature Switzerland. doi:10.48550/arXiv.2203.02314.

[BCM$^+$21] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. *J. ACM*, 68(5), August 2021. doi:10.1145/3441309.

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the Computational Hardness Needed for Quantum Cryptography. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:21, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2023.24, doi:10.4230/LIPIcs.ITCS.2023.24.

[BJY97] Mihir Bellare, Markus Jakobsson, and Moti Yung. Round-optimal zero-knowledge arguments based on any one-way function. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 280–305. Springer, 1997. doi:10.1007/3-540-69053-0_20.

[BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 361–379. Springer, Berlin, Heidelberg, August 2013. doi:10.1007/978-3-642-40084-1_21.

[CFP22] Benjamin Chan, Cody Freitag, and Rafael Pass. Universal reductions: Reductions relative to stateful oracles. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 151–180, Cham, 2022. Springer Nature Switzerland. doi:10.1007/978-3-031-22368-6_6.

[CLMP12] Kai-Min Chung, Edward Lui, Mohammad Mahmoody, and Rafael Pass. Unprovable security of 2-message zero knowledge. Cryptology ePrint Archive, Paper 2012/711, 2012. URL: https://eprint.iacr.org/2012/711.

[CLMP13] Kai-Min Chung, Huijia Lin, Mohammad Mahmoody, and Rafael Pass. On the power of nonuniformity in proofs of security. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 389–400, 2013. doi:10.1145/2422436.2422480.

[FFS87] Uriel Fiege, Amos Fiat, and Adi Shamir. Zero knowledge proofs of identity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 210–217, 1987. doi:10.1145/28395.28419.

[Gol06] Oded Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, USA, 2006. doi:10.1017/CBO9780511546891.

[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 99–108, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993636.1993651.

[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. arXiv:https://doi.org/10.1137/S0097539793244708, doi:10.1137/S0097539793244708.

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, Mar 1963. doi:10.1080/01621459.1963.10500830.

[HY20] Akinori Hosoyamada and Takashi Yamakawa. Finding collisions in a quantum world: Quantum black-box separation of collision-resistance and one-wayness. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 3–32, Cham, 2020. Springer International Publishing. doi:10.1007/s00145-024-09517-2.

[IL89] R. Impagliazzo and M. Luby. One-way functions are essential for complexity based cryptography. In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, 1989. doi:10.1109/SFCS.1989.63483.

[ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, page 12–24, New York, NY, USA, 1989. Association for Computing Machinery. doi:10.1145/73007.73009.

[Imp95] R. Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th Annual Structure in Complexity Theory Conference (SCT'95)*, SCT '95, pages 134–147, USA, 1995. IEEE Computer Society. doi:10.1109/SCT.1995.514853.

[IMS12] Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. On efficient zero-knowledge pcps. In *Theory of Cryptography: 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings 9*, pages 151–168. Springer, 2012. doi:10.1007/978-3-642-28914-9_9.

[IR90] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *Advances in Cryptology — CRYPTO' 88*, pages 8–26, New York, NY, 1990. Springer New York. doi:10.1145/73007.73012.

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 126–152, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-96878-0_5.

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1589–1602, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585225.

[Kre21] William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2021.2, doi:10.4230/LIPIcs.TQC.2021.2.

[LM22] Alex Lombardi and Fermi Ma. Quantum rewinding tutorial, part 1: Motivation and early rewinding techniques. In *Quantum and Lattices Joint Reunion Workshop, Simons Institute, Berekly*, 2022. URL: https://simons.berkeley.edu/talks/quantum-rewinding-tutorial-part-1-motivation-early-rewinding-techniques.

[LMQW22] Alex Lombardi, Ethan Mook, Willy Quach, and Daniel Wichs. Post-quantum insecurity from lwe. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 3–32, Cham, 2022. Springer Nature Switzerland. doi:10.1007/978-3-031-22318-1_1.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 269–295, Cham, 2022. Springer Nature Switzerland. doi:10.1007/978-3-031-15802-5_10.

[Nao03] Moni Naor. On cryptographic assumptions and challenges. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, pages 96–109, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. doi:10.1007/978-3-540-45146-4_6.

[NY89] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *21st ACM STOC*, pages 33–43. ACM Press, May 1989. doi:10.1145/73007.73011.

[Pas11] Rafael Pass. Limits of provable security from standard assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 109–118, 2011. doi:10.1145/1993636.1993652.

[Rom90] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, STOC '90, page 387–394, New York, NY, USA, 1990. Association for Computing Machinery. doi:10.1145/100216.100269.

[RTV04] Omer Reingold, Luca Trevisan, and Salil Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *Theory of Cryptography*, pages 1–20, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. doi:10.1007/978-3-540-24638-1_1.

[Sha20] Ronen Shaltiel. Is it possible to improve yao's xor lemma using reductions that exploit the efficiency of their oracle? In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. doi:10.1007/s00037-023-00238-9.

[Sho94] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994. doi:10.1109/SFCS.1994.365700.

[Son14] Fang Song. A note on quantum security for post-quantum cryptography. In Michele Mosca, editor, *Post-Quantum Cryptography*, pages 246–265, Cham, 2014. Springer International Publishing. doi:10.48550/arXiv.1409.2187.

[Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, pages 135–152, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-29011-4_10.

[Wat06] John Watrous. Zero-knowledge against quantum attacks. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 296–305, New York, NY, USA, 2006. Association for Computing Machinery. doi:10.1145/1132516.1132560.

[Wat08] John Watrous. Quantum computational complexity, 2008. URL: https://arxiv.org/abs/0804.3401, arXiv:0804.3401.

[Wic13] Daniel Wichs. Barriers in cryptography with weak, correlated and leaky sources. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, page 111–126, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2422436.2422451.

[WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, Oct 1982. doi:10.1038/299802a0.

[YZ20] Takashi Yamakawa and Mark Zhandry. A note on separating classical and quantum random oracles. *Cryptology ePrint Archive*, 2020.