Quantum One-Time Memories from Stateless Hardware

Anne Broadbent* Sevag Gharibian† Hong-Sheng Zhou‡

November 4, 2015

A central tenet of theoretical cryptography is the study of the minimal assumptions required to implement a given cryptographic primitive. One such primitive is the one-time memory (OTM), introduced by Goldwasser, Kalai, and Rothblum [CRYPTO 2008], which is a classical functionality modeled after a non-interactive 1-out-of-2 oblivious transfer, and which is complete for one-time classical and quantum programs. It is known that secure OTMs do not exist in the standard model in both the classical and quantum settings. Here, we show how to use quantum information, together with the assumption of stateless (i.e., reusable) hardware tokens, to build statistically secure OTMs. This is in sharp contrast with the classical case, where stateless hardware tokens alone cannot yield OTMs. In addition, our scheme is technologically simple. We prove security in the quantum universal composability framework, employing semi-definite programming results of Molina, Vidick and Watrous [TQC 2013] and combinatorial techniques of Pastawski *et al.* [Proc. Natl. Acad. Sci. 2012].

1 Introduction

The study of theoretical cryptography is centered around the question of building cryptographic primitives secure against adversarial attacks. In order to allow a broader set of such primitives to be implemented, one often considers restricting the power of the adversary. For example, one can limit the *computing* power of adversaries to be polynomial bounded [Yao82, BM82], restrict the *storage* of adversaries to be bounded or noisy [Mau92, CM97, DFSS05], or make *trusted setups* available to honest players [Kil88, BFM88, Can01, CLOS02, IPS08, PR08, LPV09, MPR09, MPR10, MR11, KMQ11, KMPS14], to name a few. One well-known trusted setup is *tamper-proof hardware* [Kat07, GKR08], which is assumed to provide a specific input-output functionality, and which can only be accessed in a "black box" fashion. The hardware can maintain a state (i.e., is *stateful*) and possibly carry out complex functionality, but presumably may be difficult or expensive to implement or manufacture. This leads to an interesting research direction: Build cryptography primitives using the *simplest* (and hence easiest and cheapest to manufacture) hardware.

In this respect, two distinct simplified notions of hardware have captured considerable interest. The first is the notion of a *one-time memory (OTM)* [GKR08], which is arguably the simplest possible notion of *stateful* hardware. An OTM, modeled after a non-interactive 1-out-of-2 oblivious transfer, behaves as follows: first, a player (called the *sender*) embeds two values s_0 and s_1 into the OTM, and then gives the OTM to another player (called the *receiver*). The receiver can now read his choice of precisely one of s_0 or s_1; after this "use" of the OTM, however, the unread bit is lost forever. Interestingly, OTMs are complete for implementing *one-time* use programs (OTPs): given access to OTMs, one can implement statistically secure OTPs for any

*Department of Mathematics and Statistics, University of Ottawa, Ontario, Canada. Email: abroadbe@uottawa.ca.
†Department of Computer Science, Virginia Commonwealth University, Virginia, USA. Email: sgharibian@vcu.edu.
‡Department of Computer Science, Virginia Commonwealth University, Virginia, USA. Email: hszhou@vcu.edu.

efficiently computable program in the universal composability (UC) framework [GIS+10]. (OTPs, in turn, have applications in software protection and one-time proofs [GKR08].) In the quantum UC model, OTMs enable *quantum* one-time programs [BGS13]. (This situation is analogous to the case of *oblivious transfer* being complete for two-party secure function evaluation [Kil88, IPS08].) Unfortunately, OTMs are inherently *stateful*, and thus represent a very strong cryptographic assumption — any physical implementation of such a device must somehow maintain internal knowledge between activations, i.e. it must completely "self-destruct" after a single use.

This brings us to a second important simplified notion of hardware known as a *stateless* token [CGS08], which keeps no record of previous interactions. On the positive side, such hardware is presumably easier to implement. On the negative side, an adversary can run an experiment with stateless hardware as many times as desired, and each time the hardware is essentially "reset". (Despite this, stateless hardware has been useful in achieving *computationally secure* multi-party computation [CGS08, GIS+10, CKS+14], and *statistically secure* commitments [DS13].) It thus seems impossible for stateless tokens to be helpful in implementing any sort of "self-destruct" mechanism. Indeed, classically stateful tokens are trivially more powerful than stateless ones, as observed in, *e.g.* [GIS+10]. This raises the question:

> *Can* quantum *information, together with a classical stateless token, be used to simulate "self destruction" of a hardware token?*

In particular, a natural question along these lines is whether quantum information can help implement an OTM. Unfortunately, it is known that quantum information *alone* cannot implement an OTM (or, more generally, any one-time program) [BGS13]; see also Section 4 below. We thus ask the question: What are the minimal cryptographic assumptions required in a quantum world to implement an OTM?

Contributions and summary of techniques. Our main contribution is to show that, in the quantum model, OTMs can be constructed from stateless hardware tokens. We thus show a quantum reduction from stateful to stateless hardware. This is in sharp contrast with the classical case, in which such a reduction is known to be trivially impossible.

CONSTRUCTION. The construction is inspired by Wiesner's idea for *conjugate coding* [Wie83]: the quantum portion of the protocols consists in n quantum states chosen uniformly at random from $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$ (note this encoding is independent of the classical bits of the OTM functionality). We then couple this n-qubit quantum state, $|\psi\rangle$, with a *classical* stateless hardware token, which takes as inputs a choice bit b, together with an n-bit string y. If $b = 0$, the hardware token verifies that the bits of y that correspond to *rectilinear* ($|0\rangle$ or $|1\rangle$, i.e. Z basis) encoded qubits of $|\psi\rangle$ are consistent with the measurement of $|\psi\rangle$ in the computational basis, in which case the bit s_0 is returned. If $b = 1$, the hardware token verifies that the bits of y that correspond to *diagonal* ($|+\rangle$ or $|-\rangle$, i.e. X basis) encoded qubits of $|\psi\rangle$ are consistent with the measurement of $|\psi\rangle$ in the diagonal basis, in which case the bit s_1 is returned.

ASSUMPTIONS. First, crucially, the hardware token is specified to accept *classical* input only (i.e. it cannot be queried in superposition). Although this may seem a strong assumption, in Section 4.1 we show that any token which can be queried in superposition cannot be used to construct a secure OTM (with respect to our setting in which the adversary is allowed to apply arbitrary quantum operations). Similar classical-input hardware has previously been considered in, e.g., [Unr13, BGS13]. Second, we assume in this work that the sender is honest (i.e. $|\psi\rangle$ and the hardware tokens are honestly produced according to the specified protocol).

SECURITY AND INTUITION. The intuition underlying security for our scheme is clear: in order for a receiver to extract a bit s_b as encoded in the OTM, she must perform a complete measurement of the qubits of $|\psi\rangle$ in order to obtain a classical key for s_b (since, otherwise, she would likely fail the test as imposed by the hardware token). But such a measurement would invalidate the receiver's chance of extracting the bit s_{1-b}! This is exactly the "self-destruct"-like property we require in order to implement an OTM. This intuitive

notion of security was already present in Wiesner's proposal for quantum money[1] [Wie83], and is often given a physical explanation in terms of the no-cloning theorem [WZ82], or the Heisenberg uncertainty relation [Hei27].

Formally, we show statistical or information-theoretic security in the quantum *Universal Composability* (UC) framework [Unr10], which allows us to make strong security claims in terms of the *composability* of our protocol within others. Such a security proof requires the construction of a simulator, such that for any "quantum environment" wishing to interact with the OTM, the environment statistically cannot tell whether it is interacting with the *ideal* OTM functionality or the *real* OTM instance provided by our scheme. The security of this simulator requires a statement of the following form: Given access to a (randomly chosen) "quantum key" $|\psi_k\rangle$ and corresponding stateless token V_k, it is highly unlikely for an adversary to successfully extract keys for *both* the secret bits s_0 and s_1 held by V_k, *even if* the adversary is allowed to interact with V_k multiple times. The proof of this statement proceeds in two parts. First, in Appendix B.2, we show that the probability of an adversary succeeding *with* interaction with V_k is polynomially related to that of succeeding *without* interaction with V_k. This part is shown using a combinatorial technique of Pastawski *et al.* [PYJ+12] from the setting of quantum money. Second, in Appendix B.3, we show that the probability of succeeding without interaction is exponentially small in n, the number of qubits in the quantum key $|\psi_k\rangle$. This is shown using semidefinite programming results of Molina, Vidick and Watrous [MVW13], also from the context of quantum money. Both of these proof steps are sketched in Section 3.3.

Summarizing, we show the following.

Main Theorem (informal). *There exists a protocol Π, which together with a classical stateless token and the ability to randomly prepare single qubits in one of four pure states, implements the OTM functionality with statistical security in the UC framework against a corrupted receiver of the OTM.*

FURTHER IMPLICATIONS. When combined with prior results, by using the *quantum lifting* technique of Unruh [Unr10], our theorem above implies: quantum one-time *classical programs* [GIS+10], and quantum one-time *quantum programs* [BGS13] (in both cases, with security against a corrupt receiver only).

Related work. Our work contributes to the growing list of functionalities achievable with quantum information, yet unachievable classically. This includes: unconditionally secure key expansion [BB84], physically uncloneable money [Wie83, MVW13, PYJ+12], a reduction from oblivious transfer to bit commitment [BBCS92, DFL+09] and to other primitives such as "cut-and choose" functionality [FKS+13], and revocable time-release quantum encryption [Unr14]. Importantly, these protocols all make use of the technique of conjugate coding [Wie83], which is also an important technique used in protocols for OT in the bounded quantum storage and noisy quantum storage models [DFSS05, WST08] (see [BS15] for a survey).

A number of proof techniques have been developed in the context of conjugate coding, including entropic uncertainty relations [WW10]. In the context of QKD, another successful technique is the use of de Finetti reductions [Ren08] (which exploit the symmetry of the scheme in order to simplify the analysis). Recently, semidefinite programming approaches have been applied to analyze security of conjugate coding [MVW13]. This is the approach that we adopt for the "non-interactive" portion of our proof (Section B.3). Reference [PYJ+12] has also made use of Gavinsky's [Gav12] quantum retrieval games framework.

Continuing with proof techniques, somewhat similar to [PYJ+12], Aaronson and Christiano [AC12] have studied quantum money schemes in which one interacts with a verifier. They introduce an "inner product adversary method" to lower bound the number of queries required to break their scheme. (In contrast, we use the combinatorial technique of [PYJ+12].)

We remark that References [PYJ+12] and [MVW13] have studied schemes based on conjugate coding similar to ours, but in the context of quantum money. In contrast to our setting, the schemes of [PYJ+12]

[1] Intuitively, quantum money aims to construct currency which is impossible to counterfeit by the laws of quantum mechanics.

and [MVW13] (for example) involve dynamically chosen random challenges from a verifier to the holder of a "quantum banknote", whereas in our work here the "challenges" are fixed (i.e. measure all qubits in the Z or X basis to obtain secret bit s_0 or s_1, respectively), and the verifier is replaced by a stateless token.

Also, we note that prior work has achieved oblivious transfer using quantum information, together with some assumption (*e.g.* bit commitment [BBCS92] or bounded quantum storage [DFSS05]). These protocols typically use an interaction phase similar to the "commit-and-open" protocol of [BBCS92]; because we are working in the non-interactive setting, these techniques appear to be inapplicable.

Finally, we mention related work by Liu [Liu14a, Liu14b, Liu15], which establishes stand-alone secure OTMs using quantum information. In contrast to our setting, in which an adversary's allowed quantum gate set is unrestricted, Liu's results are set in the *isolated-qubit model*, which assumes that an adversary can perform only single-qubit operations (thus, no entangling gates are permitted). We remark that, the security notion of OTMs by Liu is much weaker than the simulation-based notion that is studied in this paper, and it is unclear whether this type of OTM is composable; the main goal there is to show feasibility without using any trusted setup assumptions.

Significance. Our results show a strong separation between the classical and quantum settings, since classically, stateless tokens cannot be used to securely implement OTMs. To the best of our knowledge, our work is the first to combine conjugate coding with *stateless* hardware tokens. Moreover, while our protocol shares similarities with prior work in the setting of quantum money, building OTMs appears to be a new focus here [2].

Our protocol has a simple implementation, fitting into the single-qubit prepare-and-measure paradigm (in fact, one needs only the ability to prepares states $|0\rangle, |1\rangle, |+\rangle, |-\rangle$). Thus, our scheme is in principle amenable to experimental implementations (the quantum portion of our protocol, could, in principle, be implemented using current hardware for quantum key distribution [BB84]). In addition, from a theoretical cryptographic perspective, our protocol is attractive in that its implementation requires an assumption of a stateless hardware token, which is conceivably easier and cheaper to mass produce than a stateful token.

In terms of security guarantees, we allow *arbitrary* operations on behalf of a malicious quantum receiver in our protocol (i.e. all operations allowed by quantum mechanics), with the adversary only restricted in that the stateless token is assumed only usable as a black box. The security we obtain is statistical, with the only computational assumption being on the number of *queries* made to the token (i.e. in order to break the scheme, an adversary must make an *exponential* number of calls to the token, and moreover this is sufficient, as clearly one can try all possible keys via brute force in this time). Finally, our proofs are rigorous statements in the quantum UC framework, meaning our protocol can be easily composed with others proved secure in this framework (e.g. combining our results with [BGS13]'s protocol immediately yields UC-secure quantum OTPs against a dishonest receiver).

Finally, we close by remarking that our scheme is "tight" with respect to two impossibility results. First, the assumption that the token be queried only in the computational basis cannot be relaxed: Section 4.1 shows that if the token can be queried in superposition, then an adversary in our setting can easily break any OTM scheme. Second, our scheme has the property that corresponding to each secret bit s_i held by the token, there are exponentially many valid keys one can input to the token to extract s_i. In Section 4.2, we show that for any "measure-and-access" OTM (i.e. an OTM in which one measures a given quantum key and uses the classical measurement result to access a token to extract data, of which our protocol is an example), a polynomial number of keys implies the ability to break the scheme with inverse polynomial probability (more generally, Δ keys allows probability at least $1/\Delta^2$ of breaking the scheme).

Organization of the paper. The remainder of the paper is organized as follows. We begin in Section 2 with preliminaries, including the ideal functionalities for an OTM and stateless token. In Section 3, we give our

[2] We remark, however, that a reminiscent concept of single usage of quantum "tickets" in the context of quantum money is very briefly mentioned in Appendix S.4.1 of [PYJ+12].

construction for an OTM based on a stateless hardware token; the proof ideas for security are also provided. In Section 4, we discuss "tightness" of our construction by showing two impossibility results for "relaxations" of our scheme. Due to space constraints, we include the description of classical UC and quantum UC in Appendix A. Appendix B gives a formal proof upper bounding the maximum probability with which two accepting keys for a token V can be extracted from a single quantum key $|\psi\rangle$. (These results are used to finish the security proof in Section 3.) In addition, the security proof for a lemma in Section 4 can be found in Appendix C.

2 Preliminaries

Notation. We say two binary distributions $\mathbf{X}$ and $\mathbf{Y}$ are *indistinguishable*, denoted $\mathbf{X} \approx \mathbf{Y}$, if it holds that $|\Pr(X_n = 1) - \Pr(Y_n = 1)| \leq \mathrm{negl}(n)$. We define single-qubit $|0\rangle_+ = |0\rangle$ and $|1\rangle_+ = |1\rangle$, so that $\{|0\rangle_+, |1\rangle_+\}$ form the *rectilinear basis*. We also define $|0\rangle_\times = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle_\times = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, so that $\{|0\rangle_\times, |1\rangle_\times\}$ form the *diagonal basis*. For strings $x = x_1, x_2, \ldots x_n \in \{0,1\}^n$ and $\theta = \theta_1, \theta_2, \ldots, \theta_n \in \{+, \times\}^n$, we define $|x\rangle_\theta = \bigotimes_{i=1}^n |x_i\rangle_{\theta_i}$. Finally, H denotes the standard 2×2 Hadamard gate $H = (1\ 1;\ 1\ -1)/\sqrt{2}$ in quantum information.

Quantum universal composition (UC) framework. We consider simulation-based security in this paper. In particular, we prove the security of our construction in the quantum universal composition (UC) framework [Unr10]. Please see Appendix A for a brief description of the classical UC [Can01] and the quantum UC [Unr10]. In the next two paragraphs, we introduce two relevant ideal functionalities of one-time memory and of stateless hardware token.

One-time memory (OTM). The one-time memory (OTM) functionality $\mathcal{F}_{\mathtt{OTM}}$ involves two parties, the sender and the receiver, and consists of two phases, "Create" and "Execute". Please see Functionality 1 below for details; for the sake of simplicity, we have omitted the session/party identifiers as they should be implicitly clear from the context. We sometimes refer to this functionality $\mathcal{F}_{\mathtt{OTM}}$ as an *OTM token*.

Functionality 1 Ideal functionality $\mathcal{F}_{\mathtt{OTM}}$.

1. **Create:** Upon input (s_0, s_1) from the sender, with $s_0, s_1 \in \{0,1\}$, send create to the receiver and store (s_0, s_1).
2. **Execute:** Upon input $b \in \{0,1\}$ from the receiver, send s_b to receiver. Delete any trace of this instance.

Stateless hardware. The original work of Katz [Kat07] introduces the ideal functionality $\mathcal{F}_{\mathtt{wrap}}$ to model stateful tokens in the UC-framework. In the ideal model, a party that wants to create a token, sends the Turing machine to $\mathcal{F}_{\mathtt{wrap}}$. $\mathcal{F}_{\mathtt{wrap}}$ will then run the machine (keeping the state), when the designed party will ask for it. The same functionality can be adapted to model stateless tokens. It is sufficient that the functionality does not keep the state between two executions. A simplified version of the $\mathcal{F}_{\mathtt{wrap}}$ functionality as shown in [CGS08] (that is very similar to the $\mathcal{F}_{\mathtt{wrap}}$ of [Kat07]) is described below. Note that, again for the sake of simplicity, we have omitted the session/party identifiers as they should be implicitly clear from the context.

Although the environment and adversary are unbounded, we specify that stateless hardware can be queried only a polynomial number of times. This is necessary, since otherwise the hardware token model is vacuous (with unbounded queries, the entire input-output behavior of stateless hardware can be deduced, hence there is nothing left to hide).

Functionality 2 Ideal functionality $\mathcal{F}_{\texttt{wrap}}$.

The functionality is parameterized by a polynomial $p(\cdot)$, and an implicit security parameter n

1. **Create:** Upon input (create, M) from the sender, where where M is a Turing machine, send create to the receiver and store M.
2. **Execute:** Upon input (run, msg) from the receiver, execute $M(msg)$ for at most $p(n)$ steps, and let out be the response. Let $out := \perp$ if M does not halt in $p(n)$ steps. Send out to the receiver.

3 Feasibility of Quantum OTMs using Stateless Hardware

In this section, we present a *quantum* construction for one-time memories by using stateless hardware (Section 3.1). We also state our main theorem (Theorem 3.1). In Section 3.2, we describe the Simulator and prove Theorem 3.1 using the technical results of Appendix B. The intuition and techniques behind the proofs in Appendix B are sketched in Section 3.3.

3.1 Construction

We now present the OTM protocol Π in the $\mathcal{F}_{\texttt{wrap}}$ hybrid model, between a sender P_{s} and a receiver P_{r}. Here the security parameter is n.

- Upon receiving input (s_0, s_1) from the environment where $s_0, s_1 \in \{0,1\}$, sender P_{s} operates as follows:
  - The sender chooses random strings $x \in_R \{0,1\}^n$ and $\theta \in_R \{+,\times\}^n$, and prepares $|x\rangle_\theta$. Then the sender, based on tuple (s_0, s_1, x, θ), prepares the program M as in **Program 1**.

Program 1 Program for hardware token

Hardcoded values: $s_0, s_1 \in \{0,1\}$, $x \in \{0,1\}^n$, and $\theta \in \{+,\times\}^n$
Inputs: $y \in \{0,1\}^n$ and $b \in \{0,1\}$, where y is a claimed measured value for the quantum register, and b the evaluator's choice bit

1. If $b = 0$, check that the $\theta = +$ positions return the correct bits in y according to x. If Accept, output s_0. Otherwise output $\perp$.
2. If $b = 1$, check that the $\theta = \times$ positions return the correct bits in y according to x. If Accept, output s_1. Otherwise output $\perp$.

  - The sender sends $|x\rangle_\theta$ to the receiver.
  - The sender sends (create, M) to functionality $\mathcal{F}_{\texttt{wrap}}$, and the functionality sends create to notify the receiver.
- The receiver P_{r} operates as follows:
  Upon input b from the environment, and $|x\rangle_\theta$ from the receiver, and create notification from $\mathcal{F}_{\texttt{wrap}}$,
  - If $b = 0$, measure $|x\rangle_\theta$ in the computational basis to get string y and input $(\mathsf{run}, (y, b))$ into $\mathcal{F}_{\texttt{wrap}}$.
  - If $b = 1$, apply $\mathrm{H}^{\otimes n}$ to $|x\rangle_\theta$, then measure in the computational basis to get string y and input $(\mathsf{run}, (y, b))$ into $\mathcal{F}_{\texttt{wrap}}$.

  Return the output of $\mathcal{F}_{\texttt{wrap}}$ to the environment.
  It is easy to see that the output of $\mathcal{F}_{\texttt{wrap}}$ is s_b for both $b = 0$ and $b = 1$.

Note again that the hardware token, as defined in **Program 1**, accepts only classical input (i.e. it cannot be queried in superposition). As mentioned earlier, relaxing this assumption yields impossibility of a secure OTM implementation, as shown in Section 4. Our main theorem is now stated as follows.

Theorem 3.1. *Construction* Π *above quantum-UC-realizes* $\mathcal{F}_{\texttt{OTM}}$ *in the* $\mathcal{F}_{\texttt{wrap}}$ *hybrid model with statistical security against a corrupted receiver.*

3.2 Proof of Theorem 3.1

To prove Theorem 3.1, we must construct and analyze an appropriate simulator, which we now proceed to do.

3.2.1 The simulator

In order to prove Theorem 3.1, for any unbounded adversary $\mathcal{A}$ who corrupts the receiver, we need to build a simulator $\mathcal{S}$ (having access to the OTM functionality $\mathcal{F}_{\texttt{OTM}}$), such that for any unbounded environment $\mathcal{Z}$, the executions in the real model and that in simulation are statistically indistinguishable. Our simulator $\mathcal{S}$ is given below:

- The simulator emulates an internal copy of the adversary $\mathcal{A}$ who corrupts the receiver. The simulator emulates the communication between $\mathcal{A}$ and the external environment $\mathcal{Z}$ by forwarding the communication messages between $\mathcal{A}$ and $\mathcal{Z}$.
- The simulator $\mathcal{S}$ needs to emulate the whole view for the adversary $\mathcal{A}$. First, the simulator picks dummy inputs $\tilde{s}_0 = 0$ and $\tilde{s}_1 = 0$, and randomly chooses $x \in \{0,1\}^n$, and $\theta \in \{+,\times\}^n$, and generates program $\tilde{M}$. Then the simulator plays the role of the sender to send $|x\rangle_\theta$ to the adversary $\mathcal{A}$ (who controls the corrupted receiver). The simulator also emulates $\mathcal{F}_{\texttt{wrap}}$ to notify $\mathcal{A}$ by sending create to indicate that the hardware is ready for queries.
- For each query $(\mathsf{run}, (b, y))$ to $\mathcal{F}_{\texttt{wrap}}$ from the adversary $\mathcal{A}$, the simulator evaluates program $\tilde{M}$ (that is created based on $\tilde{s}_0, \tilde{s}_1, x, \theta$) as in the construction, and then acts as follows:
  1. If this is a rejecting input, output $\perp$.
  2. If this is the first accepting input, call the external $\mathcal{F}_{\texttt{OTM}}$ with input b, and learn the output s_b from $\mathcal{F}_{\texttt{OTM}}$. Output s_b.
  3. If this is a subsequent accepting input, output s_b (as above).

3.2.2 Analysis

We now show that the simulation and the real model execution are statistically indistinguishable. There are two cases in an execution of the simulation which we must consider:

- *Case 1: In all its queries to* $\mathcal{F}_{\texttt{wrap}}$*, the accepting inputs of* $\mathcal{A}$ *have the same choice bit* b*.* In this case, the simulation is perfectly indistinguishable.
- *Case 2: In its queries to* $\mathcal{F}_{\texttt{wrap}}$*,* $\mathcal{A}$ *produces accepting inputs for both* $b = 0$ *and* $b = 1$*.* In this case, the simulation fails (the environment can distinguish the real model from the ideal model), since the simulator is only able to retrieve a single bit from the external OTM functionality $\mathcal{F}_{\texttt{OTM}}$ (either corresponding to $b = 0$ or $b = 1$).

Thus, whereas in Case 1 the simulator behaves perfectly, in Case 2 it is in trouble. Fortunately, in Theorem B.1 we show that the probability that Case 2 occurs is exponentially small in n, the number of qubits comprising $|x\rangle_\theta$. Specifically, we show that for an arbitrary m-query strategy (i.e. any quantum strategy allowed by quantum mechanics, whether efficiently implementable or not, which queries the token at most m times), the probability of Case 2 occurring is at most $2\binom{m}{2}(\frac{1}{2} + \frac{1}{2\sqrt{2}})^n$. Since m is assumed to be polynomially-bounded, this expression is exponentially small in n. This concludes the proof.

3.3 Security analysis for the token: Intuition

Our simulation proof showing statistical security of our Quantum OTM construction of Section 3.1 relies crucially on Theorem B.1, which can informally be stated as follows: given a single copy of n-qubit state $|x\rangle_\theta$, the probability that an unbounded adversary is able to extract both bits s_0 and s_1 is exponentially small (in n), *even if* the adversary is allowed to query the token a polynomial number of times.

We first give a formal statement of this result, followed by the intuition behind its proof (the full proof is given in Appendix B). In our discussion, we shall switch to a quantum information view for the stateless token (as opposed to the cryptographic view used in Sections 2 and 3), since this is the setting in which we show Theorem B.1. Specifically, we model a stateless token V as a map from a query register $\mathcal{Q}$ to a set of three strings: 000 for "reject", $10s_0$ for "accept and return s_0", and $01s_1$ for "accept and return s_1". Furthermore, we refer to the states $|x\rangle_\theta$ (i.e. the "quantum keys") from conjugate coding as *BB84* states, for the important role that they play in the quantum key distribution protocol [BB84]. For convenience, we refer to a BB84 key $|x\rangle_\theta$ as state $|\psi_k\rangle$, with corresponding token V_k (since the verification program of V_k depends on k). Finally, for any finite dimensional complex Hilbert space $\mathcal{X}$, we use $\mathcal{D}(\mathcal{X})$ to denote the set of density operators acting on $\mathcal{X}$ (see Section B.1 for further quantum information notation).

Theorem 3.2. *Let $\mathcal{X}, \mathcal{Q}$ be finite dimensional Hilbert spaces, and let $S = \{|\psi_k\rangle, V_k\}$ denote the ensemble of BB84 states $|\psi_k\rangle \in \mathcal{X}$ and corresponding oracles $V_k : \mathcal{Q} \mapsto \{0,1\}^3$ used in Section 3.1. Then, for any interactive strategy Φ (formally, a trace-preserving, completely positive (TPCP) map $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q} \otimes \mathcal{Q})$) which queries V_k for m times,*

$$\frac{1}{|S|} \Pr[V_k \otimes V_k \text{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \text{ outputs } 10s_001s_1] \leq 2\binom{m}{2}\left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n. \quad (1)$$

Intuitively, the adversary's strategy Φ attempts to map the input BB84 key $|\psi_k\rangle$ to a pair of classical secret keys, one for s_0 and one for s_1, stored in register $\mathcal{Q} \otimes \mathcal{Q}$, which is then checked by two parallel runs of V_k, i.e. by $V_k \otimes V_k$. Note that the adversary's strategy is modeled by an arbitrary TPCP map Φ, meaning any strategy allowable by the laws of quantum mechanics (whether efficiently implementable or not) is permitted.

To prove this statement, we proceed in two steps: First, we show that the ability to adaptively query the hardware token m times can improve the success probability of extracting a pair of keys from $|\psi_k\rangle$ for s_0 and s_1 by at most a polynomial factor in m (in comparison to the setting in which no queries to the token are allowed). Second, we show that if the adversary cannot query the token, then it can extract keys from $|\psi_k\rangle$ for both s_0 and s_1 with probability at most inverse exponential in n, the number of key bits.

Step 1: Eliminating interaction with the token. To reduce the interactive setting to the non-interactive one, we apply a combinatorial technique of Reference [PYJ+12]. In [PYJ+12], this approach was used to analyze two schemes for quantum money in which a verification oracle is queried (possibly multiple times), with each query outputting a single bit (i.e. accept or reject). Our setting, however, requires a classical oracle which outputs *multiple* bits, i.e. accept/reject, along with possibly multiple *secret* bits for the accept case. For this part of the proof, we observe that the technique of [PYJ+12] can be easily generalized to any setting of the following form: One is given (according to some distribution) a state $|\psi_k\rangle$ and classical oracle V_k, and asked to measure $|\psi_k\rangle$ to determine some classical *key* to input to V_k in order to obtain some desired output from V_k. Unlike [PYJ+12], the oracles V_k can output strings of any length — the only restriction is that all V_k must share the same set of output strings (i.e. the set of output strings is independent of k, so that intuitively, no information about keys is leaked by the output of V_k).

To make this formal, we introduce the notion of a *fixed-output ensemble of states and oracles*. Specifically, let $\mathcal{X}, \mathcal{Q}$ be finite dimensional Hilbert spaces. A fixed-output ensemble $\{|\psi_k\rangle, V_k\}$ of states $|\psi_k\rangle \in \mathcal{X}$ and oracles V_k is one with the following properties. Each V_k accepts a (mixed) state $\rho \in \mathcal{D}(\mathcal{Q})$ as input, measures

ρ according to some s-outcome projective measurement $\{\Pi_i^k\}_{i=1}^s$, and outputs t bits according to these rules: To denote "reject", V_k outputs 0^t. Otherwise, V_k outputs a non-zero t-bit string y; denote the set of such "good" outputs y as $G \subseteq \{0,1\}^t$. Crucially, each V_k may have a distinct set of measurement operators $\{\Pi_i^k\}_{i=1}^s$, but all V_k share the *same* set of output strings $G \cup \{0^t\}$ (hence the name *fixed-output*). The formal statement we now show is below, and its proof follows that of Theorem 9 of [PYJ+12] closely.

Lemma 3.3 (see also Theorem 9 of [PYJ+12])**.** *Let $\mathcal{X}, \mathcal{Q}$ be finite dimensional Hilbert spaces, and let $S = \{|\psi_k\rangle, V_k\}$ be a fixed-output ensemble of states of $\mathcal{X}$ and $\mathcal{Q}$. Fix any distinct $y_1, y_2 \in G$, and suppose that for any trace-preserving, completely positive (TPCP) map $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q} \otimes \mathcal{Q})$,*

$$\frac{1}{|S|} \Pr[V_k \otimes V_k \text{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \text{ outputs } y_1 \neq y_2 \in G] \leq p \tag{2}$$

where $0 \leq p \leq 1$. Then, for any strategy Γ which, given $|\psi_k\rangle$ with probability $1/|S|$, queries V_k m times, it holds that the probability that at least two queries to V_k return distinct strings in G is at most $\binom{m}{2}|G|(|G|-1)p$.

Thus, querying the token V m times can increase the probability of success by at most a $\mathrm{poly}(m)$ factor (note that in our setting, $|G| = 2$).

Intuitively, the proof of Lemma 3.3 proceeds as follows. One first observes that the action of the adversary, Γ, can be characterized by a sequence of $m+1$ *non-interactive* TPCP maps Γ_i, where map i is determined by query answers q_1 through q_{i-1}. In general, there are hence 2^m such sequences of non-interactive maps Γ_i to consider. However, we are only interested in sequences of queries which yield at least two accepting queries extracting both s_0 and s_1. This intuitively reduces the number of sequences of maps which must be explicitly considered to $\binom{m}{2}$ (i.e. choose the two query positions in which the *first two* successful queries were made), which is polynomial in m. But by the assumption of Lemma 3.3, any fixed non-interactive map succeeds at extracting both s_0 and s_1 with probability at most p. Roughly, adding the success probabilities of all $\binom{m}{2}$ non-interactive sequences of maps we need to consider now yields the desired success probability of at most $\mathrm{poly}(m)$ times p.

Step 2: Bounding cheating probabilities in the non-interactive setting. By Step 1, we are reduced to the question: Given BB84 state $|\psi_k\rangle$ and no query access to token V_k, what is the maximum probability with which accepting keys for both s_0 and s_1 can be extracted? To analyze this, we apply the semidefinite programming based results of Molina, Vidick, and Watrous [MVW13], which studied similar BB84-based schemes in the context of quantum money. The lemma below is a special case of Lemma 5 of [MVW13] (our re-statement below is formulated with respect to the context of this paper), whose proof is essentially identical to that of Lemma 5 of [MVW13], save for some minor modifications for our setting. (The full proof is given in Appendix B.3.

Lemma 3.4 (see Lemma 5 of [MVW13])**.** *Let $\mathcal{X} = (\mathbb{C}^2)^{\otimes n}$ and $\mathcal{Q} = (\mathbb{C}^2)^{\otimes(n+1)}$, and let $S = \{|\psi_k\rangle, V_k\}$ be the fixed-output ensemble of states of $\mathcal{X}$ and $\mathcal{Q}$ corresponding to the construction of Section 3.1. Fix any distinct $y_1, y_2 \in G = \{10s_0, 01s_1\}$. Then, over all trace-preserving, completely positive (TPCP) maps $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q} \otimes \mathcal{Q})$, we have*

$$\max_{\Phi} \frac{1}{|S|} \Pr[V_k \otimes V_k \text{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \text{ outputs } y_1 \neq y_2 \in G] = \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n \approx 0.854^n.$$

The proof of Lemma 3.4 is based on semidefinite programming (SDPs) (a brief primer on SDPs is given in Appendix B.1). Specifically, one first considers the case of $n = 1$, i.e. the BB84 key is a single qubit (this means the secret key for either s_0 or s_1 is empty, but crucially the adversary does not know for which s_i this

holds). In the non-interactive setting, one can formulate the optimal probability of producing accepting keys for both s_0 and s_1 as a primal SDP via the Choi-Jamiołkowski representation of linear maps [Cho75, Jam72]. In particular, this representation encodes an arbitrary TPCP map Φ as a positive semi-definite matrix $J(\Phi)$ (satisfying additional properties which can be encoded as a linear constraint). Via the corresponding dual SDP and duality theory, one can then upper bound the optimal value of the primal SDP. In particular, one can demonstrate a dual feasible solution Y which yields a (tight) upper bound on the cheating probability of $\alpha = 1/2 + 1/(2\sqrt{2}) \approx 0.854$. The argument generalizes analogously to the setting of n BB84 key qubits where $n > 1$, in which a similar primal SDP can be formulated, and for which the dual feasible solution $Y^{\otimes n}$ yields a (tight) upper bound on the cheating probability of α^n, which is inverse exponential in n.

4 Impossibility Results

We now discuss "tightness" of our protocol with respect to impossibility results. To begin, it is easy to argue that OTMs cannot exist in the plain model (i.e. without additional assumptions) in both the classical and quantum settings: in the classical setting, impossibility holds, since software can always be copied. Quantumly, this follows by a simple rewinding argument [BGS13]. Here, we give two simple no-go results for the quantum setting which support the idea that our scheme is "tight" in terms of the minimality of the assumptions it uses: First, a stateless token which can be queried in *superposition* cannot be used to securely construct an OTM (Section 4.1). Second, for *measure and access* schemes such as ours, in order for a stateless token to allow statistical security, it must have an *exponential* number of keys per secret bit (Section 4.2).

4.1 Impossibility: Tokens which can be queried in superposition

In our construction, we require that all queries to the token be classical strings, i.e. no querying in superposition is allowed. It is easy to argue via a standard rewinding argument that relaxing this requirement yields impossibility of a secure OTM, as we now show. Specifically, let M be a quantum OTM implemented using a hardware token. Without loss of generality, we may model the token as an oracle O_f realizing a function $f : \{0,1\}^n \mapsto \{0,1\}$ in the standard way, i.e. for all $y \in \{0,1\}^n$ and $b \in \{0,1\}$,

$$O_f|y\rangle|b\rangle = |y\rangle|b \oplus f(y)\rangle.$$

Now, suppose our OTM stores two secret bits s_0 and s_1, and provides the receiver with an initial state $|\psi\rangle \in A \otimes B \otimes C$, where A, B, and C are *ancilla* (i.e. algorithm's workspace), *query* (i.e. input to O_f), and *answer* (i.e. O_f's answers) registers, respectively. By definition, an honest receiver must be able to access precisely one of s_0 or s_1 with certainty, given $|\psi\rangle$. Thus, for any $i \in \{0,1\}$, there exists a quantum query algorithm $A_i = U_m O_f \cdots O_f U_2 O_f U_1$ for unitaries $U_i \in \mathcal{U}(A \otimes B \otimes C)$ such that $A_i|\psi\rangle = |\psi'\rangle_{AB}|s_i\rangle_C$. For any choice of i, however, this implies a malicious receiver can now classically copy s_i to an external register, and then "rewind" by applying $A_i^\dagger$ to $|\psi'\rangle_{AB}|s_i\rangle_C$ to recover $|\psi\rangle$. Applying $A_{i'}$ for $i' \neq i$ to $|\psi\rangle$ now yields the second bit i' with certainty as well. We conclude that a quantum OTM which allows superposition queries to a stateless token is insecure (assuming an adversary is not restricted in the quantum operations it can apply).

Remark 4.1. *(1) It is because the token is* stateless *that we are able to model it here via the standard oracle query framework used in the quantum query complexity literature. It is easy to see that the argument above breaks down for a "stateful" token. Specifically, suppose the token is allowed to have a private memory. Then, one can design the token such that each time it is queried, it copies the n input qubits in B via CNOT gates to a fresh set of n ancilla qubits initialized to all zeroes — this effectively* forces *a measurement in the computational basis on B, which prevents A_i from being rewinded via the standard argument above. (2) Above, we assumed the OTM outputs s_i with certainty. The argument can be generalized to the setting*

in which the OTM outputs s_i with probability at least $1-\epsilon$ for small $\epsilon > 0$; in this case, Winter's Gentle Measurement Lemma [Win99] can be used to show that both bits can again be recovered with non-negligible probability.

4.2 Impossibility: Tokens with a bounded number of keys

We have observed that allowing superposition queries to the token prevents an OTM from being secure. One might next ask how simple a hardware token with classical queries can be, while still allowing a secure OTM. We now explore one such strengthening of our construction in which the token is forced to have a bounded number of keys.

To formalize this, let us define the notion of a "measure-and-access (MA)" OTM, i.e. an OTM in which given an initial state $|\psi\rangle$, an honest receiver applies a prescribed measurement to $|\psi\rangle$, and feeds the resulting classical string (i.e. key) y into the token O_f to obtain s_i. Our construction is an example of a MA memory in which each bit s_i has an *exponential* number of valid keys y such that $f(y) = s_i$. One might ask whether the construction can be strengthened such that each s_i has a bounded number (e.g. a polynomial number) of keys. We now show that such a strengthening would preclude security.

For clarity, implicitly in our proof below, we model the oracle O_f as having three possible outputs: 0, 1, or 2, where 2 is output whenever O_f is fed an invalid key y. This is required for the notion of having "few" keys to make sense (i.e. there are 2^n candidate keys, and only two secret bits, each of which is supposed to have a bounded number of keys). Note that our construction indeed fits into this framework.

Lemma 4.2. *Let M be a MA memory with oracle O_f, such that O_f cannot be queried in superposition. If a secret bit s_i has at most Δ keys y_i such that $f(y_i) = s_i$, then given a single copy of $|\psi\rangle$, one can extract both s_0 and s_1 from M with probability at least $1/\Delta^2$.*

Remark 4.3. *The proof is given in Appendix C. Lemma 4.2 shows that in the paradigm of measure-and-access memories, our construction is essentially tight — in order to bound the adversary's success probability of obtaining both secret bits by an inverse exponential, we require each secret bit to have exponentially many valid keys. Second, as in the setting of superposition queries, the above proof can be generalized to the setting in which the OTM returns the correct bit s_i with probability at least $1-\epsilon$ for small $\epsilon > 0$.*

References

[AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In *Proc. 44th Symposium on Theory of Computing (STOC) 2012*, pages 41–60, 2012. Full version availables as arXiv:1203.4740.

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.

[BBBW82] Charles H. Bennett, Gilles Brassard, Seth Breidbard, and Stephen Wiesner. Quantum cryptography, or unforgeable subway tokens. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 267–275. Plenum Press, New York, USA, 1982.

[BBCS92] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. Practical quantum oblivious transfer. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 351–366. Springer, August 1992.

[Bea91] Donald Beaver. Secure multiparty protocols and zero-knowledge proof systems tolerating a faulty minority. *Journal of Cryptology*, 4(2):75–122, 1991.

[BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.

[BGS13] Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs - (extended abstract). In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 344–360. Springer, August 2013.

[BM82] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *23rd FOCS*, pages 112–117. IEEE Computer Society Press, November 1982.

[BS15] Anne Broadbent and Christian Schaffner. Quantum cryptography beyond quantum key distribution. arXiv:1510.06120 [quant-ph], 2015.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[Can00a] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, 2000.

[Can00b] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. `http://eprint.iacr.org/2000/067`.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.

[CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, February 2007.

[CGS08] Nishanth Chandran, Vipul Goyal, and Amit Sahai. New constructions for UC secure computation using tamper-proof hardware. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 545–562. Springer, April 2008.

[Cho75] M. D. Choi. Completely positive linear maps on complex matrices. *Linear Alg. Appl.*, 10:285, 1975.

[CKS+14] Seung Geol Choi, Jonathan Katz, Dominique Schröder, Arkady Yerukhimovich, and Hong-Sheng Zhou. (efficient) universally composable oblivious transfer using a minimal number of stateless tokens. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 638–662. Springer, February 2014.

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, pages 494–503. ACM Press, May 2002.

[CM97] Christian Cachin and Ueli Maurer. Unconditional security against memory-bounded adversaries. In *Advances in Cryptology - CRYPTO 1997*, LNCS, pages 292–306. Springer, 1997.

[DFL+09] Ivan Damgård, Serge Fehr, Carolin Lunemann, Louis Salvail, and Christian Schaffner. Improving the security of quantum protocols via commit-and-open. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 408–427. Springer, August 2009.

[DFSS05] Ivan Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner. Cryptography in the bounded quantum-storage model. In *Symposium on Foundations of Computer Science - FOCS 2005*, pages 449–458. IEEE, 2005.

[DS13] Ivan Damgård and Alessandra Scafuro. Unconditionally secure and universally composable commitments from physical assumptions. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 100–119. Springer, December 2013.

[FKS+13] Serge Fehr, Jonathan Katz, Fang Song, Hong-Sheng Zhou, and Vassilis Zikas. Feasibility and completeness of cryptographic tasks in the quantum world. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 281–296. Springer, March 2013.

[Gav12] Dmitry Gavinsky. Quantum money with classical verification. In *Computational Complexity (CCC), 2012 IEEE 27th Annual Conference on*, pages 42–52, June 2012.

[GIS+10] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia. Founding cryptography on tamper-proof hardware tokens. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 308–326. Springer, February 2010.

[GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. One-time programs. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 39–56. Springer, August 2008.

[GL91] Shafi Goldwasser and Leonid A. Levin. Fair computation of general functions in presence of immoral majority. In Alfred J. Menezes and Scott A. Vanstone, editors, *CRYPTO'90*, volume 537 of *LNCS*, pages 77–93. Springer, August 1991.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

[Hei27] Werner Heisenberg. Schwankungserscheinungen und quantenmechanik. *Zeitschrift fuer Physik*, 40(7):501–506, July 1927.

[HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 411–428. Springer, August 2011.

[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, August 2008.

[Jam72] A. Jamiołkowski. Linear transformations which preserve trace and positive semi-definiteness of operators. *Rep. Math. Phys.*, 3:275, 1972.

[Kat07] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 115–128. Springer, May 2007.

[Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *20th ACM STOC*, pages 20–31. ACM Press, May 1988.

[KMPS14] Daniel Kraschewski, Hemanta K. Maji, Manoj Prabhakaran, and Amit Sahai. A full characterization of completeness for two-party randomized function evaluation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 659–676. Springer, May 2014.

[KMQ11] Daniel Kraschewski and Jörn Müller-Quade. Completeness theorems with constructive proofs for finite deterministic 2-party functions. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 364–381. Springer, March 2011.

[Liu14a] Yi-Kai Liu. Building one-time memories from isolated qubits: (extended abstract). In Moni Naor, editor, *ITCS 2014*, pages 269–286. ACM, January 2014.

[Liu14b] Yi-Kai Liu. Single-shot security for one-time memories in the isolated qubits model. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 19–36. Springer, August 2014.

[Liu15] Yi-Kai Liu. Privacy amplification in the isolated qubits model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 785–814. Springer, April 2015.

[LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 179–188. ACM Press, May / June 2009.

[Mau92] Ueli M. Maurer. Protocols for secret key agreement by public discussion based on common information. In *Advances in Cryptology - CRYPTO 1992*, volume 740 of *LNCS*, pages 461–470. Springer, 1992.

[MPR09] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. Complexity of multi-party computation problems: The case of 2-party symmetric secure function evaluation. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 256–273. Springer, March 2009.

[MPR10] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A zero-one law for cryptographic complexity with respect to computational UC security. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 595–612. Springer, August 2010.

[MR92] Silvio Micali and Phillip Rogaway. Secure computation (abstract). In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 392–404. Springer, August 1992.

[MR11] Ueli Maurer and Renato Renner. Abstract cryptography. In Bernard Chazelle, editor, *ICS 2011*, pages 1–21. Tsinghua University Press, January 2011.

[MVW13] Abel Molina, Thomas Vidick, and John Watrous. Optimal counterfeiting attacks and generalizations for wiesners quantum money. In Kazuo Iwama, Yasuhito Kawano, and Mio Murao, editors, *Theory of Quantum Computation, Communication, and Cryptography*, volume 7582 of *Lecture Notes in Computer Science*, pages 45–64. Springer Berlin Heidelberg, 2013.

[PR08] Manoj Prabhakaran and Mike Rosulek. Cryptographic complexity of multi-party computation problems: Classifications and separations. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 262–279. Springer, August 2008.

[PS04] Manoj Prabhakaran and Amit Sahai. New notions of security: Achieving universal composability without trusted setup. In László Babai, editor, *36th ACM STOC*, pages 242–251. ACM Press, June 2004.

[PW01] Birgit Pfitzmann and Michael Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *Proc. 22nd IEEE Symposium on Security & Privacy (S&P) 2001*, pages 184–200. IEEE, 2001. Full version available at `http://eprint.iacr.org/2000/066`.

[PYJ+12] Fernando Pastawski, Norman Y Yao, Liang Jiang, Mikhail D Lukin, and J Ignacio Cirac. Unforgeable noise-tolerant quantum tokens. *Proceedings of the National Academy of Sciences*, 109(40):16079 –16082, 2012.

[Ren08] Renato Renner. *Security of Quantum Key Distribution*. PhD thesis, ETH Zürich (Switzerland), September 2008.

[Unr10] Dominique Unruh. Universally composable quantum multi-party computation. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 486–505. Springer, May 2010.

[Unr13] Dominique Unruh. Everlasting multi-party computation. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 380–397. Springer, August 2013.

[Unr14] Dominique Unruh. Revocable quantum timed-release encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 129–146. Springer, May 2014.

[Wat11] J. Watrous. Lecture 7: Semidefinite programming, 2011. Latest version available at: https://cs.uwaterloo.ca/~watrous/CS766/LectureNotes/07.pdf.

[Wie83] Stephen Wiesner. Conjugate coding. *ACM SIGACT News*, 15(1):78–88, 1983. Original article written circa 1970.

[Win99] Andreas Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45:2481–2485, 1999.

[WST08] Stephanie Wehner, Christian Schaffner, and Barbara M. Terhal. Cryptography from noisy storage. *Physical Review Letters*, 100(22):220502, June 2008.

[WW10] Stephanie Wehner and Andreas Winter. Entropic uncertainty relations—a survey. *New J. Phys.*, 12(2):025009, Feb 2010.

[WZ82] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.

[Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91. IEEE Computer Society Press, November 1982.

A Universal Composition (UC) Framework

We consider simulation-based security. The Universal Composability (UC) framework was proposed by Canetti [Can01, Can00b], culminating a long sequence of simulation-based security definitions (cf. [GMW87, GL91, MR92, Bea91, Can00a]); please see also [PW01, PS04, CDPW07, LPV09, MR11] for alternative/extended frameworks. Recently Unruh [Unr10] extend the UC framework to the quantum setting. Next, we provide a high-level description of the original classical UC model by Canetti [Can01, Can00b], and then the quantum UC model by Unruh [Unr10].

A.1 Classical UC Model ([Can01, Can00b])

Machines. The basic entities involved in the UC model are players $P_1, \ldots, P_k$ where k is polynomial of security parameter n, an adversary $\mathcal{A}$, and an environment $\mathcal{Z}$. Each entity is modeled as a interactive Turing machine (ITM), where $\mathcal{Z}$ could have an additional non-uniform string as advice. Each P_i has identity i assigned to it, while $\mathcal{A}$ and $\mathcal{Z}$ have special identities $id_{\mathcal{A}} := \texttt{adv}$ and $id_{\mathcal{Z}} := \texttt{env}$.

Protocol Execution. A protocol specifies the programs for each P_i, which we denote as $\pi = (\pi_1, \ldots, \pi_k)$. The execution of a protocol is coordinated by the environment $\mathcal{Z}$. It starts by preparing inputs to all players, who then run their respective programs on the inputs and exchange messages of the form $(id_{\texttt{sender}}, id_{\texttt{receiver}}, \texttt{msg})$. $\mathcal{A}$ can corrupt an arbitrary set of players and control them later on. In particular, $\mathcal{A}$ can instruct a corrupted player sending messages to another player and also read messages that are sent to the corrupted players. During the course of execution, the environment $\mathcal{Z}$ also interacts with $\mathcal{A}$ in an arbitrary way. In the end, $\mathcal{Z}$ receives outputs from all the other players and generates one bit output. We use $\mathsf{EXEC}[\mathcal{Z}, \mathcal{A}, \pi]$ denote the distribution of the environment $\mathcal{Z}$'s (single-bit) output when executing protocol π with $\mathcal{A}$ and the P_i's.

Ideal Functionality and Dummy Protocol. Ideal functionality $\mathcal{F}$ is a trusted party, modeled by an ITM again, that perfectly implements the desired multi-party computational task. We consider an "dummy protocol", denoted $P^{\mathcal{F}}$, where each party has direct communication with $\mathcal{F}$, who accomplishes the desired task according to the messages received from the players. The execution of $P^{\mathcal{F}}$ with environment $\mathcal{Z}$ and an adversary, usually called the simulator $\mathcal{S}$, is defined analogous as above, in particular, $\mathcal{S}$ monitors the communication between corrupted parties and the ideal functionality $\mathcal{F}$. Similarly, we denote $\mathcal{Z}$'s output distribution as $\mathsf{EXEC}[\mathcal{Z}, \mathcal{S}, P^{\mathcal{F}}]$.

Definition A.1 (Classical UC-secure Emulation)**.** *We say π (classically) UC-emulates π' if for any adversary $\mathcal{A}$, there exists a simulator $\mathcal{S}$ such that for all environments $\mathcal{Z}$,*

$$\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\pi] \approx \mathsf{EXEC}[\mathcal{Z},\mathcal{S},\pi']$$

We here consider that $\mathcal{A}$ and $\mathcal{Z}$ are computationally unbounded, and we call it statistical UC-security. We require the running time $\mathcal{S}$ is polynomial in that of $\mathcal{A}$. We call this property Polynomial Simulation*.*

Let $\mathcal{F}$ be a well-formed two party functionality. We say π (classically) UC-realizes $\mathcal{F}$ if for all adversary $\mathcal{A}$, there exists a simulator $\mathcal{S}$ such that for all environments $\mathcal{Z}$, $\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\pi] \approx \mathsf{EXEC}[\mathcal{Z},\mathcal{S},P^{\mathcal{F}}]$. We also write $\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\pi] \approx \mathsf{EXEC}[\mathcal{Z},\mathcal{S},\mathcal{F}]$ if the context is clear.

UC-secure protocols admit a general composition property, demonstrated in the following universal composition theorem.

Theorem A.2 (UC Composition Theorem [Can00b])**.** *Let π,π' and σ be n-party protocols. Assume that π UC-emulates π'. Then σ^{π} UC-emulates $\sigma^{\pi'}$.*

A.2 Quantum UC Model ([Unr10])

Now, we give a high-level description of quantum UC model by Unruh [Unr10].

Quantum Machine. In the quantum UC model, all players are modeled as quantum machines. A quantum machine is a sequence of quantum circuits $\{M^n\}_{n\in\mathbb{N}}$, for each security parameter n. M^n is a completely positive trace preserving operator on space $\mathcal{H}^{\mathtt{state}}\otimes\mathcal{H}^{\mathtt{class}}\otimes\mathcal{H}^{\mathtt{quant}}$, where $\mathcal{H}^{\mathtt{state}}$ represents the internal workspace of M^n and $\mathcal{H}^{\mathtt{class}}$ and $\mathcal{H}^{\mathtt{quant}}$ represent the spaces for communication, where for convenience we divide the messages into classical and quantum parts. We allow a non-uniform quantum advice[3] to the machine of the environment $\mathcal{Z}$, while all other machines are uniformly generated.

Protocol Execution. In contrast to the communication policy in classical UC model, we consider a network $\mathbf{N}$ which contains the space $\mathcal{H}_{\mathbf{N}} := \mathcal{H}^{\mathtt{class}}\otimes\mathcal{H}^{\mathtt{quant}}\otimes_i\mathcal{H}_i^{\mathtt{state}}$. Namely, each machine maintains individual internal state space, but the communication space is shared among all . We assume $\mathcal{H}^{\mathtt{class}}$ contains the message $(id_{\mathtt{sender}}, id_{\mathtt{receiver}}, \mathtt{msg})$ which specifies the sender and receiver of the current message, and the receiver then processes the quantum state on $\mathcal{H}^{\mathtt{quant}}$. Note that this communication model implicitly ensures authentication. In a protocol execution, $\mathcal{Z}$ is activated first, and at each round, one player applies the operation defined by its machine M^n on $\mathcal{H}^{\mathtt{class}}\otimes\mathcal{H}^{\mathtt{quant}}\otimes\mathcal{H}^{\mathtt{state}}$. In the end $\mathcal{Z}$ generates a one-bit output. Denote $\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\Pi]$ the output distribution of $\mathcal{Z}$.

Ideal Functionality. All functionalities we consider in this work are classical, i.e., the inputs and outputs are classical, and its program can be implemented by an efficient classical Turing machine. Here in the quantum UC model, the ideal functionality $\mathcal{F}$ is still modeled as a quantum machine for consistency, but it only applies classical operations. Namely, it measures any input message in the computational basis to get an classical bit-string, and implements the operations specified by the classical computational task.

We consider an "dummy protocol", denoted $P^{\mathcal{F}}$, where each party has direct communication with $\mathcal{F}$, who accomplishes the desired task according to the messages received from the players. The execution of $P^{\mathcal{F}}$ with environment $\mathcal{Z}$ and an adversary, usually called the simulator $\mathcal{S}$, is defined analogous as above, in particular, $\mathcal{S}$ monitors the communication between corrupted parties and the ideal functionality $\mathcal{F}$. Similarly, we denote $\mathcal{Z}$'s output distribution as $\mathsf{EXEC}[\mathcal{Z},\mathcal{S},P^{\mathcal{F}}]$. For simplicity, we also write it as $\mathsf{EXEC}[\mathcal{Z},\mathcal{S},\mathcal{F}]$.

[3] Unruh's model only allows classical advice, but we tend to take the most general model. It is easy to justify that almost all results remain unchanged, including the composition theorem. See [HSS11, Section 5] for more discussion.

Definition A.3 (Quantum UC-secure Emulation)**.** *We say* Π *quantum-UC-emulates* Π' *if for any quantum adversary* $\mathcal{A}$*, there exists a (quantum) simulator* $\mathcal{S}$ *such that for all quantum environments* $\mathcal{Z}$*,*

$$\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\Pi] \approx \mathsf{EXEC}[\mathcal{Z},\mathcal{S},\Pi']$$

We consider here that $\mathcal{A}$ *and* $\mathcal{Z}$ *are computationally unbounded, we call it (quantum) statistical UC-security. We require the running time* $\mathcal{S}$ *is polynomial in that of* $\mathcal{A}$*. We call this property* Polynomial Simulation*.*

Similarly, (quantum) computational UC-security can be defined. Let $\mathcal{F}$ be a well-formed two party functionality. We say Π **quantum-UC-realizes** $\mathcal{F}$ if for all quantum adversary $\mathcal{A}$, there exists a (quantum) simulator $\mathcal{S}$ such that for all quantum environments $\mathcal{Z}$, $\mathsf{EXEC}[\mathcal{Z},\mathcal{A},\Pi] \approx \mathsf{EXEC}[\mathcal{Z},\mathcal{S},P^{\mathcal{F}}]$.

Quantum UC-secure protocols also admit general composition:

Theorem A.4 (Quantum UC Composition Theorem [Unr10, Theorem 11])**.** *Let* Π,Π' *and* Σ *be quantum-polynomial-time protocols. Assume that* Π *quantum UC-emulates* Π'*. Then* Σ^{Π} *quantum UC-emulates* $\Sigma^{\Pi'}$*.*

Remark A.5. *Out of the two protocol parties (the sender and the receiver), we consider security only in the case of the receiver being a corrupted party. Note that we are only interested in cases where the same party is corrupted with respect to all composed protocol. Furthermore, we only consider static corruption.*

B Security Analysis for the Token

We now provide the technical result (Theorem B.1) that is used to prove security of our Quantum OTM construction of Section 3.1. An informal statement of the following theorem is as follows: given a single copy of n-qubit $|x\rangle_\theta$, the probability that an unbounded adversary is able to extract both bits s_0 and s_1 is exponentially small (in n), *even if* the adversary is allowed to query the token a polynomial number of times. For the proofs of this section, we model a token V as outputting one of three strings: 000 for "reject", $10s_0$ for "accept and return s_0", and $01s_1$ for "accept and return s_1". Furthermore, we refer to the states (i.e. the quantum keys) from conjugate coding as *BB84* state, for the important role that they play in the quantum key distribution protocol [BB84].

Theorem B.1. *Let* $\mathcal{X},\mathcal{Q}$ *be finite dimensional Hilbert spaces, and let* $S = \{|\psi_k\rangle, V_k\}$ *denote the ensemble of BB84 states* $|\psi_k\rangle \in \mathcal{X}$ *and corresponding oracles* $V_k : \mathcal{Q} \mapsto \{0,1\}^3$ *used in Section 3.1. Then, for any interactive strategy* Φ *(formally, a trace-preserving, completely positive (TPCP) map* $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q}\otimes\mathcal{Q})$*) which queries* V_k *m times,*

$$\frac{1}{|S|}\Pr[V_k \otimes V_k \text{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \text{ outputs } 10s_001s_1] \leq 2\binom{m}{2}\left(\frac{1}{2}+\frac{1}{2\sqrt{2}}\right)^n. \tag{3}$$

Proof. The claim follows immediately by setting $G = \{10s_0, 01s_1\}$ and combining Lemmas B.2 and B.3. □

Lemmas B.2 and B.3 and their proofs are found in Sections B.2 and B.3, respectively. We begin in Section B.1 by stating the relevant notation for this section, and give brief reviews of quantum channels and semidefinite programming.

B.1 Notation, quantum channels, and semidefinite programming

Notation. Let $\mathcal{X}$ be a finite dimensional complex Hilbert space. Then, $\mathcal{L}(\mathcal{X})$, $\mathrm{Herm}(\mathcal{X})$, $\mathrm{Pos}(\mathcal{X})$, and $\mathcal{D}(\mathcal{X})$ denote the sets of linear, Hermitian, positive semidefinite, and density operators acting on $\mathcal{X}$, respectively. The notation $A \succeq B$ means $A - B$ is positive semidefinite.

Quantum channels. A linear map $\Phi : \mathcal{L}(\mathcal{X}) \mapsto \mathcal{L}(\mathcal{Y})$ is a *quantum channel* if Φ is trace-preserving and completely positive (TPCP). These are the channels which map density operators to density operators. For the semidefinite programs in Section B.3, a useful representation of linear maps $\Phi : \mathcal{L}(\mathcal{X}) \mapsto \mathcal{L}(\mathcal{Y})$ known as the Choi-Jamiołkowski matrix, $J(\Phi) \in \mathcal{L}(\mathcal{Y} \otimes \mathcal{X})$, will be used. The latter is defined (with respect to some choice of orthonormal basis $\{|i\rangle\}$ for $\mathcal{X}$) as

$$J(\Phi) = \sum_{i,j} \Phi(|i\rangle\langle j|) \otimes |i\rangle\langle j|.$$

We use the following three properties of $J(\Phi)$ [Cho75, Jam72]:

- Φ is completely positive if and only if $J(\Phi) \succeq 0$.
- Φ trace-preserving if and only if $\mathrm{Tr}_{\mathcal{Y}}(J(\Phi)) = I_{\mathcal{X}}$.
- For any input state $|\phi\rangle \in \mathcal{X}$ and $|\psi\rangle \in \mathcal{Y}$, $\mathrm{Tr}(\Phi(|\phi\rangle\langle\phi|)|\psi\rangle\langle\psi|) = \mathrm{Tr}(J(\Phi)|\psi\rangle\langle\psi| \otimes |\overline{\phi}\rangle\langle\overline{\phi}|)$ for $|\overline{\phi}\rangle$ the complex conjugate of $|\phi\rangle$.

Semidefinite programs. We present semidefinite programs (SDPs) in a form useful for quantum information, as done (e.g.) in the notes of Watrous [Wat11] or [MVW13]. For further details, a standard text on convex optimization is Boyd and Vandenberghe [BV04]. Given any 3-tuple (A, B, Φ) for operators $A \in \mathrm{Herm}(\mathcal{X})$ and $B \in \mathrm{Herm}(\mathcal{Y})$, and linear map $\Phi : \mathcal{L}(\mathcal{X}) \mapsto \mathcal{L}(\mathcal{Y})$ mapping Hermitian operators to Hermitian operators, one can state a *primal* and *dual* semidefinite program:

Primal problem (P)		Dual problem (D)	
sup	$\mathrm{Tr}(AX)$	inf	$\mathrm{Tr}(BY)$
s.t.	$\Phi(X) = B,$	s.t.	$\Phi^*(Y) \succeq A$
	$X \in \mathrm{Pos}(\mathcal{X}),$		$Y \in \mathrm{Herm}(\mathcal{Y}),$

where Φ^* denotes the *adjoint* of Φ, which is the unique map satisfying $\mathrm{Tr}(A^\dagger \Phi(B)) = \mathrm{Tr}((\Phi^*(A))^\dagger B)$ for all $A \in \mathcal{L}(\mathcal{Y})$ and $B \in \mathcal{L}(\mathcal{X})$. The remarkable power of SDPs lies in the concept of *weak duality*, which states that the optimal value of P is upper bounded by the optimal value of Q. (Note that not all SDPs have feasible solutions; in this case, we label the optimal values as $-\infty$ for P and ∞ for D, respectively.) Thus, if one can phrase a maximization problem Π as an SDP P, then there is a simple method for upper bounding the optimal value of Π — demonstrate a feasible solution for the dual program D. This is precisely the technique used to upper bound the probably of cheating in Section B.3.

B.2 Step 1: Reducing the interactive case to the non-interactive case

To reduce the interactive setting to the non-interactive one, we apply a combinatorial technique of Reference [PYJ$^+$12] (see Theorem 9 therein). In [PYJ$^+$12], this approach was used to analyze two schemes for quantum money in which a verification oracle is queried (possibly multiple times), with each query outputting a single bit (i.e. accept or reject). Our setting, however, requires a classical oracle which outputs *multiple* bits (i.e. accept/reject, along with possibly multiple secret bits for the accept case). In this section, we observe that the technique of [PYJ$^+$12] can be easily generalized to any setting of the following form: One is given (according to some distribution) a state $|\psi_k\rangle$ and classical oracle V_k, and asked to measure $|\psi_k\rangle$ to determine some classical *key* to input to V_k in order to obtain some desired output from V_k. Unlike [PYJ$^+$12], the oracles V_k can output strings of any length — the only restriction is that all V_k must share the same set of output strings (i.e. the set of output strings is independent of k, so that intuitively, no information about keys

is leaked by the output of V_k). We give a formal statement of the generalized presentation of the technique below, and a proof which follows that of Theorem 9 of [PYJ+12] closely.

To do so, we first introduce the notion of a *fixed-output ensemble of states and oracles*. Specifically, let $\mathcal{X}, \mathcal{Q}$ be finite dimensional Hilbert spaces. A fixed-output ensemble $\{|\psi_k\rangle, V_k\}$ of states $|\psi_k\rangle \in \mathcal{X}$ and oracles V_k is one with the following properties. Each V_k accepts a (mixed) state $\rho \in \mathcal{D}(\mathcal{Q})$ as input, measures ρ according to some s-outcome projective measurement $\{\Pi_i^k\}_{i=1}^s$, and outputs t bits according to these rules: To denote "reject", V_k outputs 0^t. Otherwise, V_k outputs a non-zero t-bit string y; denote the set of such "good" outputs y as $G \subseteq \{0,1\}^t$. Note that $s = |G| + 1$, and $t \geq \lceil \log_2(s) \rceil$. Crucially, each V_k may have a distinct set of measurement operators $\{\Pi_i^k\}_{i=1}^s$, but all V_k share the *same* set of output strings $G \cup \{0^t\}$ (hence the name *fixed-output*).

Lemma B.2 (see also Theorem 9 of [PYJ+12])**.** *Let $\mathcal{X}, \mathcal{Q}$ be finite dimensional Hilbert spaces, and let $S = \{|\psi_k\rangle, V_k\}$ be a fixed-output ensemble of states of $\mathcal{X}$ and $\mathcal{Q}$. Fix any distinct $y_1, y_2 \in G$, and suppose that for any trace-preserving, completely positive (TPCP) map $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q} \otimes \mathcal{Q})$,*

$$\frac{1}{|S|} \Pr[V_k \otimes V_k \text{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \text{ outputs } y_1 \neq y_2 \in G] \leq p \tag{4}$$

where $0 \leq p \leq 1$. Then, for any strategy Γ which, given $|\psi_k\rangle$ with probability $1/|S|$, queries V_k m times, it holds that the probability that at least two queries to V_k return distinct strings in G is at most $\binom{m}{2} |G| (|G| - 1) p$.

To apply Lemma B.2 to our setting, we can model our token so that it returns 3 bits: The first two bits are set to 00 for reject, 10 for acceptance of a key for the first choice of secret bit, and 01 for acceptance for the second secret bit. The third bit encodes the desired secret bit. Thus, we define G in Lemma B.2 as $G = \{10s_0, 01s_1\}$ for secret bits $s_0, s_1 \in \{0,1\}$. This generalizes straightforwardly if one wishes to encode more secret bits or even secret strings of longer length.

Proof of Lemma B.2. We follow [PYJ+12] closely. We can model Γ as follows. For each query $i \in [m]$, there is a query space $\mathcal{Q}_i$, such that in the ith query, V_k maps $\mathcal{D}(\mathcal{Q}_i)$ to $\{0,1\}^t$ via a projective measurement. This measurement is given by projectors $M = \{\Pi_i^k\}_{i=1}^s$, where $s = |G| + 1$, and the Π_i^k are labelled by outcomes in $G \cup \{0^t\}$. Since V_k acts in tensor product across different $\mathcal{Q}_i$, we may assume without loss of generality that Γ performs the queries in sequence, reading the outcome of query i before performing query $i+1$. Thus, we can view Γ as a sequence of TPCP maps Γ_i as follows. The first map satisfies $\Gamma_1 : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{X} \otimes \mathcal{Q}_1)$. For $1 < i < m+1$, we have $\Gamma_i : \mathcal{D}(\mathcal{X}) \otimes \mathcal{D}(\mathcal{C}_i) \mapsto \mathcal{D}(\mathcal{X} \otimes \mathcal{Q}_i)$, where $\mathcal{C}_i$ stores the t-bit classical outcome of query i. Finally, $\Gamma_{m+1} : \mathcal{D}(\mathcal{X}) \otimes \mathcal{D}(\mathcal{C}_m) \mapsto \mathcal{D}(\mathcal{X})$. After each Γ_i is applied, the oracle is applied to $\mathcal{Q}_i$, placing output $x_i \in \{0,1\}^t$ in $\mathcal{C}_i$.

Now, conditioned on $x_1 \cdots x_i$, the map Γ_{i+1} is determined (for $i \geq 1$). Thus, any setting of $\mathcal{C}_1 \otimes \cdots \otimes \mathcal{C}_m$ to $x = x_1 \cdots x_m \in \{0,1\}^t \times \cdots \times \{0,1\}^t$ uniquely determines a sequence of maps $\Gamma_x = \Gamma_{m+1} \circ \cdots \circ \Gamma_1$. Note that in general, $\Gamma_x(|\psi_k\rangle\langle\psi_k|)$ is not equal to the state produced by applying Γ to $|\psi_k\rangle\langle\psi_k|$, since Γ involves postselection due to oracle queries, which make the reduced state on $\mathcal{X}$ consistent with string x. However, since Γ_x does not act on register $\mathcal{Q}_i$ after it applies map Γ_i, we can recover the output of Γ on $|\psi_k\rangle\langle\psi_k|$ if we defer V_k's measurement until the end of the computation, and postselect on outcomes x. In other words, for any fixed pair $(|\psi_k\rangle, V_k)$, the probability that applying Γ to $|\psi_k\rangle\langle\psi_k|$, up to and including query q, will yield $x \in (\{0,1\}^t)^{\times q}$ is

$$\Pr(x_1 \cdots x_q \mid (|\psi_k\rangle, V_k)) = \mathrm{Tr}[(\Gamma_q \circ \cdots \circ \Gamma_1)(|\psi_k\rangle\langle\psi_k|)\Pi_{x_1}^k \otimes \cdots \otimes \Pi_{x_q}^k],$$

where $\Pi_{x_i}^k \in M$ is the projector which outputs measurement outcome x_i.

Let Y denote the set of all strings of form $x_1 \cdots x_q$ satisfying properties (1) $q \leq m$ (i.e. x corresponds to $q \leq m$ queries), (2) $x_i \in \{0,1\}^t$ for $1 \leq i \leq q$, (3) $x_q \in G$, and (4) there exists precisely one index $i \neq j$ such that $x_i \neq x_q$ and $x_i \in G$. Note that the elements of Y can have different lengths, and that $|Y| = \binom{m}{2} |G| (|G| - 1)$. Then, the statement of the claim wishes to upper bound the quantity

$$\frac{1}{|S|} \sum_k \sum_{\substack{x_1 \cdots x_m \in \left(\{0,1\}^t\right)^{\times m} \\ \text{s.t. } \exists\, i \neq j \text{ with } x_i \neq x_j \in G}} \Pr(x \mid (|\psi_k\rangle, V_k)) = \sum_{x \in Y} \frac{1}{|S|} \sum_k \Pr(x \mid (|\psi_k\rangle, V_k)). \tag{5}$$

Above, the second equality crucially uses the fact that the ensemble $(|\psi_k\rangle, V_k)$ is *fixed-output*, which yields that Y is independent of k, and so the sums commute. As done in [PYJ+12], observe now that for any $x = x_1 \cdots x_q \in Y$ such that $x_i \neq x_q \in G$,

$$\begin{aligned}
\frac{1}{|S|} \sum_k \Pr(x \mid (|\psi_k\rangle, V_k)) &= \frac{1}{|S|} \sum_k \mathrm{Tr}[(\Gamma_q \circ \cdots \circ \Gamma_1)(|\psi_k\rangle\langle\psi_k|) \Pi^k_{x_1} \otimes \cdots \otimes \Pi^k_{x_q}] \\
&\leq \frac{1}{|S|} \sum_k \mathrm{Tr}[(\Gamma_q \circ \cdots \circ \Gamma_1)(\rho) I_{x_1} \otimes \cdots I_{x_{i-1}} \otimes \Pi_{x_i} \otimes I_{x_{i-1}} \cdots \otimes I_{x_{q-1}} \otimes \Pi_{x_q}] \\
&\leq p,
\end{aligned}$$

where the last inequality follows by Equation (4) of the claim. Combining this with Equation (5) and the fact that $|Y| = \binom{m}{2} |G| (|G| - 1)$, the claim follows. □

B.3 Step 2: The non-interactive case

Lemma B.2 yields that in order to prove that it is unlikely for a corrupt receiver to extract both secret bits from the stateless token, it suffices to consider the following question: Given only the initial state $|\psi_k\rangle$, what is the maximum probability with which accepting keys (with respect to oracle V_k) for both secret bits can be extracted from $|\psi_k\rangle$ (i.e. no queries to V_k are allowed)? In the terminology of Lemma B.2, we have $G = \{10s_0, 01s_1\}$ for secret bits $s_0, s_1 \in \{0,1\}$, with 000 the string output by V_k to denote "reject". To analyze this, we apply the semidefinite programming based results of Molina, Vidick, and Watrous [MVW13], which studied similar BB84-based schemes in the context of quantum money. The lemma below is a special case of Lemma 5 of [MVW13] (our re-statement below is formulated with respect to the context of this paper), whose proof is essentially identical to that of Lemma 5 of [MVW13], save for some minor modifications for our setting. For completeness, we give a full proof below. Any modifications to the proof of [MVW13] are explicitly noted; among these is the fact that we require distinct outputs $y_1 \neq y_2 \in G$ below, which results in a better security bound in Lemma B.3 than that obtained in [MVW13] (details in the proof below).

Lemma B.3 (see Lemma 5 of [MVW13])**.** *Let $\mathcal{X} = (\mathbb{C}^2)^{\otimes n}$ and $\mathcal{Q} = (\mathbb{C}^2)^{\otimes (n+1)}$, and let $S = \{|\psi_k\rangle, V_k\}$ be the fixed-output ensemble of states of $\mathcal{X}$ and $\mathcal{Q}$ corresponding to the construction of Section 3.1. Fix any distinct $y_1, y_2 \in G = \{10s_0, 01s_1\}$. Then, over all trace-preserving, completely positive (TPCP) maps $\Phi : \mathcal{D}(\mathcal{X}) \mapsto \mathcal{D}(\mathcal{Q} \otimes \mathcal{Q})$, we have*

$$\max_{\Phi} \frac{1}{|S|} \Pr[V_k \otimes V_k \textit{ applied to } \Phi(|\psi_k\rangle\langle\psi_k|) \textit{ outputs } y_1 \neq y_2 \in G] = \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n \approx 0.854^n. \tag{6}$$

We note that the maximum value of $(\frac{1}{2} + \frac{1}{2\sqrt{2}})^n = \cos^2(\frac{\pi}{8})$ is a common constant in quantum information processing, and is attained by the tensor product of single-qubit measurements in the *Breidbart basis* [BBCS92, BBBW82].

Proof of Lemma B.3. Recall that in the construction of Section 3.1, our protocol picks $|\psi_k\rangle \in (\mathbb{C}^2)^{\otimes n}$ by locally setting each qubit (uniformly and independently at random) to one of $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. The corresponding oracle V_k then accepts a string $0y \in \{0,1\}^{n+1}$ (resp. $1y \in \{0,1\}^{n+1}$) if the encoded bits for all the Z-basis (resp. X-basis) qubits of $|\psi_k\rangle$ are correct; in this case, V_k outputs $10s_0$ (resp. $01s_1$). We now produce the semidefinite programming-based (SDP) proof of Lemma 5 of [MVW13] (with very minor modifications, indicated below), which goes as follows: First, the case of $n = 1$ (i.e. a single-qubit BB84 key) is analyzed. Based on this, a result for the parallel repetition of this 1-qubit scheme n times easily follows.

Assume without loss of generality that the adversary (i.e. Φ) places its candidate Z-basis and X-basis keys into the first and second copies of $\mathcal{Q}$, respectively, since recall the goal is for the adversary to obtain distinct outputs $y_1 \neq y_2 \in G$. (In comparison, [MVW13] decides which question to ask for each copy of $\mathcal{Q}$ uniformly at random. Thus, with probability $1/2$, it is optimal to place the same key in both copies of $\mathcal{Q}$, obtaining identical outputs $y_1 = y_2 \in G$. This also yields a weaker security bound in [MVW13].) Then, the primal semidefinite program for the expression in Equation (6) is given by

$$\begin{aligned} \max \quad & \mathrm{Tr}(XA) \\ \text{s.t.} \quad & \mathrm{Tr}_{\mathcal{Q}\otimes\mathcal{Q}}(X) = I_{\mathcal{X}} \\ & X \in \mathrm{Pos}(\mathcal{Q}\otimes\mathcal{Q}\otimes\mathcal{X}), \end{aligned}$$

where the objective function operator A is (for H the Hadamard gate)

$$\begin{aligned} A &= \frac{1}{4}\left[\sum_{b\in\{0,1\}} |0\rangle\langle 0| \otimes |b\rangle\langle b| \otimes |1\rangle\langle 1| \otimes I \otimes |b\rangle\langle b| + \sum_{c\in\{0,1\}} |0\rangle\langle 0| \otimes I \otimes |1\rangle\langle 1| \otimes |c\rangle\langle c| \otimes H|c\rangle\langle c|H\right] \\ &= \frac{1}{4}\sum_{b,c\in\{0,1\}} |0\rangle\langle 0| \otimes |b\rangle\langle b| \otimes |1\rangle\langle 1| \otimes |c\rangle\langle c| \otimes (|b\rangle\langle b| + H|c\rangle\langle c|H). \end{aligned}$$

Intuitively, X corresponds to the Choi-Jamiołkowski matrix $J(\Phi)$ of some linear map $|\Phi\rangle : \mathcal{L}(\mathcal{X}) \mapsto \mathcal{L}(\mathcal{Q}\otimes\mathcal{Q})$. The conditions $\mathrm{Tr}_{\mathcal{Q}\otimes\mathcal{Q}}(X) = I_{\mathcal{X}}$ and $X \in \mathrm{Pos}(\mathcal{Q}\otimes\mathcal{Q}\otimes\mathcal{X})$ enforce that Π is trace-preserving and completely positive, respectively. Finally, A acts on five registers (in [MVW13], A acts on three registers): The first and third of these are single qubits which simply hold the values $|0\rangle$ and $|1\rangle$, which are the single-bit prefixes (i.e. choice bit b) required by V_k for Z-basis and X-basis queries, respectively (these two registers are absent in [MVW13]). The last register holds the BB84 state which was prepared (i.e. either $|0\rangle$, $|1\rangle$, $|+\rangle$, or $|-\rangle$), and the second and fourth registers hold the corresponding encoded bit for the Z- and X-bases, respectively. Thus, e.g. when the last register has value $|+\rangle$, then the second register reads I since any value is allowed here, but the fourth register must read $|0\rangle$, since this is the encoded bit. Finally, for convenience, let us define A' on registers 2, 3, and 4 so that we can rewrite A as

$$A = |0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes A', \tag{7}$$

i.e. we have simply factored out registers 1 and 3 in the expression for A.

For brevity, define $V_{b,c} = (|b\rangle\langle b| + H|c\rangle\langle c|H)$. Since $\langle b|H|c\rangle = 1/\sqrt{2}$ for all $b, c \in \{0,1\}$, by considering $\mathrm{Tr}(V_{b,c})$ and $\mathrm{Tr}(V_{b,c}^2)$, one easily obtains that the eigenvalues of $V_{b,c}$ are $1 \pm (1/\sqrt{2})$. Suppose the corresponding eigenvectors for $V_{0,0}$ are $|\psi_+\rangle$ and $|\psi_-\rangle$, respectively. Then, a primal feasible solution is given by

$$X = |0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes (|00\rangle\langle 00| \otimes |\psi_+\rangle\langle\psi_+| + |11\rangle\langle 11| \otimes |\psi_-\rangle\langle\psi_-|) =: |0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes X'. \tag{8}$$

(In [MVW13], the first and third registers are omitted above.) Observing that $V_{11} = 2I - V_{00}$, this solution obtains objective value $1/2 + 1/(2\sqrt{2})$. This is now shown tight by demonstrating a matching dual solution.

Specifically, the dual SDP is given by

$$\begin{aligned} \max \quad & \mathrm{Tr}(Y) \\ \text{s.t.} \quad & I_{\mathcal{Q}\otimes\mathcal{Q}} \otimes Y \succeq A \\ & Y \in \mathrm{Herm}(\mathcal{X}). \end{aligned}$$

Since A is block-diagonal, the first condition can be simplified from $I_{\mathcal{Q}\otimes\mathcal{Q}} \otimes Y \succeq A$ to $\forall\ b,c\ Y \succeq \frac{1}{4}V_{b,c}$. Since the largest eigenvalue of any $V_{b,c}$ is $1+(1/\sqrt{2})$, it follows that $Y = (1/4 + 1/(4\sqrt{2}))I$ is a feasible solution for the dual SDP with objective value $1/2 + 1/(2\sqrt{2})$, matching the primal feasible solution's value. This completes the analysis for the $n = 1$ case.

The case of $n > 1$ now follows easily: The SDPs above generalize straightforwardly, and for optimal solutions X and Y for the $n = 1$ case above, the generalized solutions are $|0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes (X')^{\otimes n}$ and $Y^{\otimes n}$, respectively. Formally, define $\mathcal{X}^n = \mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n$, and analogously for $\mathcal{Q}^n$. For brevity, set $\mathcal{S}^n_{qqx} = (\mathcal{Q}_1 \otimes \mathcal{Q}_1 \otimes \mathcal{X}_1) \otimes \cdots \otimes (\mathcal{Q}_n \otimes \mathcal{Q}_n \otimes \mathcal{X}_n)$. The primal and dual SDPs of the n-fold repetition are now

Primal problem (P)

$$\begin{aligned} \max \quad & \mathrm{Tr}(X(U_\pi|0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes (A')^{\otimes n}U_\pi^\dagger)) \\ \text{s.t.} \quad & \mathrm{Tr}_{\mathcal{Q}^n\otimes\mathcal{Q}^n}(X) = I_{\mathcal{X}^n} \\ & X \in \mathrm{Pos}(\mathcal{Q}^n \otimes \mathcal{Q}^n \otimes \mathcal{X}^n), \end{aligned}$$

Dual problem (D)

$$\begin{aligned} \max \quad & \mathrm{Tr}(Y) \\ \text{s.t.} \quad & I_{\mathcal{Q}^n\otimes\mathcal{Q}^n} \otimes Y \succeq U_\pi(|0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes (A')^{\otimes n})U_\pi^\dagger \\ & Y \in \mathrm{Herm}(\mathcal{X}^n), \end{aligned}$$

where U_π is a permutation aligning space $\mathcal{Q}^n \otimes \mathcal{Q}^n \otimes \mathcal{X}^n$ with $\mathcal{S}^n_{qqx}$. If the optimal solution to the $n = 1$ primal SDP from Equation (8) is $|0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes X'$ with value α, it is clear that $|0\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_3 \otimes (X')^{\otimes n}$ obtains value α^n for P above. Similarly, for optimal solution Y to the $n = 1$ dual SDP with value α, $Y^{\otimes n}$ achieves α^n for D above; this follows since for any operators C and D, we have that $\mathrm{Tr}(C \otimes D) = \mathrm{Tr}(C)\mathrm{Tr}(D)$ and that $C \succeq D \succeq 0$ implies $C^{\otimes n} \succeq D^{\otimes n}$. □

C Proof of Lemma 4.2

Proof. Observe first that an honest receiver Alice wishing to extract s_i acts as follows. She applies a unitary $U_i \in \mathcal{U}(A \otimes B)$ to get state

$$|\phi_1\rangle := U_i|\psi\rangle_{AB}|0\rangle_C. \tag{9}$$

She then measures B in the computational basis and postselects on result $y \in \{0,1\}^n$, obtaining state

$$|\phi_2\rangle := |\phi_y\rangle_A|y\rangle_B|0\rangle_C. \tag{10}$$

She now treats y as a "key" for s_i, i.e. she applies O_f to $B \otimes C$ to obtain her desired bit s_i, i.e.

$$|\phi_3\rangle := |\phi_y\rangle_A|y\rangle_B|s_i\rangle_C. \tag{11}$$

A malicious receiver Bob wishing to extract s_0 and s_1 now acts similarly to the rewinding strategy for superposition queries. Suppose without loss of generality that s_0 has at most Δ keys. Then, Bob first applies

U_0 to prepare $|\phi_1\rangle$ from Equation (9), which we can express as

$$|\phi_1\rangle = \sum_{y\in\{0,1\}^n} \alpha_y |\psi_y\rangle_A |y\rangle_B |0\rangle_C. \tag{12}$$

for $\sum_y |\alpha_y|^2 = 1$. Since measuring B next would allow us to retrieve s_0 in register C with certainty, we have that all y appearing in the expansion above satisfy $f(y) = s_i$. Moreover, since s_0 has at most Δ keys, there exists a key y' such that $|\alpha_{y'}|^2 \geq 1/\Delta$. Bob now measures B in the computational basis to obtain $|\phi_2\rangle$ from Equation (10), obtaining y' with probability at least $1/\Delta$. Feeding y' into O_f yields s_1. Having obtained y', we have that $|\langle\phi_1|\phi_2\rangle|^2 \geq 1/\Delta$, implying

$$\left|\langle\psi|U_0^\dagger|\phi_{y'}\rangle|y'\rangle\right|^2 \geq 1/\Delta,$$

i.e. Bob now applies $U_0^\dagger$ to recover a state with "large" overlap with initial state $|\psi\rangle$.

To next recover s_1, define $|\psi_{\text{good}}\rangle := U_1|\psi\rangle$ and $|\psi_{\text{approx}}\rangle := U_1 U_0^\dagger |\phi_{y'}\rangle|y'\rangle$. Bob applies U_1 to obtain

$$|\psi_{\text{approx}}\rangle = \beta_1|\psi_{\text{good}}\rangle + \beta_2|\psi_{\text{good}}^\perp\rangle,$$

where $\sum_i |\beta_i|^2 = 1$, $\langle\psi_{\text{good}}|\psi_{\text{good}}^\perp\rangle = 0$, and $|\beta_1|^2 \geq 1/\Delta$. Define $\Pi_{\text{good}} := \sum_{y\in\{0,1\}^n \text{ s.t. } f(y)=s_1} |y\rangle\langle y|$. Then, the probability that measuring B in the computational basis now yields a valid key for s_1 is

$$\langle\psi_{\text{approx}}|\Pi_{\text{good}}|\psi_{\text{approx}}\rangle \geq |\beta_1|^2 \geq \frac{1}{\Delta},$$

where we have used the fact that $\Pi_{\text{good}}|\psi_{\text{good}}\rangle = |\psi_{\text{good}}\rangle$ (since an honest receiver can extract s_1 with certainty). We conclude that Bob can extract both s_0 and s_1 with probability at least $1/\Delta^2$. □