

Constructive Post-Quantum Reductions*

Nir Bitansky¹, Zvika Brakerski², and Yael Tauman Kalai³

¹ Tel Aviv University, Israel. nirbitan@tau.ac.il

² Weizmann Institute of Science, Israel. zvika.brakerski@weizmann.ac.il

³ Microsoft Research and Massachusetts Institute of Technology, MA, USA.
yael@microsoft.com

Abstract. Is it possible to convert classical reductions into post-quantum ones? It is customary to argue that while this is problematic in the interactive setting, non-interactive reductions do carry over. However, when considering quantum auxiliary input, this conversion results in a *non-constructive* post-quantum reduction that requires duplicating the quantum auxiliary input, which is in general inefficient or even impossible. This violates the win-win premise of provable cryptography: an attack against a cryptographic primitive should lead to an algorithmic advantage.

We initiate the study of constructive quantum reductions and present positive and negative results for converting large classes of classical reductions to the post-quantum setting in a constructive manner. We show that any non-interactive non-adaptive reduction from assumptions with a polynomial solution space (such as decision assumptions) can be made post-quantum constructive. In contrast, assumptions with super-polynomial solution space (such as general search assumptions) cannot be generally converted.

Along the way, we make several additional contributions:

1. We put forth a framework for reductions (or general interaction) with *stateful* solvers for a computational problem, that may change their internal state between consecutive calls. We show that such solvers can still be utilized. This framework and our results are meaningful even in the classical setting.
2. A consequence of our negative result is that quantum auxiliary input that is useful against a problem with a super-polynomial solution space cannot be generically “restored” post-measurement. This shows that the novel rewinding technique of Chiesa et al. (FOCS 2021) is tight in the sense that it cannot be extended beyond a polynomial measurement space.

1 Introduction

The notion of provable security in cryptography has had a great impact on the field and has become a de-facto gold standard in evaluating the security of cryptographic primitives. A provably secure cryptographic primitive is stated in

*A full version of this work is available [3].

the form of a computational problem P , whose hardness is related by means of *reduction* to that of another problem Q which is either by itself considered intractable or in turn can be further reduced down the line. The reduction is an algorithm that solves the problem Q provided that it is given access to an algorithm that solves the problem P .

This gives rise to the “win-win principle” which stands as one of the main motivations for using provably secure cryptography. The logic is the following. Either an algorithmic solution for P cannot be found, i.e. the cryptographic primitive P is secure for all intents and purposes, or one can find an algorithmic solution for P which would imply an algorithmic solution for Q , thus contributing to the state of the art in algorithms design. Indeed, cryptographic reductions are the main working tool for the theoretical cryptographer. Numerous reductions between cryptographic primitives are known and hundreds of such reductions are published in the cryptographic literature every year.

The emergence of the quantum era in computing poses a new challenge to provable security and the win-win principle. Many existing reductions in the “pre-quantum” world implicitly or explicitly relied on the P -algorithm being classical. These reductions are thus a-priori invalid when considering quantum algorithms. A central line of investigation in the domain of post-quantum security is thus dedicated to the following question.

To what extent can pre-quantum reductions be ported to the post-quantum setting?

Such conversion may not always be possible. This is particularly a concern when considering *interactive* problems, i.e. ones where the solution to P involves multiple messages being exchanged with the solver algorithm. Indeed, one of the most prominent techniques for proving security in the interactive setting, namely the notion of *rewinding*, does not directly translate to the quantum setting and moreover one can explicitly show cases where pre-quantum reductions exist but post-quantum ones do not. In fact, this property was actually used to construct *proofs of computational quantumness* [5] in which a party proves that it is quantum by succeeding in a task for which there is a classical impossibility result (under computational assumptions). In a nutshell, the reason is that a quantum algorithm may keep a quantum state between rounds of interaction, and this quantum state is measured and thus potentially destroyed in order to produce the next message of interaction. It is therefore not possible to naively “rewind” the interaction back to a previous step as is customary in many classical proofs.

The focus of this work, therefore, is on *non-interactive cryptographic assumptions*. These are problems P whose syntax contains a (randomized) instance generator which generates some instance x , and a verifier that checks whether solutions y are valid (with respect to x or more generally the randomness that was used to generate x). The role of the solver algorithm in this case is simply to take x as input and produce a y that “verifies well” (we avoid getting into the exact formalism at this point).

Contrary to the interactive case, it is customary to postulate (often without proof) that classical reductions to non-interactive cryptographic assump-

tions carry over straightforwardly to the post-quantum setting since there is no rewinding. There is a simple challenge-response interface that on the face of it “does not care” whether the underlying P -solver is implemented classically or quantumly. This viewpoint, however, is overly simplistic, since the P solver may use *quantum auxiliary input*: a quantum state $|s\rangle$ that is used as a resource for solving P . The state $|s\rangle$ can be the result of some natural process upon which we have no control, or a result of some exhaustive preprocessing, or generated in the course of execution of some protocol. At any rate, the means to produce $|s\rangle$ are often not at our disposal, we just get a copy of the state.

In this case, similarly to the interactive setting, the quantum state is measured whenever the P -solver is called, and therefore, it potentially precludes us from calling the P solver more than once. This issue is often addressed in the literature by noticing that providing many copies of $|s\rangle$ would allow to call the P solver multiple times – namely there exists a quantum state $|s\rangle^{\otimes t}$ that allows to solve Q given access to the P solver. Therefore, the existence of a classical reduction still implies that if Q is intractable *even given arbitrary auxiliary input*, then the same holds for P .

We argue that the aforementioned common “solution” for post-quantum reductions in the presence of quantum auxiliary input is unsatisfactory. First and foremost, this solution violates the win-win principle. While the argument above indeed implies that (some form of) intractability for P follows from (some form of) intractability for Q , it *does not* allow to convert an auxiliary-input algorithm for P into an auxiliary-input algorithm for Q in a constructive manner, since the transformation $|s\rangle \rightarrow |s\rangle^{\otimes t}$ is not an efficient one. An additional related concern is the *durability* of such reductions. Namely, that if we wish to execute the reduction more than once (i.e. solve multiple instances of Q) then we need to duplicate the state $|s\rangle$ an a-priori unbounded number of times.

Given this state of affairs, the question we are facing is the following.

To what extent can pre-quantum reductions to non-interactive assumptions be ported to the post-quantum setting constructively and durably?

Naturally, we do not wish to redo decades of cryptographic work in re-proving each result individually. Instead, we would like to identify the broadest class of pre-quantum reductions that can be generically converted into the post-quantum regime, and at the same time characterize the limitations where such generic conversion is not possible. This is the focus of this work, and indeed we show a generic transformation for a very broad class of reductions. Along the way we develop an adversarial model for *stateful* adversaries that may be of interest in its own right, even in the *classical* setting.

1.1 Our Main Results

We prove a general positive result for converting classical reductions into post-quantum ones. In particular we consider *non-adaptive reductions*. In such reductions, the set of queries to the oracle is determined before any query is made. It

turns out that an important parameter in our positive as well as our negative result is the size of the *solution space* of the computational problem P (“the cryptographic primitive”). Our positive results apply to cases where the solution space is polynomial.⁴ One notable example the case where P is a “decision assumption”, namely the P solver is a distinguisher that returns a single bit as output. Another notable example is the case where P is an NP search problem, with unique solutions (e.g., injective one-way functions or unique signatures). An informal result statement follows.

Theorem 1.1 (Positive result, informal). *There exists an efficient transformation for converting any classical non-adaptive black-box reduction from assumption Q to assumption P , where P is a non-interactive assumption with a polynomial solution space, into a constructive and durable post-quantum reduction from Q to P .*

We prove a complementary negative result, for the case where P has a large solution space. The negative result relies on the existence of classical indistinguishability obfuscation which is secure against quantum adversaries.

Theorem 1.2 (Negative result, informal). *Assume the existence of post-quantum secure indistinguishability obfuscation. Then there exist non-interactive assumptions P, Q , where P has a super-polynomial solution space and the following hold. There exists a classical non-adaptive black-box reduction from assumption Q to assumption P , but there is no such constructive post-quantum reduction.*

As explained above, in order to address the question of constructiveness, we need to develop a new adversarial model and a host of tools to address this question. An account of these intermediate contributions appears in the technical overview below.

1.2 Our Techniques and Additional Contributions

Known approaches fall short of achieving constructiveness and durability since they regard quantum auxiliary input similarly to its classical counterpart, despite the inherent difference of the inability to duplicate or reuse quantum information. We assert that the process of making multiple calls to an algorithm with quantum side information is inherently *stateful*. Namely, the internal state of the “oracle” changes and evolves over time. In this work we put forth a framework for stateful solvers, namely algorithms that change their internal state and thus their behavior over time.

In the post-quantum setting, reductions start from *one-shot* solvers. That is, ones that have an initial state that allows them to provide an answer for a single

⁴They in fact apply even if the solution space is polynomial *per instance* (but the space is not the same for all instances), and a certain natural verifiability property is satisfied (see Definition 3.2).

instance of P successfully, but afterwards all bets are off. It seems natural (and, as we show, turns out to be useful) to consider stateful solvers that propagate their P -solving property throughout an execution, we call this property *persistence*. Persistent solvers evolve their state in an arbitrary way subject to being able, at any point in their evolution, to successfully answer a P -query (with some noticeable advantage).

A Framework for Stateful Solvers. Section 3 is dedicated to formally defining the notion of a (potentially stateful) solver and quantifying its success probability in solving a problem P . We accordingly provide definitions for a post-quantum reduction in this setting, and more specifically the notion of a post-quantum black-box reduction. The standard notion of a classical black-box reduction is recovered as a special case of our definition, when specializing to so-called *stateless* P -solvers.

Using our new formalism, the task at hand is to convert a reduction that expects to be interacting with a *stateless* solver, into one that is successful even when given a *one-shot stateful* solver.

One-Shot Solvers Imply Persistent Solvers. One-shot solvers may seem quite useless, since on the face of it they may only successfully respond to a single query. However, our first technical result, in Section 4, is that they can in fact be converted generically (but in a non-black-box manner) into persistent solvers. Namely, ones that can answer an *a-priori unbounded* number of queries and maintain roughly the same success probability. The persistent solver has a state of length that is polynomially related to that of the one-shot solver. The running time of the persistent solver increases with each query it is being asked. That is, the time complexity of answering the t -th query scales with $\text{poly}(t)$ for a fixed polynomial. This still ensures that for any polynomial-length sequence of queries, the total time to answer all queries is bounded by a fixed polynomial. The persistent value of the resulting solver (i.e. the value that is maintained for an a-priori unbounded number of times) is itself a random variable that is determined during the conversion process. The expectation of the persistent value is equal to the one-shot value of the solver we start from. (We note that it is inherently impossible to achieve a non-probabilistic behavior, i.e. to ensure a persistent value that is always above some threshold.⁵)

Our transformation is an extension of the techniques in the recent work of Chiesa, Ma, Spooner and Zhandry [7], that can be interpreted as showing such a transformation for “public-coin” cryptographic assumptions (ones where the instances are uniformly distributed and the verification requires only the in-

⁵To see this, consider the case where the one-shot auxiliary input $|s\rangle$ is a superposition giving weight $\sqrt{1-\varepsilon}$ to a value $|\perp\rangle$ that always makes the P -solver fail, and giving weight $\sqrt{\varepsilon}$ to a state that makes the P -solver perfectly successful. Then, by trace-distance considerations, any processing of $|s\rangle$ must be ε -statistically-indistinguishable from a case where $|s\rangle = |\perp\rangle$. Therefore, with probability at least $1-\varepsilon$ the persistent value will be trivial. Nevertheless, using a Markov argument, if we start from a one-shot solver with a non-negligible advantage, we recover, with a non-negligible probability, a many-shot solver with a non-negligible persistent advantage.

stance and the solution, and not the randomness that was used to generate the instance). It is only in this step that we have the restriction that the solution space of the problem needs to be polynomial, due to limitations of the [7] technique. Our negative result (further discussed below) proves that these limitations are inherent.

The conversion from one-shot to persistent is the only transformation that uses the solver in a non-black-box manner. In the rest of our (positive) results we take a persistent P -solver and a bound on the length of its auxiliary quantum state and only make black-box use of this solver, i.e., provide instances as input and receive solutions as output. We do not further intervene with the evolution of the state between consecutive calls to the solver.

Once we transformed our solver to being persistent, we are guaranteed that we can make multiple P queries, and each one will be answered by a “successful” solver. It may seem that our mission is complete. However, this is far from being the case. While all queries are answered by a successful solver, these solvers may be arbitrarily correlated. For example, thinking about a simple linearity test where a reduction queries $x_1, x_2, x_3 = x_1 \oplus x_2$ and checks whether a linear relation holds. It may be the case that for each query x_i we get a response y_i from an approximately-linear function, and yet the solver “remembers” that x_1, x_2 were previously made as queries, and deliberately fails on $x_1 \oplus x_2$ in the next query. Another example, that will be quite useful to illustrate our transformation is that of the Goldreich-Levin (GL) hardcore bit [12], where queries take the form $(f(x), r_i)$, always with the same $f(x)$, and with additional correlations between the r_i values across different queries. In particular, it may be the case that once a query with some value $f(x)$ has been made, the solver refuses to meaningfully answer any additional queries with the same $f(x)$.⁶

We note that attributing adversarial behavior to the solver is done for purposes of analysis. Our transformation from one-shot to persistent appears quite “innocent” and we do not know whether it can actually generate such pathological behavior that will prevent reductions from running. However, we cannot rule it out and therefore we consider a worst-case adversarial model.

When described in this way, it seems that only very specialized reductions can be carried over to the post quantum setting. For example, ones that employ a strong form of random self reduction when making solver queries. One such case is the search-to-decision reduction for the learning with errors problem [20]. However, as the GL hardcore bit example demonstrates, this doesn’t even extend to all search to decision reductions. We must therefore find a new way to utilize stateful solvers. Indeed, the handle that we use is that while the solver may change its behavior adversarially, its adversarial behavior is constrained by the length of the auxiliary state $|s\rangle$ that it uses. We will indeed leverage the fact

⁶We note that while the *classical* GL reduction, falls under our umbrella of non-adaptive reductions, in this specific case, it is in fact known how to devise a single-query *quantum* reduction [1]. This, however, does not resolve the question of durability, and more importantly does not provide a general framework for all non-adaptive reductions.

that this state is polynomially bounded to limit the adversarial powers of the solver and handle more general reductions.

Before moving on to describe our techniques in this context, we notice that while this adversarial model (of black-box access to a persistent solver) emerged as a by-product of our work on quantum reductions, it is nevertheless a valid model in its own right in both the quantum and classical setting. We may consider interacting with an adversary/solver that is *only* guaranteed to be noticeably successful at every point in time but, unlike the standard notion of an “oracle”, may change its behavior over time. In our case, we allow the behavior to change arbitrarily, so long that the amount of information carried over between executions is bounded (in our case, by the length of the state, which is polynomially bounded).

Memoryless Persistent Solvers. Our next step, in Section 5, is to show that a persistent solver, even with adversarial behavior, can be effectively converted into a more predictable form of solver that we call *memoryless* (note that this is different from our final goal which is to achieve a *stateless* solver). A memoryless solver keeps track of the sequence number of the question it is asked (e.g. it knows that it is now answering query number 4) but it is not allowed to remember any information about the actual content of the previous queries that were made.

We show that a combination of a non-adaptive reduction and a persistent solver induce a memoryless (persistent) solver (more accurately a distribution over memoryless solvers). These memoryless solvers are accessible using a simulator that, given access to the reduction and the original solver, efficiently simulates the interaction of the reduction with the induced memoryless solver, up to inverse-polynomial statistical distance. Note that we require that the reduction is non-adaptive. Namely, its queries to the solver can be arbitrarily correlated (as in the GL case), but the identity of the queries must not depend on the answers to previous queries.

The transformation relies on the fact that the solver has a bounded amount of memory, say ℓ qubit of state that is propagated through the execution. Our strategy is to dazzle the solver with an abundance of i.i.d dummy queries, that are sampled from the marginal distribution of the “real” queries (for example, in the GL case, each dummy query will have the form $(f(x_i), r_i)$ where x_i, r_i are both random). In between the dummy queries, in random locations, we plant our real queries, in random order. We prove that the solver, having only ℓ qubits of state, must answer our real queries as if they were dummy queries. This requires us to develop a proper formalism and to prove a new lemma (Plug-In Lemma) using tools from quantum information theory. See Section ?? for the full details.

Stateless Solvers at Last. Finally, we show in Section 6 that memoryless solvers imply stateless solvers. This is again shown by means of simulation via a similar formalism to the previous result. Recall that a stateless solver must answer all queries according to the same distribution. This transformation again relies on the non-adaptive nature of the reduction, namely on the ability to generate all solver-queries ahead of time. To do this, we notice that we can think of a memoryless solver simply as a sequence of stateless solvers that can be

queried one at a time. Therefore, we can consider the induced stateless solver that at every query picks a random solver from this collection and executes it on the query. This indeed will result in a stateless solver. The solving probability of the induced stateless solver is simply the average success probability of solvers in the collection, which is concentrated due to persistence. Moreover, this behavior can be simulated by randomly permuting the *queries*, while still calling the solvers according to their order in the sequence.⁷

This way, asking the queries in a permuted order to the memoryless solver will (almost) mimic the action of sampling a solver from the collection independently for each query. The only reason why this mimic is not perfect is that permuted queries are sampling “without repetition”, i.e. none of the solvers in the sequence defined by the memoryless solver will be queried twice, whereas in the ideal strategy we described above, it is possible that the same solver from the sequence will be sampled more than once. We deal with this by making the number of solvers in the sequence so big, that the probability of hitting the same solver twice becomes very small (inversely polynomial for a polynomial of our choice). We simply add to our queries of interest a large number of dummy “0 queries”, and perform a random permutation on this extended set of queries.

Putting Things Together. In Section 7 we put all of the components together and prove our main positive result, that any classical non-adaptive reduction which relies on a non-interactive polynomial-solution-space assumption can be made post quantum. This requires putting together the components in a careful manner.

The fact that the first step in our transformation was to produce a persistent P , allows us to continue using it even after having solved a Q instance. This means that we can solve additional instances of Q , or use it to solve additional instances of P or any other problem Q' for which a non-adaptive reduction to P exists. In particular, this property implies that our reduction is *durable*.

A Negative Result for Search Assumptions. We show in Section 8 that a generic conversion from classical to constructive quantum reductions is not always possible, even for the case of non-adaptive reductions to non-interactive assumptions. In particular, if P is an assumption with a large solution space (intuitively, a search assumption) this may not be possible.

We show our negative result by relying on a recently introduced primitive known as tokenized signatures [2]. These are signature schemes with the standard classical syntax, but for which it is possible to produce a quantum *signature token*. The signature token allows to generate a single classical signature for a message of the signer’s choice, but only one such signature can be created. Tokenized signatures have been constructed relative to a classical oracle [2] or based on cryptographic assumptions [8].

⁷Remember that we have access to the memoryless solver which only allows to make queries in order.

We can define an “assumption” which is essentially the task of signing a random message using a tokenized signature scheme.⁸ In the classical world, there is a trivial reduction between the task of signing one random message and the task of signing two random messages. However, if we consider a quantum solver that holds the token as auxiliary input, then by definition it should not be possible to use it to obtain two signatures for two different messages. Our negative result holds for any conversion process that is constructive, and in particular does not obtain any implicit non-uniform advice about the assumption.

1.3 Other Related Work

The question of which reductions can be translated from the classical to the post-quantum setting also received significant attention in the context of the random-oracle model (ROM), starting from the work of Boneh et al. [4]. The question asked in these works is whether it is possible to convert reductions in the classical ROM into ones in the quantum ROM (QROM, where the adversary is allowed to make quantum queries to the oracle). There are several results proving that specific schemes that are secure in the ROM are also secure in the QROM [23, 21, 14, 15, 24, 11, 17, 10, 16]. Recently, a more general “lifting theorem” was given in [22], showing how to convert a proof in the ROM to one in the QROM for any “search-type game” where a challenger makes only a *constant* number of queries to the random oracle. This work also presented a negative result, showing that there are schemes that are secure in the ROM yet are insecure in the QROM. While the general motivation in these works is similar to ours, the question they ask is quite different from ours. In the ROM/QROM, the solver is allowed to make queries to the oracle (which is simulated by the reduction), which is more similar to the setting where interactive-assumptions are used.

Our memoryless transformation (Section 5) relies heavily on the state of the solver being bounded in length. The idea that bounded quantum memory can be used to restrict an otherwise all powerful adversary is at the core of the bounded quantum storage model. It can be shown (see, e.g., [9]) that it is possible to achieve cryptographic abilities against strong adversaries while relying only on a limit on the amount of quantum storage they can use. This setting is quite different from ours, though, since the quantum bounded storage model allows an unbounded amount of classical memory, which in our setting would make it impossible to achieve any result. Indeed, the bounded storage model requires quantum communication (whereas our reduction-solver communication is completely classical), and thus the set of tools and techniques that are used in both settings are completely different.

⁸The assumption is instantiated by a verification key which we can think of as non-uniformity of the assumption, see discussion in Section 8.

2 Preliminaries and Tools

We say that a given function $f(x_1, \dots, x_k)$ is $\text{poly}(x_1, \dots, x_k)$, if there exist constants c, C such that $(x_1 \cdot x_2 \cdot \dots \cdot x_k)^c \leq f \leq (x_1 \cdot x_2 \cdot \dots \cdot x_k)^C$.

We denote by TD the trace distance between two matrices.

Algorithms. By default, when referring to an *algorithm* we mean a classical probabilistic (resp. quantum) algorithm. Algorithms may be uniform or non-uniform, meaning that they have *classical* advice related to the input size (we specify when uniformity matters). An *efficient* algorithm is also polynomial time.

Quantum Notation. We use standard quantum information in Dirac notation. We denote quantum variables in boldface \mathbf{x} and classical variables in lowercase x . The density matrix of \mathbf{x} is denoted $\rho_{\mathbf{x}}$. Classical variables may also have (diagonal) density matrices. Quantum variables \mathbf{x}, \mathbf{y} have a joint density matrix $\rho_{\mathbf{x}, \mathbf{y}}$ if they can be jointly produced by an experiment. As usual, \mathbf{x}, \mathbf{y} are independent if $\rho_{\mathbf{x}, \mathbf{y}} = \rho_{\mathbf{x}} \otimes \rho_{\mathbf{y}}$. We never assume that quantum variables are independent unless we explicitly say so. Quantum registers are denoted in capital letters. We also sometimes use capital letters to denote distributions, where it is clear from the context. For a finite Hilbert space \mathcal{H} we denote by $\mathbf{S}(\mathcal{H})$ the set of density matrices over quantum states in \mathcal{H} .

A quantum procedure is a general quantum algorithm that can apply unitaries, append ancilla registers in 0 state, perform measurements in the computational basis and trace out registers. The complexity of F is the number of *local* operations it performs (say, operations on up to 3 qubits are considered local). If F is a quantum procedure then we denote by $F(\mathbf{x})$ the application of F on \mathbf{x} . Any unitary induces a quantum procedure that implements this unitary, which does not perform measurements or trace out registers, we call this procedure “a unitary quantum circuit”.

Purification of Quantum Procedures and States. A quantum procedure may introduce new ancilla qubits, perform intermediate measurement throughout its computation and discard registers or parts thereof. However, any quantum procedure can be *purified* into unitary form without much loss in complexity [19]. This is formally stated below.

Proposition 2.1. *Let C be a general quantum procedure of complexity s . Then it is possible to efficiently generate a unitary quantum circuit \widehat{C} of size $O(s)$, such that for any quantum state (\mathbf{x}, \mathbf{a}) , setting $(\mathbf{y}, \mathbf{z}) = \widehat{C}(\mathbf{x}, \mathbf{0})$, it holds that (\mathbf{y}, \mathbf{a}) has identical density matrix to $(C(\mathbf{x}), \mathbf{a})$.*

Likewise, any quantum state can be viewed as a reduced density matrix of the output of a unitary (which may be inefficient to implement).

Proposition 2.2. *Let \mathbf{x} be a variable with density matrix $\rho_{\mathbf{x}}$. Then there exists a unitary U over registers XY such that applying $U(\mathbf{0}, \mathbf{0})$, the reduced density matrix of the value in the X register has density matrix $\rho_{\mathbf{x}}$.*

2.1 The Plug-In Lemma

The following lemma is another manifestation of information incompressibility in the quantum setting. Specifically, we are interested in an experiment in which an all powerful compressing procedure attempts to compress t samples which are arbitrarily distributed into ℓ quantum bits. We show that this is infeasible even in the weak sense in which a decoder receives the compressed value, and a $(j - 1)$ -prefix of the sequence, and is required to identify the j -th element. We show that as t increases, the probability of succeeding in the experiment drops. A formal statement follows.

Lemma 2.3 (Plug-In Lemma). *Let $\vec{Y} = (Y_1, \dots, Y_t)$ be a joint distribution over t classical random variables. Let \vec{y} be distributed according to \vec{Y} . Let \mathbf{s} be an ℓ -qubit random variable that has arbitrary dependence on \vec{y} . We let \vec{y}_i denote the prefix $\vec{y}_i = (y_1, \dots, y_i)$ for $1 \leq i \leq t$, and \vec{y}_0 is the empty vector (and likewise for \vec{Y}). Let J be the uniform distribution over $[t]$ and let $j \leftarrow J$. Define $y' \leftarrow Y_j | (\vec{Y}_{j-1} = \vec{y}_{j-1})$. Then it holds that*

$$\text{TD}((j, \vec{y}_{j-1}, y_j, \mathbf{s}), (j, \vec{y}_{j-1}, y', \mathbf{s})) \leq \sqrt{\ell/(2t)}. \quad (1)$$

Note that the above two distributions are *not* identical even though (j, \vec{y}_{j-1}, y_j) and (j, \vec{y}_{j-1}, y') are identically distributed. The reason is that in both cases, \mathbf{s} is always generated as a function of \vec{y} , i.e. using y_j and not y'_j .

The lemma is proven in the full version of this paper [3].

3 Assumptions, Stateful Solvers, and Reductions

In this section, we formally define the concepts of non-interactive cryptographic assumptions, stateful solvers, and their value and advantage in breaking an assumption.

3.1 Non-Interactive Assumptions

We define the notion of a non-interactive (falsifiable) cryptographic assumption as in [18, 13]. While we frame the notion as “cryptographic”, it can be viewed more generally as a notion for average-case problems where the solution can be verified.

Definition 3.1 (Non-Interactive Assumption). *A non-interactive assumption is associated with polynomials $d(\lambda), n(\lambda), m(\lambda)$ and a tuple $P = (G, V, c)$ with the following syntax. The generator G takes as input 1^λ and $r \in \{0, 1\}^d$, it returns $x \in \{0, 1\}^n$. The verifier V takes as input 1^λ and $(r, y) \in \{0, 1\}^d \times \{0, 1\}^m$ and returns a single bit output. (Both G and V are deterministic.) $c(\lambda)$ is the assumption’s threshold.*

We say that P is falsifiable if G, V are uniform polynomial-time algorithms (in their input size).

We also define a property called *verifiably-polynomial image* that roughly speaking requires that any instance has at most polynomial many solutions and that this can be verified in some weak sense. The property in particular captures problems where the solution space $\{0, 1\}^m$ is of polynomial size such as decision problems (where $m = 1$), and problems in **NP** where there are a few solutions per instance (such as injective one-way functions).

Definition 3.2 (Verifiably-Polynomial Image). *A non-interactive assumption P has a verifiably-polynomial image if there exists an efficient verifier K and a polynomial $k = \text{poly}(\lambda)$, such that for any instance $x \in \{0, 1\}^n$, the set $Y_x := \{y : K(1^\lambda, x, y) = 1\}$ of K -valid solutions is of size at most k and for any valid instance $x = G(1^\lambda, r)$ and solution y such that $V(1^\lambda, r, y) = 1$, it holds that $y \in Y_x$.*

The traditional notion of the advantage in solving an assumption P is measured in terms of the distance between the solving probability (which we term the value) and the threshold c .

Definition 3.3 (Value and Advantage of Classical Functions). *Let $P = (G, V, C)$ be a non-interactive assumption and let $f = \{f_\lambda : \{0, 1\}^n \rightarrow \{0, 1\}^m\}_\lambda$ be a family of (possibly randomized) functions. For every $\lambda \in \mathbb{N}$, we define the corresponding value and advantage:*

$$\text{val}_P[f](\lambda) := \Pr \left[V(1^\lambda, r, y) = 1 \mid \begin{array}{l} r \leftarrow \{0, 1\}^d \\ x = G(1^\lambda, r) \\ y \leftarrow f_\lambda(x) \end{array} \right], \text{a}_P[f](\lambda) := |\text{val}_P[f](\lambda) - c(\lambda)|,$$

where the probability is also above the randomness of f_λ in case it is randomized.

3.2 Stateful Solvers

The premise of our work is that in the quantum setting, one ought to think about *stateful* solvers, which generalizes the standard treatment of a solver as a one-shot algorithm. We now define this formally.

Definition 3.4 (Stateful Solvers: Syntax). *Let P be a non-interactive assumption.*

Let $\ell = \ell(\lambda)$ be a function. A classical (resp. quantum) ℓ -stateful solver $\mathcal{B} = (B, \text{state}_0 = \{\text{state}_{\lambda,0}\}_\lambda)$ is defined as follows.

- B is a classical (resp. quantum) algorithm that takes as input $1^\lambda, 1^t, x \in \{0, 1\}^n$ and state which is an ℓ -bit (resp. qubit) string, and outputs a value $y \in \{0, 1\}^m$ and state' which is an ℓ -bit (resp. qubit) next-state. We let $B(\dots)_y$ denote the y output and $B(\dots)_{\text{st}}$ denote the state' output.
- $\text{state}_0 = \{\text{state}_{\lambda,0}\}_\lambda$ is a sequence of classical (resp. quantum) states consisting of $\ell = \ell(\lambda)$ bits (resp. qubits).

We say that \mathcal{B} is efficient if B runs in time $\text{poly}(\lambda, t, n)$; i.e., in polynomial time in the lengths of its inputs.

Remark 3.5 (Non-uniformity). The algorithm B may have a non-uniform classical advice. It does not have any additional quantum advice.

Remark 3.6 (Dependence on Runtime). Our definition allows the running time of efficient stateful solvers to depend polynomially on the “iteration” t . In particular, for any polynomial number of solving attempts $t = \text{poly}(\lambda)$, the overall running of the solver is polynomial. One could also consider a more stringent definition that requires that each call runs in fixed polynomial time independently of the iteration number t . Jumping forward, we will show how a solver can preserve its solving ability through time, but at the cost of running for longer in each step. Doing this according to the more stringent time-independent definition remains an open question.

It will be useful to define some properties of solvers with respect to an extension of the solver’s execution transcript. The extension corresponds to the would-be transcript of a purified version of the solver, running on a purified version of the initial state. This will allow us to get a precise well-defined handle on the evolution of quantum states throughout the lifetime of the solver. The extended transcript will only be used for purposes of definition and analysis and will never be required algorithmically.

Definition 3.7 (Stateful Solvers: Purifying Values). *Consider a solver $\mathcal{B} = (B, \text{state}_0 = \{\text{state}_{\lambda,0}\}_\lambda)$. Let $B_{\lambda,t,x}$ denote the quantum procedure that takes \mathbf{s} as input and produces $B(1^\lambda, 1^t, x, \mathbf{s})$ over registers SY . By Proposition 2.1, we can consider its purification $\widehat{B}_{\lambda,t,x}$ which acts on registers $SY\widehat{Y}$ and takes as input $(\mathbf{s}, \mathbf{0}, \mathbf{0})$. Then define $\widehat{B}(1^\lambda, 1^t, x, \mathbf{s})$ as the algorithm that computes $(\mathbf{s}', \mathbf{y}, \widehat{\mathbf{y}}) = \widehat{B}_{\lambda,t,x}(\mathbf{s}, \mathbf{0}, \mathbf{0})$, measures $(\mathbf{y}, \widehat{\mathbf{y}})$ in the computational basis to obtain (y, \widehat{y}) , and then outputs \mathbf{s}' as the state output, y as the solution output, and \widehat{y} as the purifying output.*

In addition, by Proposition 2.2, there exists a (possibly inefficient) unitary $\widehat{B}_{0,\lambda}$ that operates on two registers $S\widehat{Y}$ such that when applying $(\mathbf{s}_0, \widehat{\mathbf{y}}_0) \leftarrow \widehat{B}_{0,\lambda}(\mathbf{0}, \mathbf{0})$, the reduced density matrix of \mathbf{s}_0 is identical to that of state_0 . Then define $\widehat{B}_0(1^\lambda)$ as the quantum procedure that computes $(\mathbf{s}_0, \widehat{\mathbf{y}}_0) \leftarrow \widehat{B}_{0,\lambda}(\mathbf{0}, \mathbf{0})$, measures $\widehat{\mathbf{y}}_0$ in the computational basis, and then outputs \mathbf{s}_0 as state_0 and $\widehat{\mathbf{y}}_0$ as the purifying initial value.

We refer to the collection $\widehat{\mathcal{B}} = \{\widehat{B}_{i,\lambda,x}\}$ as a purification of \mathcal{B} (it is not unique).

Remark 3.8. We note that the purifying values can be arbitrarily long. These values will only be used for analysis purposes and are never produced in an actual execution, and hence we do not require any bound whatsoever on the length of the purifying values or the complexity of producing them.

We now define the concept of a *solver interaction*, which captures the process of repeatedly invoking a stateful solver by a given algorithm.

Definition 3.9 (Solver Interaction). Let $P = (G, V, c)$ be a non-interactive assumption. For any stateful solver $\mathcal{B} = (B, \text{state}_0)$ and corresponding purification $\widehat{\mathcal{B}}$, and any algorithm A with input $z \in \{0, 1\}^*$, we consider the process $A^{\mathcal{B}}(1^\lambda, z)$ of the algorithm interacting with the solver. We define this process in two different yet equivalent manners: one which is efficient given the ability to execute B , and one which may be inefficient but implies an identical output distribution. The latter will include a production of all purifying values (Definition 3.7) which will be useful for definitions and analysis.

- We let state_0 be as defined in \mathcal{B} .
Equivalently: We let $(\text{state}_0, \hat{y}) \leftarrow \widehat{B}_0(1^\lambda)$.
- A is invoked on input $(1^\lambda, z)$ and at every step $i \geq 1$:
 1. A submits a query $x_i \in \{0, 1\}^n$.
 2. $(y_i, \text{state}_i) \leftarrow B(1^\lambda, 1^i, x_i, \text{state}_{i-1})$ is invoked.
Equivalently: $(\hat{y}_i, y_i, \text{state}_i) \leftarrow \widehat{B}(1^\lambda, 1^i, x_i, \text{state}_{i-1})$ is invoked.
 3. A obtains y_i , and proceeds to the next step.
- At the end of the interaction A may produce an output w .

We sometimes refer to A as a solver-aided algorithm and use the shorthand $A_z^{\mathcal{B}}$ for the solver interaction and $A_z^{\widehat{\mathcal{B}}}$ for the purified solver interaction. We refer to the random variables $\text{state}_0, \text{state}_1, \text{state}_2, \dots$ as the state random variables of the interaction. We refer to the list of pairs of generated instances and solutions (x_i, y_i) as the transcript of the interaction and denote it by ts . We also define the extended transcript $\widehat{\text{ts}}$ of the execution as consisting of the value \hat{y}_0 followed by a list to triples (x_i, y_i, \hat{y}_i) . Given an extended transcript $\widehat{\text{ts}}$, we can produce the standard transcript ts by removing all purifying values. We call this action redaction and say that ts is the redacted transcript induced by $\widehat{\text{ts}}$. Generating an extended transcript according to the purified solver interaction $A_z^{\widehat{\mathcal{B}}}$ and redacting it produces an identical distribution to the generation of the redacted transcript by direct interaction $A_z^{\mathcal{B}}$. The length of a transcript/extended-transcript is the number of pairs/triples it contains (this means that an extended transcript of length 0 is not empty since it still contains \hat{y}_0). The i -prefix of a transcript/extended-transcript is denoted $\text{ts}_i/\widehat{\text{ts}}_i$ and contains the first i pairs/triples (and also \hat{y}_0 in the extended case).

We show that the purifying values indeed purify the entire solver interaction, in the sense that they determine all states state_i as pure states for any solver interaction.

Proposition 3.10. Let $\mathcal{B} = (B, \text{state}_0)$ be a solver with purification $\widehat{\mathcal{B}}$ and consider the extended transcript $\widehat{\text{ts}}$ of the solver interaction $A_z^{\widehat{\mathcal{B}}}$ and let t be its length. Then for all $i \leq t$, the state state_i is pure conditioned on $\widehat{\text{ts}}_i$. Specifically, it has density matrix $|s_{\widehat{\text{ts}}_i}\rangle\langle s_{\widehat{\text{ts}}_i}|$ that is completely determined by $\widehat{\text{ts}}_i$ (and therefore by the classical string $\widehat{\text{ts}}$) and does not depend on any other parameter of the execution.

Proof. We consider the purifying description of the solver interaction $A_z^{\widehat{\mathcal{B}}}$ and prove by induction. For $t = 0$, we recall that the pair $(\text{state}_0, \hat{y})$ is generated by applying $\widehat{B}_{0,\lambda}$ on the zero state, followed by measuring the \hat{Y} register. The pre-measurement state over registers $S\hat{Y}$ is therefore pure, and can always be written as

$$\sum_{\hat{y}} \alpha_{\hat{y}} |s_{\hat{y}}\rangle_S \otimes |\hat{y}\rangle_{\hat{Y}} , \quad (2)$$

where $\alpha_{\hat{y}}$ are non-negative real values with $\sum_{\hat{y}} \alpha_{\hat{y}}^2 = 1$, and $|s_{\hat{y}}\rangle$ are fully specified unit vectors. Therefore, post-selecting on having measured the value \hat{y}_0 in register \hat{Y} , we have that the state in register S is exactly $\text{state}_0 = |s_{\hat{y}_0}\rangle \langle s_{\hat{y}_0}|$, which completes the base step of the proof.

Now assume that the above holds for all $i < t$. Consider a transcript $\widehat{\mathbf{ts}}$ of length t s.t. $\widehat{\mathbf{ts}} = \widehat{\mathbf{ts}}_{t-1} \parallel (x, y, \hat{y})$ for some $\widehat{\mathbf{ts}}_{t-1}, x, y, \hat{y}$.

Let us consider the state of the system right before the t -th query to the solver. At this point, $\widehat{\mathbf{ts}}_{t-1}$ was already determined, and thus by induction we know that $\text{state}_{t-1} = |s_{\widehat{\mathbf{ts}}_{t-1}}\rangle \langle s_{\widehat{\mathbf{ts}}_{t-1}}|$ is a pure state. At this point x has also been determined.

By definition, state_t is produced by executing a unitary $\widehat{B}_{\lambda,t,x}$ (that acts on registers $S\hat{Y}\hat{Y}$) on $(\text{state}_{t-1}, \mathbf{0}, \mathbf{0})$, which is pure by the induction hypothesis, and measuring the $\hat{Y}\hat{Y}$ registers. The analysis here is similar to the base case. The pre-measurement state is pure (since it is induced by applying a unitary on a pure state) and thus can always be written as

$$\sum_{y, \hat{y}} \alpha_{y, \hat{y}} |s_{y, \hat{y}}\rangle_S \otimes |y, \hat{y}\rangle_{\hat{Y}\hat{Y}} , \quad (3)$$

and as above α_y are non-negative real values with $\sum_y \alpha_y^2 = 1$, and $|s_{y, \hat{y}}\rangle$ are fully specified unit vectors. Post selecting on y, \hat{y} leaves us with register S containing $\text{state}_t = |s_{y, \hat{y}}\rangle \langle s_{y, \hat{y}}|$, which completes the proof.

We are now ready to define the concepts of value and advantage of stateful solvers. Traditionally, when thinking about stateless solvers, we consider their *one shot value*, namely the probability that they solve the problem on a random instance. Since they are stateless this probability does not change over time. In the case of stateful solvers, this probability may change over time. Our definition of the *many shot values* aims to capture exactly this. For any solver interaction $A_z^{\mathcal{B}}$, the value at time t , captures the probability that the solver \mathcal{B} successfully solves a random instance at this time, after a given t -round interaction with A_z . This value is, in fact, a random variable that depends on the history of the interaction. To make this precise, we consider any purification $\widehat{\mathcal{B}}$, and define these values as a function of the extended transcript.

Definition 3.11 (Stateful Solvers: Value and Advantage). *Let P be a non-interactive assumption, $\mathcal{B} = (B, \text{state}_0)$ be a corresponding stateful solver, $\widehat{\mathcal{B}}$ a corresponding purification, and A a solver-aided algorithm with input z . For*

every $\lambda, i \in \mathbb{N}$, let state_i be the i -th pure state random variable of the solver interaction $A_z^{\widehat{\mathcal{B}}}$ (determined by $\widehat{\text{ts}}_i$). The corresponding value random variables are:

$$\text{val}_P[i, A_z^{\widehat{\mathcal{B}}}] (\lambda) := \Pr \left[V(1^\lambda, r, y) = 1 \left| \begin{array}{l} r \leftarrow \{0, 1\}^d \\ x = G(1^\lambda, r) \\ (\hat{y}_{i+1}, y, \text{state}_{i+1}) \leftarrow \widehat{B}(1^\lambda, 1^i, x, \text{state}_i) \end{array} \right. \right],$$

where the probability is over the choice of r and the measurement of \hat{y}_{i+1}, y .

The one-shot value of \mathcal{B} is

$$\text{val}_P[0, \mathcal{B}] (\lambda) := \Pr \left[V(1^\lambda, r, y) = 1 \left| \begin{array}{l} r \leftarrow \{0, 1\}^d \\ x = G(1^\lambda, r) \\ (y, \text{state}_1) \leftarrow B(1^\lambda, x, \text{state}_0) \end{array} \right. \right],$$

where the probability is over the choice of r , measurements of B , and (the possibly mixed) state_0 . Note that this is in fact a number, independent of any A or the choice of purification $\widehat{\mathcal{B}}$.

The corresponding advantage random variables are:

$$\mathbf{a}_P[i, A_z^{\widehat{\mathcal{B}}}] (\lambda) := \left| \text{val}_P[i, A_z^{\widehat{\mathcal{B}}}] (\lambda) - c(\lambda) \right| \quad \mathbf{a}_P[0, \mathcal{B}] (\lambda) := \left| \text{val}_P[0, \mathcal{B}] (\lambda) - c(\lambda) \right| .$$

For a distribution \mathbb{B} on solvers $\{\mathcal{B}_\alpha\}_\alpha$, we define the one-shot value of the distribution as:

$$\text{val}_P[0, \mathbb{B}] (\lambda) = \mathbb{E}_{\alpha \leftarrow \mathbb{B}} [\text{val}_P[0, \mathcal{B}_\alpha] (\lambda)] .$$

The corresponding advantage is $\mathbf{a}_P[0, \mathbb{B}] (\lambda) = |c(\lambda) - \text{val}_P[0, \mathbb{B}] (\lambda)|$.

As the solver's state evolves over time, its advantage in solving an assumption may reduce or disappear altogether. This is in particular relevant to the quantum setting, where when a solver is invoked its internal state is disturbed. Aiming to capture solvers that remain useful over time, we next define the notion of solvers with *persistent value*, namely, solvers whose value in solving a given assumption is preserved through time. We define it more generally for distributions over solvers; single solvers are a special case.

Definition 3.12 (Persistent Value). *Let P be a non-interactive assumption. A distribution \mathbb{B} on solvers $\{\mathcal{B}_\alpha\}_\alpha$ is η -persistent if there exist purifications $\{\widehat{\mathcal{B}}_\alpha\}_\alpha$ such that for any algorithm A with input z , with probability $1 - \eta$ over the choice of solver $\alpha \leftarrow \mathbb{B}$ and over an extended transcript $\widehat{\text{ts}}$ in the solver interaction process $A_z^{\widehat{\mathcal{B}}_\alpha}$, there exists a value p such that:*

$$\max_i \left| \text{val}_P[i, A_z^{\widehat{\mathcal{B}}_\alpha}] - p \right| \leq \eta . \quad (4)$$

We call p a persistent value. Given a random variable $p^*(\alpha) \subseteq [0, 1]$, we say that a solver distribution is (p^*, η) -persistent if the condition holds for $p^*(\alpha)$.

We next define the notion of a persistent advantage. This aims to capture the case that solvers maintain a lower bound on their advantage through time.

Definition 3.13 (Persistent Advantage). *Let P be a non-interactive assumption with threshold c . A distribution \mathbb{B} on solvers $\{\mathcal{B}_\alpha\}_\alpha$ has ε -persistent advantage if there exist purifications $\{\widehat{\mathcal{B}}_\alpha\}_\alpha$ such that for any algorithm A with input z :*

$$\mathbb{E} \left[\min_i \text{val}_P[i, A_z^{\widehat{\mathcal{B}}_\alpha}] \right] \geq c + \varepsilon , \quad (5)$$

where the expectation is over the choice of solver $\alpha \leftarrow \mathbb{B}$ and over an extended transcript $\widehat{\text{ts}}$ in the solver interaction process $A_z^{\widehat{\mathcal{B}}_\alpha}$.

In the above, We require that the advantage has a consistent sign (for simplicity, positive). Intuitively, the reason we focus on persistence of the positive advantage $v_t - c$ at time t , rather than the absolute advantage $|v_t - c|$, is that if the sign of $v_t - c$ arbitrarily changes after each solver invocation, then the solver may not be as useful. (As a simple example, take a deterministic distinguisher and turn it into a stateful distinguisher that flips the output of the original distinguisher at random with each invocation, deeming it useless.) We note that η persistent solvers in particular preserve the sign of their advantage (up to η).

Memoryless and Stateless Solvers. A special case of the above definitions is that of *memoryless and stateless solvers*.

Definition 3.14. *A solver $\mathcal{B} = (B, \text{state}_0)$ is memoryless if the size of its state is $\ell = 0$. The solver is stateless if in addition (to being memoryless), the algorithm B does not depend on 1^t (in functionality or runtime).*

Remark 3.15 (Persistent Value for Stateless and Memoryless Solvers). Note that in the case of stateless solvers, successive invocations of the solver will always result in the same output distribution. Here the one-shot (and many-shot) advantage coincide with the standard notion of advantage for functions (Definition 3.3) and values are persistent (Definition 3.12). Accordingly, stateless solvers exactly capture the traditional notion of classical solvers, given by a randomized function.

Moreover, even for memoryless solvers, when considering the definition of persistent solvers the value $\text{val}_P[i, A_z^{\widehat{\mathcal{B}}_\alpha}]$ does not depend on A_z at all (only on i), and therefore it is a fixed number rather than a random variable. It follows that for (p, η) -persistent memoryless solvers, Eq. (4) holds with probability 1.

3.3 Reductions

We now define the notion of a reduction. A reduction is a way to prove a claim of the form “if there exists a successful solver for assumption P then there exists a successful solver for an assumption Q ”. We consider *constructive reductions*

in the sense that they are an explicit uniform algorithm that takes as input a successful solver for P and efficiently solves the problem Q .

The default notion of a reduction in the literature is *one shot*. In particular, a given quantum P -solver is only assumed to have a meaningful one-shot advantage in solving P , and there is no a priori guarantee on its advantage in any many shot solving process, in particular there may not be any value persistence. Likewise, the produced solver for the assumption Q is only required to have a meaningful one-shot advantage. Below we define both the default notion of one-shot reductions as well as the stronger notion of *durable reductions* requiring that the resulting Q -solver also has persistent advantage, meaning that with noticeable probability, the reduction can go on solving for an arbitrary polynomial number of times.

Definition 3.16 (Reduction). *A reduction from classically (resp. quantumly) solving a non-interactive assumption Q to classically (resp. quantumly) solving a non-interactive assumption P is an efficient classical (resp. quantum) uniform algorithm \mathcal{R} with the following guarantee. For any solver $\mathcal{B}_P = (B_P, \text{state}_0)$ for P with one-shot advantage ε and running time T , let $\text{state}'_0 = (\text{state}_0, B_P, 1^{1/\varepsilon}, 1^T)$. Then $\mathcal{B}_Q = (\mathcal{R}, \text{state}'_0)$ is a solver for Q with one-shot advantage $\varepsilon' = \text{poly}(\varepsilon, T^{-1}, \lambda^{-1})$ and running-time $\text{poly}(T, \varepsilon^{-1}, \lambda)$. We say that the reduction is durable if \mathcal{B}_Q has $\text{poly}(\varepsilon, T^{-1}, \lambda^{-1})$ -persistent advantage.*

We refer to a reduction from solving Q to classically (resp. quantumly) solving P as a classical-solver (resp. quantum-solver) reduction.

Remark 3.17 (Many Shot Reductions). There could be several conceivable extensions of the above definition that also account for the *many-shot advantage*. One such natural extension is requiring that the reduction works only given a solver with a persistent value (as in Definition 3.12). Jumping ahead, in section 4, we show that under certain conditions, persistent solving can in fact be reduced to one-shot solving, even in the quantum setting.

Remark 3.18 (The Loss). We allow for a (fixed) polynomial loss in the advantage and running time. One could naturally extend it to more general relations.

Classical Black-Box Reductions. In this work, we prove that several general classes of *classical reductions* that a priori are only guaranteed to work for classical solvers, can be enhanced *efficiently* to also work for quantum solvers. Our focus is on black-box reductions; that is, reductions that are oblivious of the representation and inner workings of the solver that they use (in contrast to the above Definition 3.16, where the reduction obtains the full description of the solver \mathcal{B}_P).

We next formally define such black box reductions, using the terminology we have already developed. Specifically, we capture the notion of a classical solver for a given problem P as a stateless (classical) solver.

Definition 3.19 (Classical Black-Box Reduction). *A classical black-box reduction, from solving a non-interactive assumption Q to solving a non-interactive*

assumption P , is an efficient classical solver-aided uniform algorithm \mathcal{R} with the following syntax and guarantee. \mathcal{R} takes as input a security parameter 1^λ , parameter $1^{1/\varepsilon}$, and instance $x \in \{0, 1\}^{n_Q}$ of Q . It interacts with a solver \mathcal{B} for P (per Definition 3.9) and produces an output $y \in \{0, 1\}^{m_Q}$. We require that for any distribution \mathbb{B} over **stateless classical solvers** $\{\mathcal{B}_\alpha\}_\alpha$ such that \mathbb{B} has advantage at least ε in solving P , the corresponding solver distribution \mathbb{R} over solvers $\{\mathcal{R}^{\mathcal{B}_\alpha}(1^\lambda, 1^{1/\varepsilon}, \cdot)\}_\alpha$ has advantage at least $\text{poly}(\varepsilon, \lambda^{-1})$ in solving Q . The advantage of \mathcal{R} is positive if its value is always at least c_Q (above the assumption Q 's threshold), regardless of any P -solver.

We further say that the reduction \mathcal{R} is *non-adaptive* if \mathcal{R} produces all of its oracle queries $x_1, \dots, x_k \in \{0, 1\}^{n_P}$ to \mathcal{B} in one shot, obtains all answers y_1, \dots, y_k , and then produces its output y .

Remark 3.20. In our definition of solver interaction, a given solver \mathcal{B} is only ever invoked for the instance size $n_P(\lambda)$. Accordingly, the above definition restricts attention to classical reductions that in order to solve problem Q for instance size $n_Q(\lambda)$ make queries to a P -solver on a specific related input size $n_p(\lambda)$. While this is not without loss of generality, it does capture natural reductions. (In fact, we are not aware of important reductions that do not adhere to this.)

Remark 3.21 (Deterministic Solver Reductions, Positive Advantage, and Repeated Queries). We consider classical reductions that ought to work when given a stateless solver from a distribution \mathbb{B} over solvers $\{\mathcal{B}_\alpha\}$. (As a matter of fact in our model, even once a stateless solver \mathcal{B}_α is fixed, the process of answering any given query is randomized, but this can be modeled as sampling a deterministic stateless solver from another distribution \mathbb{B} with the same advantage.) A weaker notion of classical reductions only requires that the reduction works for deterministic solvers. In the classical setting, this is typically not an issue, as long as the reduction has the power to fix the solver's randomness and repeatedly replace it as needed. Jumping forward, when considering quantum reductions, the randomness of a given solver may arise from the quantum nature of the solving process, and the reduction may not be able to control it. Accordingly, in our transformations from classical-solver reductions to quantum-solver reductions, we will naturally need the classical reduction we start from to also be able to deal with distributions over solvers.

We note that for typical assumptions Q such as search problems (with trivial threshold $c = 0$) or decision problems (with solution length $m = 1$ and trivial threshold $c = 1/2$), a classical reduction \mathcal{R} from Q -solving to deterministic P -solving implies a classical reduction \mathcal{R}' from Q -solving to distributional P -solving. Here two subtleties should be addressed. The first issue that could prevent \mathcal{R} from working for distributional P -solvers is that the sign of the advantage of $\mathcal{R}^{\mathcal{B}}$ as a Q -solver may depend on the randomness of \mathcal{B} and may cancel out in expectation. For search assumptions Q , where $c = 0$, this cannot happen as any advantage is positive. For decision problems, this can be avoided by slightly augmenting \mathcal{R} to make sure that the advantage is always positive using standard black-box techniques [6]. This incurs only a polynomial overhead in

solving queries, or even just a single query, at the cost of quadratically decreasing the advantage. The second issue concerns the running time of the reduction. Specifically a reduction that works for deterministic oracles, excepts to get their advantage $1^{1/\varepsilon}$ as input, where ε is the P -solver's advantage. When executing such a reduction with a solver distribution, we are given $1^{1/\varepsilon}$, where ε is the average advantage. Nevertheless, we can run the original reduction with input $1^{2/\varepsilon}$. Note that the probability that the advantage of a sampled oracle is at least $\varepsilon/2$ is at least $\varepsilon/2$, and since the reduction has positive advantage, we are overall guaranteed to maintain a noticeable advantage.

Following the above, for typical assumptions Q , we can in particular assume w.l.o.g positive advantage. For simplicity, we also assume throughout that classical reductions We do not repeat queries. This is w.l.o.g as given a deterministic oracle, the reduction can simply store previous answers and answer consistently by itself.

4 Persistent Solvers in the Quantum Setting

In this section, invoking state restoration techniques from [7], we prove that any one-shot solver for an assumption P with a verifiably-polynomial image (in particular, decision problems) can be converted into a persistent solver for P .

Theorem 4.1 (Persistence Theorem). *Let P be a non-interactive falsifiable assumption with a verifiably-polynomial image. For any inverse polynomial function η , there exist efficient quantum algorithms S, R with the following syntax and guarantee. $S_B(\text{state}_0)$ takes as input a quantum algorithm B and state state_0 and outputs a state state_0^* and a value $p^* \in [0, 1]$. $R_B(1^\lambda, 1^i, x, \text{state}_{i-1}^*)$ takes as input B , a security parameter 1^λ , step 1^i , input $x \in \{0, 1\}^n$, and state state_{i-1}^* and outputs a solution $y \in \{0, 1\}^m$ and state state_i^* .*

For any solver $\mathcal{B} = (B, \text{state}_0)$ with one-shot value $p = \text{val}_P[0, \mathcal{B}]$, considering the random variable $(\text{state}_0^, p^*) \leftarrow S_B(\text{state}_0)$, it holds that:*

1. $\mathbb{E}[p^*] = p$.
2. $\mathcal{R}^* = (R_B, \text{state}_0^*)$ sampled in this process is a distribution over efficient stateful solvers that is (p^*, η) -persistent.

Remark 4.2. The efficiency of the algorithms S, R is also polynomial in the running time of B . We avoid passing explicitly the running time bound as input to simplify notation.

The proof relies on techniques from [7] and can be found in the full version of this work [3].

5 Stateful Solvers To Memoryless Solvers

The following theorem shows that it is possible to convert stateful solvers into memoryless solvers with the same value, albeit with a few caveats. First, the

distribution of queries that is to be made to the memoryless solver needs to be known ahead of time (i.e. it needs to be decided upfront in a non-adaptive manner). Second, the resulting memoryless solver might not be efficiently executable. Instead, we provide a simulator that can emulate its behavior, but only once, and only on an input that comes from the prescribed distribution. The simulator only manages to simulate the execution up to some statistical error, and its running time is polynomial in the inverse of this error. A formal theorem statement follows.

Theorem 5.1. *There exists a polynomial time oracle-aided simulator SimMemless with the following properties. Let \mathcal{B} be a (p, η) -persistent ℓ -stateful solver for a falsifiable non-interactive assumption P and let $D = \{D_\lambda\}_\lambda$ be an efficiently samplable distribution ensemble over k -tuples of P instances. Finally, let δ be some parameter. Then there exists a (p, η) -persistent (but possibly inefficient) distribution over memoryless solvers $\mathcal{B}' = \mathcal{B}'_{\ell, D, \delta} = (\mathcal{B}', \emptyset)$ for P such that the following holds.*

Consider sampling $\vec{x} \leftarrow D_\lambda$, and let $\mathcal{B}'(1^\lambda, \vec{x})$ be the transcript of the process that feeds the elements of \vec{x} into \mathcal{B}' one-by-one in order (i.e. executes $\mathcal{B}'(1^\lambda, 1^i, x_i, \emptyset)$ in order). Then $\text{SimMemless}^{\mathcal{B}, D}(1^\lambda, 1^\ell, 1^{1/\delta}, \vec{x})$ makes non-adaptive black-box access to \mathcal{B} and produces a distribution that is within at most δ statistical distance from $\mathcal{B}'(1^\lambda, \vec{x})$.

We note that our simulator is “almost” a black-box algorithm in \mathcal{B} in the sense that it takes the size of the state 1^ℓ as input, but otherwise it only makes black-box queries to \mathcal{B} . We also emphasize that the simulator does not depend at all on p, η or any other property of \mathcal{B} (other than ℓ).

5.1 The Simulator SimMemless

We start by describing the simulator that will be used to prove Theorem 5.1. The simulator SimMemless simply “floods” the solver \mathcal{B} with queries from a fixed distribution, and plants the elements of \vec{x} in random positions.

Specifically, $\text{SimMemless}^{\mathcal{B}, D}(1^\lambda, 1^\ell, 1^{1/\delta}, \vec{x})$ works as follows. Let t be such that $k\sqrt{\ell/2t} \leq \delta$, i.e. $t = O(\ell(k/\delta)^2)$. The simulator is also going to generate a non-adaptive sequence of queries. We start by defining our “flooding” distribution.

Definition 5.2 (Random Marginal). *Let D be a distribution over X^k , i.e. k -tuples over a domain X . Then the random marginal distribution D_U over X is a distribution obtained by sampling (x_1, \dots, x_k) according to D , sampling a random i in $[k]$, and outputting x_i as the final sample.*

The simulator starts by sampling the following values.

1. A vector \vec{z} of $k \cdot t$ samples $z_{j,i} \leftarrow D_U$, where $j \in [k]$, $i \in [t]$.
2. k uniform samples $i_j \leftarrow [t]$, where $j \in [k]$.
3. A uniform permutation π over $[k]$.

It then generates a sequence of queries \vec{z}^* by taking the vector \vec{z} and, for all $j \in [k]$, replacing z_{j,i_j} with $x_{\pi(j)}$. Namely, thinking of \vec{z} as containing k sequences of queries of length t each, we plug in a random element from \vec{x} in a random location in each sequence.

The simulator then calls \mathcal{B} on the queries in \vec{z}^* in order, to obtain a sequence of responses \vec{y} . Let $y_{j,i}$ be the (j,i) element in this sequence. We define $y_j^* = y_{j,i_j}$. The simulator returns the transcript $((x_1, y_{\pi^{-1}(1)}^*), \dots, (x_k, y_{\pi^{-1}(k)}^*))$. Namely, we output a transcript that pairs each x_i with the response that \mathcal{B} produces when introduced to the query $z_{j,i_j} = x_i$, namely $\pi(j) = i$.

5.2 Proving Theorem 5.1

We now turn to prove the theorem. We start by defining a hybrid distribution which is defined with respect to purifying executions of \mathcal{B} . This will allow us to make claims about extended transcripts, and finally to redact to standard transcript and derive the proof of the theorem.

A Hybrid Distribution. To prove the theorem, we define the hybrid distribution \mathcal{S}_h , defined for every $h \in \{0, 1, \dots, k\}$.

1. Sample a uniform permutation π over $[k]$.
2. For all $j \in [k]$, sample a random index $i_j \in [t]$.
3. Sample \vec{x} from D .
4. Generate a sequence of queries $z_{j,i}$ for all $j \in [k]$, $i \in [t]$ as follows.
 - (a) For all $j > h$, set $z_{j,i_j} = x_{\pi(j)}$.
 - (b) Otherwise sample z_{j,i_j} from D_U .
5. Generate the extended transcript $\widehat{\mathbf{ts}}$ of executing \mathcal{B} (in a purifying manner) on the entries $z_{j,i}$ in lexicographic order (i.e. starting with $(1, 1), \dots, (1, t)$ and concluding with $(k, 1), \dots, (k, t)$). We let $\widehat{\mathbf{ts}}_{j,i}$ denote the prefix of the transcript prior to making the (j, i) query. We let $|s_{j,i}\rangle$ denote the solver state respective to $\widehat{\mathbf{ts}}_{j,i}$, as guaranteed by Proposition 3.10. Notice that $|s_{1,1}\rangle$ is the initial state \mathbf{state}_0 of \mathcal{B} conditioned on $\widehat{\mathbf{ts}}_0 = \widehat{y}_0$.
6. The output of the hybrid \mathcal{S}_h then consists the following values, for all $j \in [k]$:
 - (a) The values $i_j, \pi(j)$.
 - (b) The quantum state in the beginning of the j -th run: $|s_{j,1}\rangle$.
 - (c) The quantum state right before the i_j -th query in the j -th sequence is made: $|s_{j,i_j}\rangle$.
 - (d) The value $x_{\pi(j)}$, which is the i_j -th query in the j -th sequence if $j > h$.
 - (e) An answer $(y_{\pi(j)}, \widehat{y}_{\pi(j)})$ computed as follows.
 - If $j > h$ then set $(y_{\pi(j)}, \widehat{y}_{\pi(j)}) = (y_{j,i_j}, \widehat{y}_{j,i_j})$ (i.e. the (y, \widehat{y}) -part of the (j, i_j) -th triple in $\widehat{\mathbf{ts}}$).
 - Otherwise generate $(y_{\pi(j)}, \widehat{y}_{\pi(j)})$ as $\widehat{B}(1^\lambda, 1^{t(j-1)+i_j}, x_{\pi(j)}, |s_{j,i_j-1}\rangle)_{y, \widehat{y}}$.

In what follows, we will prove that the distributions induced by the first and last hybrids are close in trace distance, as formalized below.

Lemma 5.3. *It holds that $\text{TD}(\mathcal{S}_0, \mathcal{S}_k) \leq k\sqrt{\ell/(2t)}$.*

Before proving Lemma 5.3, we argue that it implies the validity of Theorem 5.1. Indeed, we observe that the output of the simulator `SimMemless` can be extracted from \mathcal{S}_0 by simply outputting all of the pairs $((x_1, y_1), \dots, (x_k, y_k))$. Applying the same extraction procedure on the last hybrid \mathcal{S}_k will lead to a sequence $((x_1, y_1), \dots, (x_k, y_k))$ in which $y_{\pi(j)} = B(1^\lambda, 1^{t(j-1)+i_j}, x_{\pi(j)}, |s_{j,i_j-1}\rangle)_y$. However, in the hybrid \mathcal{S}_k , the transcript $\widehat{\text{ts}}$, and therefore all states $|s_{j,i}\rangle$, are generated independently of \vec{x} . Therefore, for every values of $\pi, \widehat{\text{ts}}$ one could define a memoryless adversary $\mathcal{B}' = (B'_{\pi, \widehat{\text{ts}}}, \emptyset)$, defined by

$$B'_{\pi, \widehat{\text{ts}}}(1^\lambda, 1^j, x, \emptyset) = B(1^\lambda, 1^{t(j-1)+i_{j'}}, x, |s_{j',i_{j'}-1}\rangle)_y, \quad (6)$$

with $j' = \pi^{-1}(j)$. Note that the sequence of states is hard-wired into B' and it does not require to propagate a state throughout the execution.

We therefore indeed have that the solver \mathcal{B}' is a distribution over memoryless solvers indicated by sampling $\pi, \widehat{\text{ts}}$ from their respective distributions and executing $B'_{\pi, \widehat{\text{ts}}}$. Since \mathcal{B} is (p, η) -persistent, we have that with probability $1 - \eta$ over $\widehat{\text{ts}}$, all invocations of $B(1^\lambda, 1^{t(j-1)+i_{j'}}, x, |s_{j',i_{j'}-1}\rangle)$ have value $p \pm \eta$, which would imply that $(B'_{\pi, \widehat{\text{ts}}}, \emptyset)$ is (p, η) -persistent. Therefore, the distribution \mathcal{B}' is also, by definition, (p, η) -persistent.

The proof of Lemma 5.3 will follow from a standard hybrid argument, given by the following lemma.

Lemma 5.4. *For all $h \in \{0, 1, \dots, k-1\}$ it holds that*

$$\text{TD}(\mathcal{S}_h, \mathcal{S}_{h+1}) \leq \sqrt{\ell/(2t)}. \quad (7)$$

Proof. We will show that the lemma holds true even when conditioning both $\mathcal{S}_h, \mathcal{S}_{h+1}$ on any value for $\widehat{\text{ts}}_{h,1}$ (the $(h \cdot t)$ -prefix of the transcript $\widehat{\text{ts}}$).

We will show that the lemma follows from the following claim.

Claim 5.5. *Conditioning on any value of $\widehat{\text{ts}}_{h,1}$ for both $\mathcal{S}_h, \mathcal{S}_{h+1}$, the joint distribution of:*

$$(i_h, \widehat{\text{ts}}_{h,i_h}, |s_{h+1,1}\rangle, (x_{\pi(h)}, y_{\pi(h)}, \hat{y}_{\pi(h)})) \quad (8)$$

is within trace distance $\sqrt{\ell/(2t)}$ between $\mathcal{S}_h, \mathcal{S}_{h+1}$.

Given Claim 5.5, Lemma 5.4 follows since all other elements of the two distributions $\mathcal{S}_h, \mathcal{S}_{h+1}$ can be sampled given $\widehat{\text{ts}}_{h,1}$ and $(i_h, \widehat{\text{ts}}_{h,i_h}, |s_{h+1,1}\rangle, (x_{\pi(h)}, y_{\pi(h)}, \hat{y}_{\pi(h)}))$, as follows.

1. Sample the permutation π and the query vector \vec{x} conditioned on the value $x_{\pi(h)}$.
2. For every $j \in [k] \setminus \{h\}$, sample i_j uniformly in $[t]$.

3. For all $j < h$, the transcript prefix $\widehat{\mathbf{ts}}_{h,1}$ determines all states $|s_{j,i}\rangle$ (for all $i \in [t]$), which in turn, together with \vec{x} , determines the distribution of $y_{\pi(j)}, \hat{y}_{\pi(j)}$ for all $j < h$ (since this distribution is specified by applying the solver B on $x_{\pi(j)}$ with quantum state that is determined by the h -prefix).
4. For all $j > h$ the outputs of both $\mathcal{S}_h, \mathcal{S}_{h+1}$ are determined as the outcomes of an identical quantum process applied to the state $|s_{h+1,1}\rangle$ (the initial state of the $(h+1)$ -th sequence), considering that π and \vec{x} have been determined.

We now proceed to prove Claim 5.5, and focus on the distribution of $(i_h, \widehat{\mathbf{ts}}_{h,i_h}, |s_{h+1,1}\rangle, (x_{\pi(h)}, y_{\pi(h)}, \hat{y}_{\pi(h)}))$ in the two hybrids, given that $\widehat{\mathbf{ts}}_{h,1}$ is fixed. The claim follows straightforwardly from our information theoretic Plug-In Lemma (Lemma 2.3), where the classical values y_i in the lemma corresponds to pairs $(z_{h,i}, y_{h,i}, \hat{y}_{h,i})$ generated in the h 'th round in the hybrid experiment. Note that since we fixed $\widehat{\mathbf{ts}}_{h,1}$, the distribution over these classical values is also fixed, and indeed the value $\mathbf{s} = |s_{h+1,1}\rangle$ depends on this sequence of t values. The triple $(x_{\pi(h)}, y_{\pi(h)}, \hat{y}_{\pi(h)})$ differs between \mathcal{S}_h and \mathcal{S}_{h+1} since in the former it is exactly equal to the i_{h+1} element in the h -th sequence, and in the latter it is sampled from the marginal distribution of this element. We can therefore apply the plug-in lemma directly to obtain the $\sqrt{\ell/(2t)}$ bound on the trace distance as Claim 5.5 requires. This completes the proof of the claim and thus also of the lemma.

6 Memoryless Solvers To Stateless Solvers

Theorem 6.1. *There exists a polynomial-time oracle-aided simulator SimStateless with the following properties. Let \mathcal{B} be a (p, η) -persistent memoryless solver for a falsifiable non-interactive assumption P and let $\{D_\lambda\}_\lambda$ be an efficiently samplable distribution ensemble over k -tuples of P instances. Let δ be some parameter.*

Then there exists a (p, η) -persistent (but possibly inefficient) stateless solver $\mathcal{B}'' = \mathcal{B}''_\delta = (\mathcal{B}'', \emptyset)$ for P such that the following holds. Consider sampling $\vec{x} \leftarrow D_\lambda$, and let $\mathcal{B}''(1^\lambda, \vec{x})$ be the transcript of the process that feeds the elements of \vec{x} into \mathcal{B}'' (i.e. executes $B''(1^\lambda, x_i, \emptyset)$ for all x_i). Then $\text{SimStateless}^\mathcal{B}(1^\lambda, 1^{1/\delta}, \vec{x})$ makes non-adaptive black-box access to \mathcal{B} and produces a distribution that is within at most δ statistical distance from $\mathcal{B}''(1^\lambda, \vec{x})$.

Proof. The simulator $\text{SimStateless}^\mathcal{B}$ runs as follows. Given \vec{x} as input, it generates a query vector \vec{x}' of length $t = k^2$ as follows. It samples, without repetitions, k indices i_1, \dots, i_k and sets $x'_j = x_{i_j}$. All other values of x' are set to 0 (or some other fixed value).

After making the queries in \vec{x}' to \mathcal{B} and receiving an output vector \vec{y}' , the simulator sets $y_j = y_{i_j}$ returns $((x_1, y_1), \dots, (x_k, y_k))$.

Let us now define the stateless adversary \mathcal{B}'' . On input x , $B''(1^\lambda, x)$ samples $j \leftarrow [t]$ uniformly, and outputs $y = B(1^\lambda, 1^j, x, \emptyset)_y$. The solver \mathcal{B}'' is also (p, η) -persistent; indeed, its value is the average of values, which are all η -close to

p . (Recall Remark 3.15 about persistent values for stateless and memoryless solvers.)

To bound the statistical distance between $\text{SimStateless}^{\mathcal{B}}(1^\lambda, 1^{1/\delta}, \vec{x})$ and $\mathcal{B}''(1^\lambda, \vec{x})$, we consider the case where in the course of the execution of $\mathcal{B}''(1^\lambda, \vec{x})$, all j 's that are sampled are distinct. This happens with probability at least $1 - k^2/t = 1 - \delta$. Conditioned on this event, $\mathcal{B}''(1^\lambda, \vec{x})$ is identically distributed as $\text{SimStateless}^{\mathcal{B}}(1^\lambda, 1^{1/\delta}, \vec{x})$. It follows that in general the statistical distance is bounded by δ .

We conclude with a corollary that combines Theorem 5.1 and Theorem 6.1.

Corollary 6.2. *There exists a polynomial-time simulator Sim with the following properties. Let \mathcal{B} be a (p, η) -persistent ℓ -stateful solver for a falsifiable non-interactive assumption P and let $\{D_\lambda\}_\lambda$ be an efficiently samplable distribution ensemble over k -tuples of P instances. Finally, let δ be some parameter.*

Then there exists a (p, η) -persistent (but possibly inefficient) distribution over stateless solvers $\mathcal{B}'' = \mathcal{B}''_{\ell, D, \delta} = (\mathcal{B}'', \emptyset)$ for P . Consider sampling $\vec{x}^ \leftarrow D_\lambda$, and let $\mathcal{B}''(1^\lambda, \vec{x}^*)$ be the transcript of the process that feeds the elements of \vec{x}^* into \mathcal{B}'' (i.e. executes $\mathcal{B}''(1^\lambda, x_i^*, \emptyset)$ for all x_i^*). Then $\text{Sim}^{\mathcal{B}, D}(1^\lambda, 1^\ell, 1^{1/\delta}, \vec{x}^*)$ makes non-adaptive black-box access to \mathcal{B} and produces a distribution that is within at most δ statistical distance from $\mathcal{B}''(1^\lambda, \vec{x}^*)$.*

The proof can be found in the full version of the paper [3].

7 Classical Non-Adaptive Reductions and Quantum Solvers

In this section, we show that a wide class of classical reductions can be translated to the quantum setting. Specifically we start from any non-adaptive black-box reductions from classically solving P with a verifiably-polynomial image (Definition 3.2), to classically solving Q . We transform it into a quantum reduction from quantumly solving P to quantumly solving Q .

Theorem 7.1. *Assume there exists a classical non-adaptive black-box reduction from solving a non-interactive assumption Q to solving a non-interactive assumption P with a verifiably-polynomial image. Then there exists a quantum reduction from solving Q to quantumly solving P . This reduction is durable if the original classical reduction has positive advantage.*

Proof. Let \mathcal{R} be a classical non-adaptive black-box reduction from solving a non-interactive assumption $Q = (G_Q, V_Q, c_Q)$ to solving a non-interactive assumption $P = (G_P, V_P, c_P)$. We present a quantum reduction \mathcal{R}' from solving Q to quantumly solving P . We start by describing and analyzing \mathcal{R}' with a one-shot advantage, and then extend it to address durability in the case that \mathcal{R} has positive advantage. We assume w.l.o.g that \mathcal{R} never makes the same query twice to its oracle function (see Remark 3.21).

Recalling Definition 3.16, \mathcal{R}' takes as input $(1^\lambda, 1^t, x_Q, \text{state})$, where $x_Q \in \{0, 1\}^{n_Q}$ is potentially an instance of Q , and its initial state is $\text{state}'_0 = (\text{state}_0, B, 1^{1/\varepsilon}, 1^T)$, where we are guaranteed that $\mathcal{B} = (\text{state}_0, B)$ is a P solver with advantage at least ε that runs in time at most T .

We let ε' denote the advantage of \mathcal{R} in solving Q when given access to an oracle that solves P with advantage at least $\varepsilon/2$. We are guaranteed that $\varepsilon' = \text{poly}(\varepsilon, \lambda^{-1})$. We set $\delta = \varepsilon'/2$ and $\eta = \min\{\varepsilon/4, \varepsilon'/2\}$.

We define a distribution D over $(\{0, 1\}^{n_P})^k$ as the distribution over the set of oracle queries produced by first sampling a uniform r'_Q and using it to generate $x'_Q = G_Q(1^\lambda, r'_Q)$, and finally executing $\mathcal{R}(1^\lambda, 1^{4/\varepsilon}, x'_Q)$ to produce a k -tuple of P -instances.

Having all of these definitions in place, we can now introduce the execution of $\mathcal{R}'(1^\lambda, 1^0, x_Q, \text{state}'_0)$. Namely, we start by analyzing the one-shot execution of \mathcal{R}' (the case $t = 0$).

1. Let R, S be the state restoration algorithms with respect to P as guaranteed by Theorem 4.1, with parameter η as defined above. Set $(\text{state}^*_0, p^*) \leftarrow S_B(\text{state}_0)$. Define $\mathcal{B}_0 = \mathcal{R}^* = (R_B, \text{state}^*_0)$ and recall that \mathcal{B}_0 is (p^*, η) -persistent, and that $\mathbb{E}[p^*] = p$.
2. Execute $\mathcal{R}(1^\lambda, 1^{3/\varepsilon}, x_Q)$ to obtain the sequence of queries \vec{x} .
3. Recall the simulator Sim guaranteed by Corollary 6.2. Execute $\text{Sim}^{\mathcal{B}_0, D}(1^\lambda, 1^\ell, 1^{1/\delta}, \vec{x})$ to obtain a transcript ts .
4. Extract the responses to \vec{x} from ts and resume the execution \mathcal{R} from step 2 with these responses. Once the execution of \mathcal{R} completes and a value y_Q is output, output y_Q as the output of \mathcal{R}' .

To analyze the one-shot value and advantage of \mathcal{R}' , we start by analyzing the performance of \mathcal{R}' conditioned on obtaining a fixed value p^* in step 1 of the execution. In this case \mathcal{B}_0 is (p^*, η) -persistent, and we can invoke Corollary 6.2 to conclude that there exists a (p^*, η) -persistent distribution over stateless adversaries \mathcal{B}''_{p^*} s.t. the output of \mathcal{R}' is within statistical distance δ from the execution of $\mathcal{R}^{\mathcal{B}''_{p^*}}(1^\lambda, 1^{4/\varepsilon}, x_Q)$.

In turn, the execution of $\mathcal{R}^{\mathcal{B}''_{p^*}}(1^\lambda, 1^0, x_Q, \text{state}'_0)$ is equivalent to executing $\mathcal{R}^{\mathcal{B}''}(1^\lambda, 1^{4/\varepsilon}, x_Q)$, where \mathcal{B}'' is a distribution over stateless solvers defined as follows. First sample p^* from its designated distribution, then sample \mathcal{B}''_{p^*} from the (p^*, η) -persistent distribution of stateless solvers. Recall that with probability $1 - \eta$ over the sampling of \mathcal{B}''_{p^*} , it holds that the outcome is a (single) (p^*, η) -persistent stateless solver and therefore that $|\text{val}_P[0, \mathcal{B}''_{p^*}] - p^*| \leq \eta$. It follows that with probability at least $1 - \eta$:

$$\begin{aligned} |\mathbb{E}[\text{val}_P[0, \mathcal{B}'']] - p| &= |\mathbb{E}[\text{val}_P[0, \mathcal{B}''_{p^*}] - p^*]| \\ &\leq \mathbb{E}[|\text{val}_P[0, \mathcal{B}''_{p^*}] - p^*|] \\ &\leq \eta. \end{aligned}$$

It follows that \mathcal{B}'' has advantage at least $\varepsilon - 2\eta \geq \varepsilon/2$ in solving P . We have therefore that $\mathcal{R}^{\mathcal{B}''}$ has advantage at least ε' in solving Q . Since the output of

\mathcal{R}' is within $\delta = \varepsilon'/2$ statistical distance from $\mathcal{R}^{\mathcal{B}''}$, we conclude that \mathcal{R}' has advantage at least $\varepsilon'/2$. We therefore established the one-shot value of \mathcal{R}' .

In the full version [3] we further show that \mathcal{R}' is durable.

8 An Impossibility Result for Search Assumptions

Our result in Section 7 transforms a classical non-adaptive reduction \mathcal{R} from solving Q to classically solving P into a reduction \mathcal{R} to *quantumly* solving P . It is restricted to assumptions P with a verifiably-polynomial image. While this captures a large class of assumptions, such as all decision assumptions, it certainly does not capture all assumptions of interest. In particular, it does not capture *search assumptions* where the number of possible solutions per instance could be super polynomial, such as say the hardness of inverting a one-way function where the preimage size could be super-polynomial.

In this section we show that this is somewhat inherent. We prove that for search assumptions, such a transformation cannot exist as long as the resulting reduction \mathcal{R}' is explicit in the assumptions P, Q . In particular, it may obtain as input the code of the algorithms describing P, Q , but does not get any implicit non-uniform advice regarding these assumptions. Indeed, the transformation in Section 7 (from classical non-adaptive reductions to quantum ones) as well as the Persistence Theorem 4.1 on which it relies, the resulting quantum reduction \mathcal{R}' is in fact black-box in the assumptions P, Q , and in particular explicit.

Definition 8.1 (Assumption Pair Collection). *An assumption pair collection \mathcal{PQ} consists of pairs of assumptions (P, Q) , each given by its corresponding (possibly non-uniform) algorithms (G_P, V_P, c_P) and (G_Q, V_Q, c_Q) .*

Definition 8.2 (Explicit Reduction). *An explicit quantum reduction for assumption pair collection \mathcal{PQ} is an efficient algorithm \mathcal{R} with the following guarantee. For any $(P, Q) \in (\mathcal{P}, \mathcal{Q})$ and any quantum solver $\mathcal{B}_P = (B_P, \text{state}_0)$ for P with one-shot advantage ε and running time T , let $\text{state}'_0 = (\text{state}_0, (P, Q), B_P, 1^{1/\varepsilon}, 1^T)$. Then $\mathcal{B}_Q = (\mathcal{R}, \text{state}'_0)$ is a solver for Q with one-shot advantage $\text{poly}(\varepsilon, T^{-1}, \lambda^{-1})$ and running-time $\text{poly}(T, \varepsilon^{-1}, \lambda)$.*

We say that the reduction is strongly explicit, instead of being given the explicit description of (P, Q) as part of its input, it is given oracle access to its corresponding algorithms.

Note that in the above definition state'_0 is formally a sequence

$$\text{state}'_{0,\lambda} = (\text{state}_{0,\lambda}, (P, Q)_\lambda, B_{P,\lambda}, 1^{1/\varepsilon(\lambda)}, 1^{T(\lambda)}) ,$$

where $(P, Q)_\lambda$ consist of their corresponding algorithms (possibly along with their corresponding non-uniform advice) restricted to security parameter λ (w.l.o.g circuits).

Restating our result from Section 7, we proved that for any pair collection \mathcal{PQ} , if for any $(P, Q) \in \mathcal{P}, \mathcal{Q}$, P has verifiably-polynomial image, and there exists

a classical non-adaptive black-box reduction $\mathcal{R}_{P,Q}$ from solving Q to solving P , then there also exists a strongly explicit quantum reduction \mathcal{R}' for \mathcal{PQ} . We prove that if P does not have a verifiably-polynomial image this may not be the case.

Theorem 8.3. *There exists an assumption pair collection \mathcal{PQ} , such that for any $(P, Q) \in \mathcal{P}, \mathcal{Q}$, there exists a classical non-adaptive black-box reduction $\mathcal{R}_{P,Q}$ from solving Q to solving P , but there is no strongly explicit reduction \mathcal{R}' for \mathcal{PQ} . Assuming also post-quantum indistinguishability obfuscation, there also does not exist an explicit reduction \mathcal{R}' .*

The theorem is proven in the full version of the paper [3].

Acknowledgements

Nir Bitansky is a member of the checkpoint institute of information security and was supported by the Israel Science Foundation (Grants No. 484/18 and 2137/19, by Len Blavatnik and the Blavatnik Family Foundation, by the Blavatnik Interdisciplinary Cyber Research Center at Tel Aviv University, and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482). Zvika Brakerski was supported by the Israel Science Foundation (Grant No. 3426/21), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701). Yael Tauman Kalai's work supported by DARPA under Agreement No. HR001120200. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

References

1. Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In Helmut Alt and Afonso Ferreira, editors, *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science, Antibes - Juan les Pins, France, March 14-16, 2002, Proceedings*, volume 2285 of *Lecture Notes in Computer Science*, pages 323–334. Springer, 2002.
2. Shalev Ben-David and Or Sattath. Quantum tokens for digital signatures. *CoRR*, abs/1609.09047, 2016.
3. Nir Bitansky, Zvika Brakerski, and Yael Tauman Kalai. Constructive post-quantum reductions. *CoRR*, abs/2203.02314, 2022. <https://arxiv.org/abs/2203.02314>.
4. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, pages 41–69, 2011.
5. Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 320–331, 2018.

6. Zvika Brakerski and Oded Goldreich. From absolute distinguishability to positive distinguishability. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, volume 6650 of *Lecture Notes in Computer Science*, pages 141–155. Springer, 2011.
7. Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments. *CoRR*, abs/2103.08140, 2021.
8. Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584. Springer, 2021.
9. Ivan Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner. Cryptography in the bounded-quantum-storage model. *SIAM J. Comput.*, 37(6):1865–1890, 2008.
10. Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round fiat-shamir and more. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, pages 602–631, 2020.
11. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*, pages 356–383, 2019.
12. Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In David S. Johnson, editor, *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 25–32. ACM, 1989.
13. Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In Omer Reingold, editor, *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*, pages 202–219. Springer, 2009.
14. Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. Indcca-secure key encapsulation mechanism in the quantum random oracle model, revisited. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, pages 96–125, 2018.
15. Shuichi Katsumata, Shota Yamada, and Takashi Yamakawa. Tighter security proofs for GPV-IBE in the quantum random oracle model. In *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part II*, pages 253–282, 2018.
16. Juliane Krämer and Patrick Struck. Encryption schemes using random oracles: From classical to post-quantum security. In *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020, Paris, France, April 15-17, 2020, Proceedings*, pages 539–558, 2020.

17. Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*, pages 326–355, 2019.
18. Moni Naor. On cryptographic assumptions and challenges. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer, 2003.
19. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016.
20. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.
21. Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the fujisaki-okamoto and OAEP transforms. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*, pages 192–216, 2016.
22. Takashi Yamakawa and Mark Zhandry. Classical vs quantum random oracles. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 568–597. Springer, 2021.
23. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, pages 758–775, 2012.
24. Mark Zhandry. How to record quantum queries, and applications to quantum indifferenciability. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*, pages 239–268, 2019.