Semi-Quantum Money<br>Roy Radian ${ }^{1}$ and Or Sattath ${ }^{1}$<br>${ }^{1}$ Computer Science Department, Ben-Gurion University of the Negev

October 21, 2020


#### Abstract

Quantum money allows a bank to mint quantum money states that can later be verified and cannot be forged. Usually, this requires a quantum communication infrastructure to transfer quantum states between the user and the bank. Gavinsky Gav12 introduced the notion of classically verifiable quantum money, which allows verification through classical communication. In this work we introduce the notion of classical minting, and combine it with classical verification to introduce semi-quantum money. Semi-quantum money is the first type of quantum money to allow transactions with completely classical communication and an entirely classical bank. This work features constructions for both a public memory-dependent semi-quantum money scheme and a private memoryless semi-quantum money scheme. The public construction is based on the works of Zhandry [Zha19] and Coladangelo Col19, and the private construction is based on the notion of Noisy Trapdoor Claw Free Functions (NTCF) introduced by Brakerski et al. $\left.\mathrm{BCM}^{+} 18\right]$.


In terms of technique, our main contribution is a perfect parallel repetition theorem for NTCF.

## 1 Introduction

Introduced by Wiesner circa 1969, quantum money was the precursor to what is now known as quantum cryptography Wie83. The motivation behind quantum money is to design money that is physically impossible to counterfeit, by using a variant of the (quantum) no-cloning theorem WZ82, Par70, Die82. This notion of quantum money is in sharp contrast to our current notions of bills and coins that, at least in principle, can be counterfeited.

All quantum money schemes consist of three parts: key-gen, which generates a key, mint which uses the key to issue a new quantum money state, and verify which tests whether an alleged money state is legitimate. There
are two main categories of quantum money: private and public. In a private setting, the key is required to run the verification. On the other hand, in a public quantum money scheme, key-gen generates a secret/public key-pair, where the secret key is used in mint and the public key is used in verify. In this work we deal both with private and public schemes.

A variant of quantum money called classically verifiable quantum money was introduced in Gav12] (see also [PYJ+12, MVW13, GK15, BS16a, AA17, AGKZ20 ): the money is verified via an interactive protocol between the user and the bank. This protocol requires a quantum computer for the user, a classical computer for the bank, and classical communication between them. In a classical verification, the banknote is always measured and destroyed. This destructive measurement makes sure it cannot be reused 1

In this work, we introduce a new variant of classically verifiable quantum money: semi-quantum money. In this setting, the minting also shares this property, i.e., it is a protocol that involves both the bank and the user, and requires only classical resources from the bank. In standard quantum money, in contrast, minting is a quantum algorithm run by the bank, which sends the output - the quantum money state - to the user, via a quantum channel.

In semi-quantum money, the money state is generated by the user. This concept seems somewhat counterintuitive; if banknotes are generated by the user, could the user not create as many notes as he or she pleases? The key point of the minting process is the protocol between the user and the bank: the user is supposed to generate a superposition over two registers using information provided by the bank, measure one of the registers, and report the result back to the bank. If the user will try to repeat the same procedure, the measurement outcome - as well as the post-measured state - will be different with overwhelming probability. As far as the authors are aware, no prior work considered classical minting.

The fact that semi-quantum money is also classically verifiable means that instead of sending the quantum state to the bank for verification, the user and the bank run a classical interactive verification protocol that tests the validity of the money. Semi-quantum money got its name from the fact that the minting and verification protocols require only classical resources (communication and computation) from the bank.

This introduction of a quantum money scheme where the banks are classical perhaps raises the question whether the concept could be improved, such that the bank would be quantum and the user classical. However, such a setting is inherently flawed; if the user is classical, they could not hold their own money, meaning the bank would have to hold the state of every

[^0]note of every user ${ }^{2}$. This makes the "quantumness" of the money redundant, since it would be permanently kept within the bank in any case. Thus, it would seem that the setting where the bank and communication is classical and the user is quantum is the "least quantum" a quantum money scheme could be.

In this work we introduce both a public construction and a private construction for semi-quantum money. The public construction is based on an existing public quantum money scheme which we combine with an existing tool that allows classical verification, so our public construction requires little technical work. Our private scheme, on the other hand, is based on NTCF - a tool which was designed for a different purpose entirely. Its construction, therefore, entails a much greater technical challenge. For that reason we address the public result first

Assumptions. Our results assume authenticated and noiseless classical channels (which could be realized using standard classical error-correction and authentication techniques), along with perfect quantum devices (quantum memory, quantum computer and quantum communication channels). Of course, such quantum devices are not currently available, and are not predicted to be available in the short term (especially because of the long term quantum memory inherently required for quantum money).

Prior Knowledge. Before we go any further, we discuss the accessibility of this work. The reader is assumed to have a basic understanding of classical cryptography, and we follow the definitions and conventions of Gol04] and KL14. This work is aimed at readers who are familiar with quantum computing, but is also accessible to other readers. For further reading, consult NC11] for general quantum computing, and BS16b for quantum cryptography. The two major "quantum" facts that are crucial to understand for this paper are the following: (i) A qubit is the quantum analog of a bit. Unlike bits, qubits cannot be copied due to the no-cloning theorem. (ii) To extract classical information from qubits, a measurement has to be preformed. The measurement changes the quantum state. Crucially, this process is not reversible. This is in contrast to classical systems, where rewinding is possible.

Public semi-quantum money. In a public quantum money scheme, unlike in a private scheme, any user can verify a banknote using the bank's public key without aid from the bank. There are several advantages for a public scheme: it does not require three-party quantum communication between the bank, the sender, and the receiver. The only requirement is

[^1]a quantum channel between the sender and the receiver ${ }^{3}$. Public schemes have a major advantage over private schemes also in terms of privacy: since the bank is not involved in the transactions, the bank cannot track all transactions of the note. However, it is much harder to construct a secure public scheme - see the related works paragraph below.

We construct a public semi-quantum money scheme based on Zhandry's quantum lightning ( $(\overline{Z h a 19}])$, and the notion of bolt-to-certificate introduced in Coladangelo's follow-up work ([Col19]). Our classical verification based on Col19 is memory-dependent, meaning the bank has to keep a database of spent notes. We leave it as an open question whether a memoryless public semi-quantum money exists (we compare memory-dependent vs. memoryless quantum money in Appendix E).

Our main public result is:
Theorem 1 (Public Semi-Quantum Money). Assuming the existence of a secure Quantum Lightning scheme (Definitions 39 and 40) with bolt-tocertificate capability (Definition 41), and the existence of a $P Q-E U-C M A$ digital signature scheme (Definitions 36 and 37), then a secure memorydependent (Definition (6) public semi-quantum money scheme exists (Definition 9).

Quantum lightning ([Zha19]) is a type of public quantum money such that each quantum banknote (called a bolt) is unique: a "lightning bolt" is a quantum state, and has a serial number that is a classical string. It is hard for everyone, including the bank, to construct two valid bolts with the same serial number. This can be thought of as if someone would "freeze" and "capture" lightning bolts in a thunderstorm that have the same fingerprint (in this case, the serial number).

The notion of bolt-to-certificate was introduced in [Col19], and it describes a process of turning a quantum lightning bolt into a classical certificate. The certificate proves a lightning bolt with a certain serial number has been destroyed (the serial number is classical and thus survives the destruction of the bolt itself).

In this work, we use quantum lightning with bolt-to-certificate capability to construct a public semi-quantum scheme. In public semi-quantum money we want to allow any user to verify banknotes without destroying them, and classical verification with the bank (communication between two users is still quantum). Therefore, we facilitate a quantum verification algorithm to be used by the quantum users (that would preserve the banknote), and a classical verification protocol to be used with the classical bank (that

[^2]would destroy the banknote). The quantum verification is derived directly from quantum lightning verification (since quantum lightning is already a public money scheme), and the classical verification is derived from the bolt-to-certificate capability - the bolt is exchanged for a classical certificate that is shown to the bank to prove the note has been spent. Moreover, we introduce a slight alteration to the quantum lightning scheme such that the banknotes become user-generated; instead of the bank producing the bolt and sending it to the user, the user would generate the bolt and the bank would sign its (classical) serial number. The resulting scheme requires quantum communication between the quantum users and classical communication with the classical bank (banknotes could still be passed between users with only classical communication by going through a bank, in the same manner of a private scheme).

The security notion is as follows: no $w>\ell$ verifications can be made using $\ell$ bolts. In our case this means that, besides the fact that the same bolt cannot be used for two quantum or two classical verifications, the same bolt cannot be used for both a quantum and a classical verification (while the quantum verification does not destroy the note, it passes it to the other user). Moreover, there is an additional security concern to be considered; the idea of sabotage (introduced in the context of quantum lightning in CS20, and in the context of quantum money in [BS16a]. This notion captures the possibility of paying a user with a "sabotaged" note such that it is accepted by the receiving user (i.e., it passes the quantum verification) but will not be later accepted by another user, or by the bank (i.e., will not pass the classical verification with the bank). A quantum lightning scheme that is secure against sabotage is enough to ensure such security for our construction, and such a security proof was made in [CS20].

It should be stated that the current candidate constructions for quantum lightning are problematic. There are currently three candidate constructions: two suggested in the original work ( Zha19]) and one that was published roughly 7 years earlier in $\mathrm{FGH}^{+} 12$ but seems to be compatible with quantum lightning (the connection is made in [CS20]). One of the constructions in [Zha19] is based on "collision resistant non-collapsing hash functions", which currently do not have a candidate construction. The other has been successfully attacked by Roberts [Rob19], though he ended on a positive note with some ideas for modification to thwart the attack. The construction in $\left[\mathrm{FGH}^{+} 12\right]$ seems to be a valid quantum lightning construction (it was introduced 7 years earlier and is yet unbroken), and can be used to construct public quantum money with classical minting, but does not feature bolt-to-certificate capability which is necessary for classical verifiability, thus being unusable for a semi-quantum scheme. This means that our scheme can be based on either of Zhandry's original constructions. Moreover, only one of which (the one that is broken in some respects) was proven to be secure against sabotage in [CS20] under the same assumption. There-
fore, our construction for public semi-quantum money is on shaky ground. Nevertheless, even if the existing constructions of quantum lightning do not provide strong security guarantees, we believe the notion of a quantum lightning scheme, as well as that of bolt-to-certificate, to be plausible. In fact, a new candidate, with hardness based on lattices, was recently announced in an unpublished work by Peter Shor (see https://youtu.be/8fzLByTn8Xk).

Private semi-quantum money. In a private scheme, a banknote can be verified only with a bank. A private semi-quantum money scheme requires only classical communication with a classical bank.

Our main private result is:
Theorem 2 (Private Semi-Quantum Money). Assuming that the Learning With Errors (LWE) problem with certain sets of parameters is hard for BQP, then a secure private semi-quantum money scheme exists (Definition 23).

Our assumptions are stated in each theorem separately, and they all boil down to our LWE assumptions. There are a number of variants of the LWE problem, each one with different security parameters. We rely on different constructions from LWE which use different variants - but both assumptions are LWE assumptions. Interestingly, while private quantum money schemes can be secure without computational assumptions, they are required for semi-quantum money. In fact, there can be no information-theoretic secure quantum money scheme with classical minting (see Section 7).

The main technical tool through which to implement this scheme is the quantum secure trapdoor claw-free function (TCF - Appendix D) recently introduced in $\mathrm{BCM}^{+} 18$ (see also Mah18a, Mah18b, GV19]). Informally, a quantum secure TCF is a family of functions, where each function $f$ : $\{0,1\}^{w} \rightarrow\{0,1\}^{w}$ in the family (a) is classically efficiently computable, (b) is 2-to-1, i.e., for every $x$ there exists a unique $x^{\prime} \neq x$ such that $f(x)=f\left(x^{\prime}\right)$, and (c) has a trapdoor that, given $y$, can be used to find $x$ and $x^{\prime}$ such that $f(x)=f\left(x^{\prime}\right)=y$ (when $y$ is in the image of $f$ ), but without the trapdoor a quantum polynomial adversary cannot find any pair $x, x^{\prime}$ such that $f(x)=f\left(x^{\prime}\right)$.

In addition, we will require the adaptive hardcore bit property of a TCF that was introduced in $\left[\mathrm{BCM}^{+} 18\right]$, which is explained below. Using a quantum computer, the state $\frac{1}{\sqrt{2}}\left(|x\rangle+\left|x^{\prime}\right\rangle\right)$, where $x$ and $x^{\prime}$ are two pre-images of $y$, could be measured, and one pre-image of $y$ could be found. Moreover, by measuring the state in the Hadamard basis, a non-zero string $d$ that satisfies
$d \cdot\left(x \oplus x^{\prime}\right)=0$ could be extracted:

$$
\begin{align*}
H^{\otimes w} \frac{1}{\sqrt{2}}\left(|x\rangle+\left|x^{\prime}\right\rangle\right) & =\frac{1}{\sqrt{2^{w+1}}} \sum_{d \in\{0,1\}^{w}}(-1)^{d \cdot x}+(-1)^{d \cdot x^{\prime}}|d\rangle \\
& =\frac{1}{\sqrt{2^{w-1}}} \sum_{d \in\{0,1\}^{w} \mid d \cdot\left(x \oplus x^{\prime}\right)=0}(-1)^{d \cdot x}|d\rangle \tag{1}
\end{align*}
$$

In our construction we use the following two tests: the pre-image test (providing a pre-image of $y$ ) and the equation test (providing a non-zero $d$ that satisfies the above condition). The adaptive hardcore bit property guarantees that, even though either test can be easily passed on its own, any quantum polynomial time (QPT) adversary can successfully pass both tests with probability at most $\frac{1}{2}+\operatorname{negl}(\lambda)$. Brakerski et al. used these tests to construct a cryptographic test of quantumness (CTQ); our construction can be seen as a reinterpretation of this protocol in the quantum money setting, where the first part of the protocol can serve as the creation of a quantum money state, and the second part can serve as the its verification. The transition to quantum money introduces some challenges; mainly the need for a parallel repetition theorem for our NTCF-based primitive, and proving full-scheme security. Brakerski et al. showed a construction of a noisy trapdoor claw-free function (NTCF) that holds this adaptive hardcore property, based on the hardness of the Learning With Errors (LWE) problem $\left[\mathrm{BCM}^{+} 18\right]$. For the sake of clarity, we ignore the issues related to the noisy property in this introduction.

A TCF on its own, however, is not hard enough to construct a money scheme with; we do not want adversaries to be able to forge banknotes with probability $\frac{1}{2}$. To that end, we would like to amplify the hardness using some sort of a parallel repetition theorem (see Section 4.3). Luckily, we can rephrase this setting using the framework of weakly verifiable puzzles for which a perfect parallel repetition theorem is known [CHS05]. This perfect parallel repetition guarantees that answering both tests for $n$ puzzles correctly is as hard as trying to answer them independently, i.e., at most $\left(\frac{1}{2}\right)^{n}$ (up to negligible corrections), which is exactly our goal.

Next, we present the outline and analysis of our semi-quantum private money scheme construction. The security notion of our money scheme is rather straightforward: an adversary that receives $\ell$ banknotes, and can attempt to pass verification (polynomially) many times, cannot pass more than $\ell$ verifications. To show a construction that meets this notion, we work our way through several weaker security notions; this makes proving the security of our full scheme construction simpler. We first show how to construct a semi-quantum money scheme (Section 5) that provides a weaker level of security than a full scheme. Here, we wish to show that a counterfeiter that receives 1 quantum money state cannot create two states that will both pass verification with non-negligible probability. We call a
scheme that satisfies this weaker notion of security a 2 -of- 2 mini-scheme see Definition 25

We now describe the construction of a 2 -of- 2 mini-scheme, starting with the (honest) minting protocol. The bank picks $n$ functions $f_{1}, \ldots, f_{n}$ uniformly at random from the TCF family and sends them to the user, while keeping the trapdoors $t_{1}, \ldots, t_{n}$ private. The user creates a superposition of the form $\left|\psi_{1}\right\rangle \otimes \ldots \otimes\left|\psi_{n}\right\rangle$, where $\left|\psi_{i}\right\rangle=\frac{1}{\sqrt{2^{w}}} \sum_{x \in\{0,1\}^{w}}|x\rangle \otimes\left|f_{i}(x)\right\rangle$. The user measures all the r.h.s. registers (i.e., $\left|f_{i}(x)\right\rangle \forall 1 \leq i \leq n$ ) and sends the resulting $y_{1}, \ldots, y_{n}$ to the bank, who saves them to its databas ${ }^{4}$. Note that due to the measurement, the $i^{\text {th }}$ state collapses to $\left|\psi_{y_{i}}\right\rangle=\frac{1}{\sqrt{2}}\left(\left|x_{i}\right\rangle+\left|x_{i}^{\prime}\right\rangle\right)$, where $f_{i}\left(x_{i}\right)=f_{i}\left(x_{i}^{\prime}\right)=y_{i}$.

For verification, the bank chooses a random challenge $C_{i} \in_{R}\{0,1\}$ (which is either the pre-image or the equation challenge) for each of the $n$ registers. For the pre-image challenge, $C_{i}=0$, the user must provide a string $x_{i}$ such that $f_{i}\left(x_{i}\right)=y_{i}$. The honest user can measure $\left|\psi_{y_{i}}\right\rangle$ to find a pre-image of $y_{i}$ to pass this test with certainty. In the equation challenge, $C_{i}=1$, the user must provide a non-zero string $d_{i} \in\{0,1\}^{w}$ such that $d_{i} \cdot\left(x_{i} \oplus x_{i}^{\prime}\right)=0$. The bank can test whether the equation challenge holds by using the trapdoor $t_{i}$ to calculate both $x_{i}$ and $x_{i}^{\prime}$. An honest user can generate such a string by measuring $\left|\psi_{y_{i}}\right\rangle$ in the Hadamard basis, as described in Eq. (1). The measured $d_{i}$ will be non-zero (except with probability exponentially small in $w$ ) which will allow the user to pass this test.

We emphasize that for both the minting and the verification protocols, the bank only needs a classical computer.

We now outline the security argument. Suppose the user tries to pass verification twice. Denote by $C \in\{0,1\}^{n}$ the challenge vector in the first attempt, denote by $C^{\prime}$ the challenge vector in the second attempt, and denote by $S$ the set of coordinates in which they differ: $S=\left\{i \in[n] \mid C_{i} \neq\right.$ $\left.C_{i}^{\prime}\right\}$. With overwhelming probability, $S \neq \varnothing$, in which case for at least one coordinate the user will have to pass both challenges, and cannot succeed except with negligible probability.

The construction above is a semi-quantum 2-of-2 mini-scheme (rather than a full blown scheme). There is a slightly stronger notion of security (that is still weaker than a full blown scheme) called a mini-scheme (adapted from Aaronson and Christiano [AC13]). In a mini-scheme, the counterfeiter is given a single quantum money state and can attempt to pass verification polynomially many times. The counterfeiter succeeds if at least two of these verifications are accepted. We show in Section 5.2 that the scheme above also achieves this stronger notion.

[^3]In a full quantum money scheme, the adversary can ask for $t$ money states and must pass at least $t+1$ verifications. Aaronson and Christiano AC13] defined the notion of a public money mini-scheme and showed how such a mini-scheme can be lifted to a full-blown scheme. Ben-David and Sattath BS16a showed a similar result that lifts a private money mini-scheme to a full-blown scheme. In this work, we show how to lift an interactive private money mini-scheme to a full-blown scheme. The goal of such a mapping is to ensure that the scheme can support the issuance of multiple money states without increasing the key-size. This is done by using an authenticated encryption scheme for the mini-scheme keys and including that authenticated ciphertext as part of the money. As part of the verification, the bank can later decrypt the mini-scheme key, and use it to run the original mini-scheme verification. It is important that the encryption scheme be authenticated to prevent the adversary from altering that information (which would be possible if, for example, the encryption scheme was malleable).

Related works. The security of private quantum money schemes is generally solid, Wie83, MVW13, PYJ+12, TOI03, MS10, Gav12, GK15, JLS18. Secure public quantum money is much harder to construct. The constructions of Aaronson $\mathrm{Aar09}$ ] was broken in $\mathrm{LAF}^{+} 10$, and the construction of Aaronson-Christiano [AC13] was broken using several approaches - see the most recent attack in Ref. $\left[\mathrm{PDF}^{+} 19\right]$ and references therein. In Ref. BS16a, a construction based on quantum-secure indistinguishability obfuscation (IO) was presented, as well as a mechanism to provide classical verifiability. Zhandry later proved that the quantum money is indeed secure Zha19, though Zhandry's proof does not lend itself to the classical verifiability construction by Ben-David and Sattath. The only two constructions that are not known to be broken are by Farhi et al. $\mathrm{FGH}^{+} 12$ ] (see also Lut11), which does not have a security proof, and three constructions by Zhandry [Zha19]: The two quantum lightning constructions discussed above, and another one which proves the security of AaronsonChristiano based on quantum-secure IO. We note that the (classical) security of indistinguishability obfuscation is still ongoing research (see, https: //malb.io/are-graded-encoding-schemes-broken-yet.html for a list of constructions and their security status). As far as the authors are aware, no IO construction claims to be quantum-secure, which is required for Zhandry's scheme. To conclude, the security of public quantum money leaves much to be desired.

The work of [HS20] is somewhat relevant to what we do here, but from a device-independent point of view of private quantum money. There, the bank provides a secret key to an untrusted mint, and the mint produces the quantum banknote itself. An assumption is made on the mint regarding the dimension of the state that it outputs. While HS20 do not use the
definition of a mini-scheme as we do here (Definition 25), their work features a private quantum money construction with mini-scheme security while not proving full scheme security (Definition 24). Their scheme is also classically verifiable. To conclude, the main advantage of their scheme is that it is unconditionally secure, while the main disadvantages are that the scheme does not provide classical minting, there is an additional assumption on the dimension of the output register, and there is no security proof as a full quantum money scheme.

A very recent work of Amos et al. AGKZ20 introduced a public quantum money with completely classical communication. Their construction allows users to transfer money using classical communication only, requiring no interaction with a bank. However, the security is proved only with respect to an oracle.

A recent work of ACGH19 also shows a parallel repetition theorem, though their result is slightly different, and their proof techniques are completely different.

Modes of Operation. To fully grasp the implications of semi-quantum money, it is important to review how quantum money in its different flavors can be used, and the communication requirements (classical or quantum) of each mode of operation. In Section 2 we provide a detailed overview of the different flavors of quantum money and their respective communication requirements, showing the advantage of semi-quantum money: transactions with purely classical communication.

Our contribution. Our contribution is twofold. The first is semi-quantum money, both private and public: new models of quantum money that allow transactions with no quantum communication whatsoever (see Figs. 12 and 13), and only a classical bank. The main advantage of the new schemes compared to previous quantum money schemes is that they could be used without a quantum communication infrastructure. Classical communication has several interesting benefits over quantum communication. The most obvious one is that a classical communication infrastructure already exists; a semi-quantum money scheme - unlike previous money schemes - will not require a quantum communication infrastructure. Implementing such an infrastructure on a global scale will be expensive and challenging, and might be realized years after efficient quantum computers are commonly used. There are other benefits to classical communication, even if quantum communication infrastructure was readily available. First, due to the nocloning theorem, quantum information cannot be re-sent. In the context of quantum money, data-loss is extremely problematic - data loss means lost money. Quantum communication will naturally suffer more data-loss, at least initially. Second, for classical communication we can keep a record
(and even a signed record) which helps with matters of dispute resolution, auditing and error-handling, whereas quantum communication cannot be logged. The same argument can be made for the banks themselves; classical banks could more easily keep records and be audited.

The second contribution is the parallel repetition theorem for 1-of-2 puzzles (described earlier in the introduction). Parallel repetition (the idea of repeating a protocol polynomially many times in parallel to gain an exponential increase in soundness) seems deceptively simple, while in reality, it sometimes behaves in unexpected ways, and such proofs are usually challenging (see Raz11 and references therein for the non-cryptographic setting); BIN97 present several cases where parallel repetition surprisingly does not grant an exponential reduction in error rate in cryptographic-settings. The parallel repetition theorem for 1-of-2 puzzles could be useful in other cryptographic settings, as it builds on the TCF primitive to introduce a tool with an exponentially small error rate (rather than the constant error rate which is guaranteed in the original work).

Organization. Section 2 provides an overview on the different ways in which different flavors of quantum money can be used, along with the communication requirements of each mode of operation.

In Section 3 we deal with our proposed public semi-quantum money. Section 3.1 contain the relevant definitions, and in Section 3.2 we construct the public scheme and prove its security, proving our main public theorem, namely, Theorem 1 .

In Section 4, we deal with NTCF and 1-of-2 puzzles. In Section 4.1, we define a 1 -of-2 puzzle. In Section 4.2 we show a construction of a $\frac{1}{2}$-hard 1-of-2 puzzle based on an NTCF. We conclude Section 4 by showing, in Section 4.3, a method for constructing a strong 1-of-2 puzzle using repetition of weak 1-of-2 puzzles.

In Section 5 , we deal with our proposed private semi-quantum money. Section 5.1 contains the relevant definitions, which are adaptations of the definitions from Section 3.1 to the private setting. In Section 5.2, we construct a semi-quantum money mini-scheme and prove its security. In Section 5.3, we present a full semi-quantum money scheme construction based on any semi-quantum mini-scheme, and prove its security.

In Section 6 we combine the results of Sections 4 and 5 to prove our main private result, namely, Theorem 2. A structural overview of our private semi-quantum money result is shown in Fig. 1.

In Section 7 we show that a quantum money scheme with classical minting (and therefore any semi-quantum money scheme) cannot be informationtheoretically secure (i.e., it must rely on computational assumptions).

Appendix A is a nomenclature, a "cheat sheet" describing some of our notations. Appendix $B$ contains mainly the standard definitions of private-


Figure 1: Structure of our private scheme construction. The right-hand side of the figure shows our assumptions. The arrows point to constructions that make use of these assumptions.
key encryption and message authentication code (MAC), and can be safely skipped by readers who are familiar with these notions. Appendix C, taken almost verbatim from [CS20], comprises the definitions of quantum lightning with bolt-to-certificate. Appendix D, taken almost verbatim from Brakerski et al. $\overline{\mathrm{BCM}^{+} 18}$, comprises a definition of NTCF. Appendix E comprises a discussion of memoryless vs memory-dependent schemes. Appendix F is a simplified overview of the parallel repetition of weakly verifiable puzzles result from [CHS05] which is used for our parallel repetition of 1-of-2 puzzles result. Appendix $G$ comprises a series of figures illustrating the various ways transactions can be performed with different flavors of quantum money schemes (regular, classically-verifiable and semi-quantum) along with the types of communication each method and flavor requires.

## 2 Modes of Operation

This section discusses the different methods in which quantum money can be used. When discussing a transaction, we are interested in three parties: the payer $P$ who has money; the receiver $R$ who should receive $P$ 's money by the end of the transaction; and the bank $B$, a trusted third party. When discussing regular money, we can think of 2 different modes to execute a transaction: a direct transaction, or through the bank. Table 1 shows the communication requirements for a direct transaction and a transaction through the bank for regular (non-quantum) money.

One example of a direct transaction is online credit card payments (see Fig. 2). In this case $P$ transfers her credit card details to $R$. $R$ then contacts $B$ with $P$ 's credit card details and $B$ gives him the payment (after verifying $P$ 's credit card details). This mode requires communication between $P$ and $R$ and between $R$ and $B$.

One example of a transaction through the bank is a wire transfer (see Fig. 3). In this case $P$ contacts the bank requesting to transfer money to $R$. Bank then contacts $R$ to complete the transaction (deducting money from $P$ 's account and crediting $R$ 's account with the payment) and sends a message alerting $R$ to the transaction. This mode requires communication between $P$ and $B$ and between $R$ and $B$.

|  | Direct |  | Through Bank |  |
| :--- | :---: | :---: | :---: | :---: |
| Regular Money | $\mathcal{P} \rightarrow \mathcal{R}$ |  |  |  |
| $\mathcal{B}^{\swarrow}$ | (Fig. 22) | $\mathcal{P}_{\mathcal{B}}^{\mathcal{P}} \quad$ (Fig. 3) |  |  |

Table 1: Types of communication required for a direct transaction and a transaction through the bank with regular (non-quantum) money. $\rightarrow$ denotes classical communication.

We make the observation that different transaction methods exist for quantum money as well. In the quantum money setting this discussion is much more delicate, since there are various flavors of quantum money, and there is the question of which types of communication (classical or quantum) are required between which parties for any transaction method. An important distinction is that quantum money is held in the form of quantum states on the user's quantum computer rather than in an account kept by the bank.

The straightforward way to think of a quantum money transaction is a direct transfer in which $P$ sends $R$ the banknote directly. Afterwards, $R$ needs to verify the validity of the money, since $P$ is not trusted. In a public scheme this can be done locally on $R$ 's computer, while in a private scheme $R$ must interact with $B$ for this purpose.

There is another way to transfer quantum money: a transaction through
the bank. In such a transaction, instead of sending the money directly to $R, P$ sends the money to $B . B$ then validates the money, and proceeds to send it to $R$ (or mint a new banknote for him). This type of transaction has the following advantage: in a direct transfer, when $P$ sends her money to $R, R$ may claim that the money did not pass verification and refuse the payment. In this scenario there is usually no way for $P$ to prove that she gave $R$ valid money. This is not the case, however, when performing the transaction through $B$; since the bank $B$ is assumed to be trusted for the money to function properly, any claim made by $R$ that he did not receive valid money can be refuted by $B$ who took part in the transaction and made sure the money was valid. Along with this advantage, transactions via $B$ offer different communications requirements which could prove beneficial. These two benefits make it is interesting to review bank transactions even in the setting of public quantum money.

When discussing transactions it is interesting to talk about traceability; i.e., can any party learn any information regarding transactions that have been made between other parties? Generally, quantum money comes in 2 flavors regarding traceability: banknotes and coins. The difference is that while quantum banknotes are unique (each banknote has a serial numbers), while quantum coins are indistinguishable. Thus quantum banknotes are inherently traceable, while the traceability of coins depend on how they are used (see AMR20] for a formal definition). Apart from the differences between transaction methods discussed above, a direct transaction and a transaction through the bank also behave differently when considering traceability of quantum coins. In a direct transfer $R$ learns the identity of the $P$ and $B$ learns only the identity of $R$. In a transaction through the bank, however, $R$ does not learn the identity of $P$ and $B$ learns the identities of both $P$ and $R$. This difference gives further motivation to consider the different modes of transaction.

Table 2 lists the communication requirements of transactions for each flavor of quantum money, and Appendix $G$ contains detailed figures illustrating how transactions can be made for each flavor of quantum money. Quantum communication is denoted by $\Rightarrow$, classical communication is denoted by $\rightarrow$, and a two-sided arrow indicates two-sided communication is required.

We can see that semi-quantum money allows transactions with purely classical communication. It is an open question whether a variant of semiquantum money could be made that would allow to classically transfer quantum money between users without the aid of the bank (i.e., make a public direct transaction with classical communication alone). Amos et al. AGKZ20 showed a construction of semi-quantum money with completely classical communication between users (one-shot signature), but in the oracle model.

Remark 3. In Table 2 we state that a direct transfer using classically verifiable or semi-quantum money requires only classical communication between $P$ and $R$. This is because $R$ can initiate the verification protocol with the bank, and then act as relay between $P$ and $B$ in order to verify $P$ 's banknote. The bank then mints a new banknote for receiver to conclude the transaction (see Figs. 6 and 12). Of course, it is still possible for the $P$ to send the banknote via quantum communication, but this way requires one less quantum channel.
Remark 4. We use a stronger security notion than that used in Refs. Gav12, GK15. In these previous works a quantum money scheme is defined to be unforgeable if an adversary with $\ell$ banknotes cannot generate more than $\ell$ quantum states that would each pass the honest verification (with nonnegligible probability). In contrast, our unforgeability requirement (Definitions 8 and 24 states that an adversary with $\ell$ banknotes cannot pass more than $\ell$ potentially malicious verifications. This is more general and therefore stronger; in the existing definition the verification protocol is run honestly, while our definition allows the adversary to act maliciously during the verification protocol itself.

It is easy to see why the stronger definition is not used in the schemes above. These schemes are designed to work in the direct mode of operation, and they provide the ability to verify a banknote multiple times, meaning the scheme does not meet the stronger security definition by design. Generally, such schemes could not utilize their classical verification for a transaction through the bank; since verification does not destroy the banknote, $P$ would have to send her banknote to $B$ (through a quantum channel) so $B$ could be sure she no longer holds the banknote. This means that transactions through the bank would require all-quantum communication.

Allowing multiple verifications is beneficial for the classical verification setting, but irrelevant for the semi-quantum setting. For example, in Gavinsky's scheme [Gav12], the banknote is designed to allow a (fixed) polynomial number of verifications, after which the banknote would be sent back to the bank to be verified and destroyed, and a new one would be minted instead. If instead the banknote was destroyed after a single verification, the verifier would need the bank to mint him a new banknote after every verification, which would require quantum communication. Therefore, allowing multiple verifications without the need to re-mint allows verification with no quantum communication between $R$ and $B$ (most of the time). In the semi-quantum setting, however, minting is also classical, so allowing banknotes to pass verification without being destroyed does not give us an interesting benefit, and we can use our stronger security notion. This also allows us to use the through the bank mode of operation with only classical communication.

|  |  | Direct | Through Bank |
| :---: | :---: | :---: | :---: |
| Private | Standard | $\mathcal{P} \Rightarrow \mathcal{R}$ <br> $\mathcal{B}^{\mathbb{Z}}$ <br> (Fig. 4 ) | $\mathcal{P}_{\Downarrow_{\mathcal{B}}} \mathbb{R}^{\mathcal{R}} \quad \text { (Fig. } 5$ |
|  | Classically verifiable | $\underset{\mathcal{B}^{\mathcal{Z}}}{\underset{\sim}{\boldsymbol{Z}}} \underset{\sim}{\leftrightarrow}$ <br> (Fig. 6) | $\stackrel{P}{\mathcal{B}}_{\delta_{\mathcal{B}}}^{\mathcal{R}} \quad \text { (Fig. } 7$ |
|  | Reusable classically verifiable ${ }^{5}$ | $\begin{gathered} \mathcal{P} \Rightarrow \mathcal{B}_{\mathcal{Z}}^{\mathcal{R}} \end{gathered}$ <br> (Fig. $\qquad$ |  |
|  | Classical minting ${ }_{6}^{6}$ | $\underset{\underset{\mathcal{B}}{\mathscr{Z}} \underset{\sim}{\mathcal{X}}}{\underset{\mathcal{R}}{\mathcal{R}}}$ <br> (Fig. 10 ) | $\mathcal{P}_{\mathbb{B}_{\mathcal{Z}}} \quad \text { (Fig. } 11$ |
|  | Semi-quantum | $\underset{\mathcal{B}}{\mathcal{P}} \underset{\mathcal{Z}}{\leftrightarrow} \quad \text { (Fig. } 12$ | $\mathcal{S}_{\mathcal{B}}^{\mathcal{Z}} \quad \text { (Fig. } 13$ |
| Public | Standard | $\begin{gathered} \mathcal{P} \Rightarrow \mathcal{R} \\ \mathcal{B} \end{gathered} \quad \text { (Fig. } 14$ | $\mathcal{P}_{\Downarrow_{\mathcal{B}}}$ <br> (Fig. 5) |
|  | Classically verifiable ${ }^{7}$ | $\begin{gathered} \mathcal{P} \Rightarrow \mathcal{R} \\ \mathcal{B} \end{gathered} \quad \text { (Fig. } 14$ | $\mathcal{P}_{\Sigma_{\mathcal{B}}} \mathcal{N}^{\mathcal{R}} \quad \text { (Fig. } 7 \text { ) }$ |
|  | Reusable classically verifiable | $\begin{gathered} \mathcal{P} \Rightarrow \mathcal{R} \\ \mathcal{B} \end{gathered} \quad \text { (Fig. } 14$ |  |
|  | Classical minting ${ }^{8}$ | $\begin{gathered} \mathcal{P} \Rightarrow \mathcal{R} \\ \mathcal{B} \end{gathered} \quad \text { (Fig. } 14$ | $\underset{\mathcal{B}^{\mathcal{Z}}}{\mathcal{R}} \quad \text { (Fig. } 11$ |
|  | Semi-quantum | $\mathcal{P} \Rightarrow \mathcal{R}$ <br> (Fig. 14 ) | $\mathcal{P}_{\mathcal{S}_{\mathcal{B}}} \mathcal{R}^{\mathcal{R}} \quad \text { (Fig. 13 }$ |
|  | One-shot signatur\& ${ }^{9}$ | $\mathcal{P} \leftrightarrow \mathcal{R}$ <br> (Fig. 15 | $\mathcal{P}_{\mathcal{S}_{\mathcal{B}}^{\mathcal{R}}} \quad \text { (Fig. } 13$ |

Table 2: Types of communication required for each transaction method in each flavor of quantum money. $\Rightarrow$ denotes quantum communication and $\rightarrow$ denotes classical communication. A two-sided arrow indicates two-sided communication. Note that public schemes can also use transactions methods from private schemes.

[^4]
## 3 Quantum Lightning with Bolt-to-Certificate Implies Public Semi-Quantum Money

In this section we construct a public memory-dependent semi-quantum money scheme using Zhandry's quantum lightning from Zha19 along with Coladangelo's Bolt-to-Certificate from [Col19] and its superseding work [CS20].

### 3.1 Definitions of Public Semi-Quantum Money

Definition 5 (Interactive public quantum money). An interactive public quantum money scheme consists of a classical PPT key generation algorithm key-gen and two-party interactive QPT protocols mint and verify. key-gen ( $1^{\lambda}$ ) outputs a pair of keys $(p k, s k)$ which are the public and private keys, respectively. Both the minting protocol and the verification protocol are two-party quantum protocols: mint involves an Acquirer $A$ and a Bank B, whereas verify involves a Payer $P$ and a Receiver $R$ (in verify either party can be either a bank or a user). During both protocols, both parties receive the public key $p k$ as input, and the bank (if it participates) receives the private key sk while users do not. At the end of the honest run of mint, the user holds a quantum money state that, in general, could be a mixed state. In this work, the protocols will end with a pure state, usually denoted $|\$\rangle$. In the following sections, for the sake of clarity, we work with the pure-state formalism. The banknote that the $P$ chooses to verify is denoted in this work as the input of the verify protocol. At the end of the verification protocol, the Receiver outputs a bit b that states whether the money was accepted or not.

Correctness. The scheme is correct if there exists a negligible function $\operatorname{negl}(\lambda)$ such that:

$$
\begin{gathered}
\operatorname{Pr}\left[(p k, s k) \leftarrow \operatorname{key}-\operatorname{gen}\left(1^{\lambda}\right) ;|\$\rangle \leftarrow \operatorname{mint}_{(p k, s k)}\left(1^{\lambda}\right):\right. \\
\left.\operatorname{verify}_{p k}(|\$\rangle)=1\right]=1-\operatorname{negl}(\lambda) .
\end{gathered}
$$

Definition 6 (Memory-Dependent Quantum Money). A quantum money scheme is memory-dependent if the bank is required to maintain a state it uses throughout different runs of the quantum money protocols.

Definition 7 (Security against sabotage). An interactive public quantum money scheme $\$$ is secure against sabotage if for every QPT counterfeiter $\mathcal{A}$ there exists a negligible function negl $(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{SABOTAGE}^{-M O N E Y} \mathcal{A}_{\mathcal{A}}^{\$}(\lambda)=1\right] \leq \operatorname{negl}(\lambda) .
$$

The money sabotaging game SABOTAGE-MONEY $_{\mathcal{A}, \mathcal{S}}(\lambda)$ :

[^5]1. The challenger runs $(p k, s k) \leftarrow$ key-gen $\left(1^{\lambda}\right)$ and sends $p k$ to $\mathcal{A}$.
2. The adversary outputs a money state $|\$\rangle$ and runs verify $(|\$\rangle)$ with the challenger two consecutive times.
3. $\mathcal{A}$ wins if the first verification accepts and the second rejects, in which case the result of the game is 1 (and otherwise it is 0).

Definition 8. We say that an interactive public quantum money scheme $\$$ is secure if it is secure against sabotage (Definition 7) and if for every QPT counterfeiter $\mathcal{A}$ there exists a negligible function $\operatorname{neg}(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{COUNTERFEIT}_{\mathcal{A}, \$}(\lambda)=1\right] \leq \operatorname{negl}(\lambda)
$$

The money counterfeiting game COUNTERFEIT $_{\mathcal{A}, \$}(\lambda)$ :

1. The bank generates a key pair $(p k, s k) \leftarrow$ key-gen $\left(1^{\lambda}\right)$ such that $p k$ is publicly known.
2. The bank and the counterfeiter interact. The counterfeiter can ask the bank to run $\operatorname{mint}_{k}(\cdot)$ and verify ${ }_{k}(\cdot)$ polynomially many times, in any order the counterfeiter wishes. The counterfeiter is not bound to following his side of the protocols honestly. The counterfeiter can keep ancillary registers from earlier runs of these protocols and use them in later steps. Let $w$ be the number of successful verifications, $\ell$ the number of times that mint was called by the counterfeiter and $v$ the number of times that verify was called by the counterfeiter.
3. The bank outputs $(w, \ell, v)$.

The value of the game is 1 iff $w>\ell$. In this case we sometimes simply say that the counterfeiter wins.

We state here again that this security definition is stronger than the one previously used - this is further explained in Remark 4 .

Public semi-quantum money has two types of verifications: the first is a quantum verification algorithm we denote querify that does not destroy the banknote and can be performed between any two entities with quantum computation and communication resources, and the second is a classical verification protocol we denote cverify that destroys the money and can only be performed with a bank. querify allows money transfer between any two users, which could be performed without the aid of a bank. cverify can be thought of as depositing or spending the money with the bank - a quantum user can destroy his banknote in a way that he can prove to the bank that the note was destroyed. In the counterfeiting game above, the counterfeiter can choose which type of verification to run each time. Note that if querify
is chosen, the bank must perform the verification; i.e., the adversary must give the banknote to the bank (and the bank does not return it afterwards).

Note that the bank is no different to a user when participating in querify, since no secret information is used there - so if qverify is secure with a bank, it is also secure with other users (i.e., any counterfeiting method that would work against a user would also work against the bank). Therefore the above definition of security ensures the scheme is secure even when considering transactions between users, even though the security game is defined with the bank.

Note that if two users wish to transfer a banknote but do not share a quantum communication channel, they can do so via a bank: the Payer would perform cverify with the bank, after which the bank will run mint with the Receiver (see Fig. 13).

Definition 9 (Public semi-quantum money). We say that a protocol with the bank has classical minting (resp. verification) if the bank is classical in mint (resp. verify). We define public semi-quantum money as any secure interactive public quantum money scheme that has classical minting and two types of verification, denoted qverify and cverify, such that in the end of qverify the banknote is not destroyed (meaning it could pass further verifications), and that cverify is a classical verification which destroys the banknote and can only be performed between a user and a bank. In querify no party receives the private key, and in cverify the bank receives the private key while the user does not.

### 3.2 Construction of a Public Semi-Quantum Scheme

Following is an informal explanation of our public scheme construction, which is defined formally in Algorithm 1. The construction uses a secure quantum lightning scheme with bolt-to-certificate - a brief informal overview of both is given in Section 1 below Theorem 1, and their formal definitions (taken almost verbatim from their respective papers) can be found on Appendix C.

Let QL be a secure quantum lightning scheme with bolt-to-certificate (see Definitions 39 to 41 , and let DS be a PQ-EU-CMA digital signature scheme (see Definitions 36 and 37 ). In key-gen the bank runs the QL setup algorithm to randomly generate a set of algorithms that generate and verify bolts and certificates, and runs DS.key-gen to generate a private and public key. key-gen outputs the public digital signature key and the four QL algorithms as the scheme's public key, and the digital signature's private key as the scheme's private key. For the minting process, the user generates a bolt $|\psi\rangle$ along with its serial key $s$ using the generation algorithm gen-bolt (that is part of the public key). He sends $s$ to the bank, which sends back a signature $\sigma$ for it (using its private key). Each banknote $|\$\rangle$ consists of the bolt $|\psi\rangle$, its
serial number $s$, and its signature $\sigma$ (without the signature any user could generate by themselves notes that would pass verification).

In qverify the Payer sends to the Receiver his bolt along with its signed serial number, and the receiver verifies both the bolt and the signature. In cverify the user uses the bolt to generate a certificate and sends the certificate to the bank along with the bolt's signed serial number. The bank verifies the certificate and the signature, and checks its database for the serial number. If the serial number appears there then the certificate has been given before and so verification will fail.

Note that according to Definition 9, in a general public semi-quantum money scheme the bank uses the secret key $s k$ in cverify, but in our scheme the bank does not - meaning it is stronger in this respect. BS16a also have classical verification with only a public key, so this is not the first instance where cverify uses only the public key.

Proposition 10 (Correctness of $\$_{P}$ ). Assuming QL is a secure quantum lightning scheme with bolt-to-certificate (see Definitions 39 and 41) and DS is a digital signature scheme with perfect completeness (see Definition 36), $\$_{P}$, which is defined in Algorithm 1, is a correct (Definition 5) memorydependent (Definition 6) public semi-quantum money scheme (see Definition 9).

Proof. Let (gen-bolt, verify-bolt, gen-certificate, verify-certificate) $\leftarrow$ QL.setup $\left(1^{\lambda}\right)$ and assume an honest run of $|\$\rangle \leftarrow \$_{P} \cdot \operatorname{mint}_{(p k, s k)}()$. We need to prove correctness for both implementations of verify: we begin by proving the correctness of qverify. In an honest run of $\$_{P}$.qverify, the banknote sent on line 1 is $|\$\rangle$ generated in mint. Recall that $|\$\rangle:=(|\psi\rangle, s, \sigma)$. From the perfect completeness of DS we get:

$$
\operatorname{Pr}\left[\mathrm{DS} . \text { verify }_{p k_{\sigma}}\left(s, \mathrm{DS} . \operatorname{sign}_{s k_{\sigma}}(s)\right)=1\right]=1
$$

So the signature verification on line 2 passes since $\sigma \leftarrow \mathrm{DS} . \operatorname{sign}_{s k_{\sigma}}(s)$ on line 3 of $\$_{P}$.mint, meaning $\operatorname{Pr}\left[r_{\sigma}=1\right]=1$. From the definition of quantum lightning (Definition 39) we get:

$$
\begin{aligned}
\operatorname{Pr}[(\text { gen-bolt, verify-bolt }) & \leftarrow \mathrm{QL} . \operatorname{setup}\left(1^{\lambda}\right) ;(|\psi\rangle, s) \leftarrow \text { gen-bolt }(): \\
& \text { verify-bolt }(|\psi\rangle, s)=1] \\
& =1-\operatorname{negl}(\lambda)
\end{aligned}
$$

So the bolt verification on line 3 passes except with negligible probability since $(|\psi\rangle, s)$ is a valid bolt generated on line 1 of $\$_{P}$.mint, meaning $\operatorname{Pr}\left[r_{b}=\right.$ $1]=1-\operatorname{neg} I(\lambda)$. Therefore, from the union bound:

$$
\begin{aligned}
\operatorname{Pr}[\text { QL.qverify } & p k \\
(|\psi\rangle)=0] & \leq \operatorname{Pr}\left[r_{\sigma}=0\right]+\operatorname{Pr}\left[r_{b}=0\right] \\
& =\operatorname{negl}(\lambda)
\end{aligned}
$$

```
Algorithm 1 The Interactive Public Money Scheme \(\$_{P}\)
\(\$_{P}\). key-gen \(\left(1^{\lambda}\right)\)
    : (gen-bolt, verify-bolt, gen-certificate, verify-certificate) \(\leftarrow\) QL.setup \(\left(1^{\lambda}\right)\)
    \(\left(p k_{\sigma}, s k_{\sigma}\right) \leftarrow\) DS.key-gen \(\left(1^{\lambda}\right)\)
    \(p k \leftarrow\) (gen-bolt, verify-bolt, gen-certificate, verify-certificate, \(p k_{\sigma}\) )
    \(s k \leftarrow s k_{\sigma}\)
    return \((p k, s k)\)
\(\$_{p} \cdot\) mint \(_{(p k, s k)}\)
Acquirer Bank
    1: \(\quad(|\psi\rangle, s) \leftarrow\) QL.gen-bolt \(\left(1^{\lambda}\right)\)
    2 :
    3: \(\quad \sigma \leftarrow \mathrm{DS}^{\operatorname{sign}}{ }_{s k_{\sigma}}(s)\)
    4 :
        \(\stackrel{\sigma}{\longleftrightarrow}\)
    \(5: \quad|\$\rangle \leftarrow(|\psi\rangle, s, \sigma)\)
\(\$_{p}\).qverify \(_{p k}(|\$\rangle)\)
    1: interpret \(|\$\rangle\) as \((|\psi\rangle, s, \sigma)\)
    \(r_{\sigma} \leftarrow\) DS.verify \(_{p k_{\sigma}}(s, \sigma)\)
    \(r_{b} \leftarrow\) QL.verify-bolt \((|\psi\rangle, s)\)
    return \(r_{\sigma} \cdot r_{b}\)
\(\frac{\$_{p} . \text { cverify }_{p k}(|\$\rangle)}{\triangleright D \leftarrow \varnothing \text { before first run }}\)
Payer
                                    Bank
    1: interpret \(|\$\rangle\) as \((|\psi\rangle, s, \sigma)\)
    2: \(\quad c \leftarrow\) QL.gen-certificate \((|\psi\rangle, s)\)
    3: \(\xrightarrow{s, \sigma, c}\)
    4: \(\quad r_{\sigma} \leftarrow\) DS.verify \(_{p k_{\sigma}}(s, \sigma)\)
    5:
    \(r_{c} \leftarrow\) QL.verify-certificate \((s, c)\)
    6 :
    \(r_{d} \leftarrow s \notin D\)
    7:
    if \(r_{\sigma} \cdot r_{c} \cdot r_{d}=1\) :
        \(D \leftarrow D \cup\{s\}\)
    return \(r_{\sigma} \cdot r_{c} \cdot r_{d}\)
```

meaning $\operatorname{Pr}\left[\right.$ QL.querify $\left.{ }_{p k}(|\psi\rangle)=1\right]=1-\operatorname{neg}(\lambda)$.
We now prove correctness of $\$_{P}$.cverify. Like with querify, from the perfect completeness of DS we get that the signature verification on line 4 passes, meaning $\operatorname{Pr}\left[r_{\sigma}=1\right]=1$. From the definition of bolt-to-certificate (Definition (41) we get that the certificate verification on line 5 passes except with negligible probability, since $c$ was generated on line 2 using the valid bolt $(|\psi\rangle, s)$ that was generated in line 1 of $\$_{P}$.mint, meaning $\operatorname{Pr}\left[r_{c}=1\right]=1-$ negl $(\lambda)$. From the security of quantum lightning (Definition 40) we get that for each $s^{\prime} \in D, \operatorname{Pr}\left[s=s^{\prime}\right]$ is negligibl ${ }^{10}$, meaning $\operatorname{Pr}\left[r_{d}=1\right]=1-\operatorname{neg}(\lambda)$ assuming D is polynomial in $\lambda$. Therefore, from the union bound:

$$
\begin{aligned}
\operatorname{Pr}^{[Q L . c v e r i f y} & p k \\
(|\psi\rangle)=0] & \leq \operatorname{Pr}\left[r_{\sigma}=0\right]+\operatorname{Pr}\left[r_{c}=0\right]+\operatorname{Pr}\left[r_{d}=0\right] \\
& =\operatorname{negl}(\lambda)
\end{aligned}
$$

Theorem 11 (Security of $\$_{P}$ ). Assuming QL is a secure quantum lightning scheme (Definitions 39 and 40) with bolt-to-certificate capability (Definition 41) and that DS is a $P Q-E U-C M A$ digital signature scheme (Definitions 36 and 37), then $\$_{P}$, which is defined in Algorithm 1, is secure according to Definition 8. If QL is secure against sabotage (Definition 7), $\$_{P}$ is also secure against sabotage.

Proof. In order to prove the security of the scheme, we need to prove that no counterfeiter $\mathcal{A}$ can win the security game COUNTERFEIT $\mathcal{A}_{\mathcal{A}, \$_{P}}(\lambda)$ (Definition 8) with non-negligible probability. In the game, the counterfeiter has access to a verification oracle, meaning that in our case the counterfeiter can run either querify or cverify a polynomial amount of times, and wins if he manages to pass a total of more than $\ell$ querify and cverify verifications (when $\ell$ is the number of times mint was run).

We show that any adversary capable of breaking the security of $\$_{P}$ must break the underlying security of either QL, QL's bolt-to-certificate capability, or DS. Assume a QPT counterfeiter $\mathcal{A}$ with non-negligible success probability. Recall that $w, \ell$ and $v$ are the numbers of successful verifications, runs of $\$_{P}$.mint and runs of $\$_{P}$.verify in the counterfeiting security game, respectively ( $w$ and $v$ include runs of qverify and of cverify). This means that $w>\ell$ with non negligible probability. We assume $w$ and $v$ are polynomial in $\lambda$, otherwise the counterfeiter would not be QPT. Denote by $\left|\$_{j}\right\rangle=\left(\left|\psi_{j}\right\rangle, s_{j}, \sigma_{j}\right)$ the banknote minted in the $j^{t h}$ run of $\$_{P} \cdot m i n t{ }^{11}$. Due

[^6]to the unforgeability of DS (Definition 37), for every successful verification, $\mathcal{A}$ either sent $|\$\rangle=\left(|\psi\rangle, s_{j}, \sigma\right)$ for some $j \in[\ell],|\psi\rangle, \sigma$ on line 1 of querify or $s_{j}, \sigma, c$ for some $j \in[\ell], \sigma, c$ on line 3 of cverify ${ }^{12}$. Denote by $w_{j}$ the amount of successful verifications - either qverify or cverify - made with $s_{j}$ as input. By the pigeonhole principle, since there are only $\ell s_{j}$ 's but $w>\ell$ successful verifications, $w_{i} \geq 2$ for some $i \in[\ell]$. For that $i$, there are three possibilities: either the two verifications were querify, the two verifications were cverify, or there was one of each.

Two successful cverify runs with the same $s_{i}$ are not possible, since after each successful run of cverify the bank adds $s_{i}$ to $D$, and every subsequent verifications with $s_{i} \in D$ fails. One successful run of querify and one successful run of cverify are also not possible - if it were, we could construct a certificate forger $\mathcal{C}$ with non-negligible probability: $\mathcal{C}$ will simulate the bank and run $\mathcal{A}$ against it. With non-negligible probability, he will succeed both on a run of querify and on a run of cverify: meaning he possesses a quantum state $|\psi\rangle$ such that QL.verify-bolt $\left(|\psi\rangle, s_{i}\right)=1$ and a certificate $c$ such that QL.verify-certificate $\left(s_{i}, c\right)=1$, in contradiction to the bolt-to-certificate capability of QL (Definition 41).

The only option left is two successful runs of querify: in which case $\mathcal{A}$ is in possession of two quantum states $\left|\psi_{1}\right\rangle,\left|\psi_{2}\right\rangle$ such that QL.verify-bolt $\left(\left|\psi_{1}\right\rangle, s_{j}\right)=$ 1 and QL.verify-bolt $\left(\left|\psi_{2}\right\rangle, s_{j}\right)=1$ (he could not have passed the two verifications with the same quantum state since querify entails sending the quantum state to the bank - hence he necessarily possess two such states). We could then construct a quantum lightning adversary $\mathcal{L}$ that simulates the bank and runs $\mathcal{A}$ against it. With non-negligible probability, he will end up with some $\left|\psi_{1}\right\rangle$ and $\left|\psi_{2}\right\rangle$ that pass QL.verify-bolt with the same serial number $s_{i}$, and could use them to win $\mathrm{FORGE}^{-\mathrm{BOLT}_{\mathcal{L}, \mathrm{QL}}(\lambda) \text { (see Definition 40), in }}$ contradiction to the security of QL.

None of the three options are possible, meaning that any counterfeiter $\mathcal{A}$ has negligible success probability, i.e., $\$_{P}$ is secure.

Security against sabotage of QL directly implies security against sabotage of $\$_{P}$.

For convenience, we restate the main theorem of our public semi-quantum money result:

[^7]Theorem 1 (Public Semi-Quantum Money). Assuming the existence of a secure Quantum Lightning scheme (Definitions 39 and 40) with bolt-tocertificate capability (Definition 41), and the existence of a $P Q-E U-C M A$ digital signature scheme (Definitions 36 and 37), then a secure memorydependent (Definition 6) public semi-quantum money scheme exists (Definition 9).

We can see that Theorem 11 proves our main public result.

## 4 Trapdoor Claw-Free Families and 1-of-2 Puzzles

In this section, as the name suggests, we discuss the concepts of NTCF and 1-of-2 puzzles. For completeness, we restate the formal definition of NTCF by Brakerski et al. in Appendix D. In Section 4.1, we introduce 1-of-2 puzzles. In Section 4.2 we show how to construct a 1-of-2 puzzle using an NTCF, and in Section 4.3 we show a parallel repetition theorem for 1-of-2 puzzles that is subsequently used to construct strong 1-of-2 puzzles.

### 4.1 1-of-2 Puzzles

Definition 12 (1-of-2 puzzle). A 1-of-2 puzzle consists of four efficient algorithms: the puzzle generator $G$, an obligation algorithm $O$, a 1-of-2 solver $S$, and a verification algorithm $V . G$ is a classical algorithm, $V$ is a classical deterministic algorithm, and $O$ and $S$ are quantum algorithms.

1. $G$ outputs, on security parameter $1^{\lambda}$, a random puzzle $p$ and some verification key $v:(p, v) \leftarrow G\left(1^{\lambda}\right)$.
2. $O$ receives a puzzle $p$ as input and outputs a classical string o called the obligation and a quantum state $\rho:(o, \rho) \leftarrow O(p)$.
3. $S$ receives $p, o, \rho$ and $a$ bit $b \in\{0,1\}$ as input and outputs an answer string $a: a \leftarrow S(p, o, \rho, b)$.
4. $V$ receives $p, v, o, b, a$ as input and outputs 0 or $1: V(p, v, o, b, a) \in$ $\{0,1\}$.

Completeness: Let $\eta$ be some arbitrary function $\eta: \mathbb{N} \mapsto[0,1]$. We say that the 1-of-2 puzzle has completeness $\eta$ if there exists a negligible function negl $(\lambda)$ such that

$$
\begin{align*}
\operatorname{Pr}[(p, v) & \leftarrow G\left(1^{\lambda}\right),(o, \rho) \leftarrow O(p), b \leftarrow\{0,1\}, a \leftarrow S(p, o, \rho, b): \\
& V(p, v, o, b, a)=1]  \tag{2}\\
& \geq \eta(\lambda)-\operatorname{negl}(\lambda)
\end{align*}
$$

We define the $V_{2}$ algorithm as:

$$
\begin{equation*}
V_{2}\left(p, v, o, a_{0}, a_{1}\right)=V\left(p, v, o, 0, a_{0}\right) \cdot V\left(p, v, o, 1, a_{1}\right) . \tag{3}
\end{equation*}
$$

Hardness: Let $h: \mathbb{N} \mapsto[0,1]$ be an arbitrary function. We say that the 1-of-2 puzzle $\mathcal{Z}$ is 1 -h-hard if for any QPT 2-of-2 solver $\mathcal{T}$ there exists a negligible function negl $(\lambda)$ such that

$$
\begin{equation*}
\operatorname{Pr}\left[\operatorname{SOLVE}-2_{\mathcal{T}, \mathcal{Z}}(\lambda)=1\right] \leq h(\lambda)+\operatorname{negl}(\lambda) . \tag{4}
\end{equation*}
$$

The 2-of-2 solving game SOLVE- $2_{\mathcal{T}, \mathcal{Z}}(\lambda)$ :

1. The puzzle giver runs $(p, v) \leftarrow G\left(1^{\lambda}\right)$
2. The 2-of-2 solver $\mathcal{T}$ receives input $p$ and outputs a triple ( $o, a_{o}, a_{1}$ )
3. The puzzle giver runs $r \leftarrow V_{2}\left(p, v, o, a_{0}, a_{1}\right)$ and outputs $r$
4. $\mathcal{T}$ wins the game if and only if $r=1$, in which case the output of the game is defined to be 1 .

We say that the 1 -of-2 puzzle is strong if $\eta=1$ and $h=0$ (i.e., the puzzle is 1 -hard). We say that the 1-of-2 puzzle is weak if $\eta=1$ and $1-h$ is noticeable.

### 4.2 An NTCF Implies a 1-of-2 Puzzle

This section presents how an NTCF can be used to construct a 1-of-2 puzzle. The formal definition of an NTCF and its properties used in this section can be found in Appendix D] taken from $\left[\mathrm{BCM}^{+} 18\right]$.

Theorem 13. An NTCF implies a 1-of-2 puzzle with completeness $\eta=1$ and hardness $h=\frac{1}{2}$.

Note that the 1 -of- 2 puzzle above is a weak 1 -of- 2 puzzle.
Proof. The proof contains arguments similar to those used by Brakerski et al. $\left[\mathrm{BCM}^{+} 18\right]$.

Given an NTCF family $\mathcal{F}$ that consists of the algorithms

$$
\text { key- } \operatorname{gen}_{\mathcal{F}}, \operatorname{Inv}_{\mathcal{F}}, \mathrm{CHK}_{\mathcal{F}}, \mathrm{SAMP}_{\mathcal{F}}, J_{\mathcal{F}}
$$

we construct the 1-of-2 puzzle $\mathcal{Z}=\left(\right.$ key-gen $\left._{\mathcal{Z}}, O_{\mathcal{Z}}, S_{\mathcal{Z}}, V_{\mathcal{Z}}\right)$ as specified in Algorithm 2.

Completeness: we need to show that Eq. (2) holds for $\mathcal{Z}$ defined above. By NTCF property 3, the state $\left|\psi^{\prime}\right\rangle$ in line 2 of the algorithm $O_{\mathcal{Z}}$ is negligibly close in trace distance to:

$$
|\tilde{\psi}\rangle=\frac{1}{\sqrt{2|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}, b \in\{0,1\}} \sqrt{\left(f_{k, b}^{\prime}(x)\right)(y)}|b, x\rangle|y\rangle .
$$

```
Algorithm 2 The 1-of-2 Puzzle \(\mathcal{Z}\)
    procedure key-gen \(\mathcal{Z}_{\mathcal{Z}}(\lambda)\)
        \(\left(k, t_{k}\right) \leftarrow\) key \(-\operatorname{gen}_{\mathcal{F}}(\lambda)\)
        Set \(p \equiv k, v \equiv t_{k}\)
        return \((p, v)\)
    end procedure
    procedure \(O_{\mathcal{Z}}(p)\)
        \(\left|\psi^{\prime}\right\rangle \leftarrow \operatorname{SAMP}_{\mathcal{F}}(p,|+\rangle)\)
        Measure the last register to obtain an \(o \in \mathcal{Y}\). Denote the post-
    measurement state \(|\psi\rangle \quad \triangleright\) In the completeness we show that
    \(|\psi\rangle \approx \frac{1}{\sqrt{2}}\left(|0\rangle\left|x_{0}\right\rangle+|1\rangle\left|x_{1}\right\rangle\right)\).
        return \((o,|\psi\rangle)\)
    end procedure
    procedure \(S_{\mathcal{Z}}(p, o,|\psi\rangle, b) \triangleright p\) and \(o\) are not used in this construction.
        if \(b=0\) then
            Measure both registers of \(|\psi\rangle\) to obtain a result \(i \in\{0,1\}\) and
    \(x \in \mathcal{X}\).
            Set \(a \equiv(i, x)\)
        else if \(b=1\) then
            Evaluate the function \(J\) on the second register of \(|\psi\rangle\), and apply
    Hadamard transform on both registers.
            Measure both registers to obtain the result \(i \in\{0,1\}\) and \(d\).
            Set \(a \equiv(i, d)\)
        end if
        return \(a\)
    end procedure
    procedure \(V_{\mathcal{Z}}(p, v, o, b, a)\)
        Set \(x_{0} \equiv \operatorname{INV}_{\mathcal{F}}(v, 0, o)\) and \(x_{1} \equiv \operatorname{INV}_{\mathcal{F}}(v, 1, o) \triangleright\) Recall that \(v\) is the
    trapdoor, and \(o\) is an image of the NTCF.
        if \(b=0\) then
            Interpret \(a\) as \(i, x\)
            if \(x=x_{i}\) then
                    return 1
            else
                    return 0
            end if
        else if \(b=1\) then
            Interpret \(a\) as \(i, d\).
            if \(d \in G_{p, 0, x_{0}} \cap G_{p, 1, x_{1}}\) and \(d \cdot\left(J\left(x_{0}\right) \oplus J\left(x_{1}\right)\right)=i\) then \(\triangleright\) This
    membership test uses \(\mathrm{CHK}_{\mathcal{F}}\)
                return 1
            else
                return 0
            end if
        end if
    end procedure
```

For the sake of the analysis, therefore, we can replace $|\psi\rangle$ with $|\tilde{\psi}\rangle$, and the algorithm will behave the same, up to a negligible probability. By NTCF property 2 b , the post-measurement state $|\psi\rangle$ generated by $O_{\mathcal{Z}}$ is $|\psi\rangle=$ $\frac{1}{\sqrt{2}}\left(|0\rangle\left|x_{0}\right\rangle+|1\rangle\left|x_{1}\right\rangle\right)$, where $\left(x_{0}, x_{1}\right) \in \mathcal{R}_{p}$. Since $o$ was the outcome of the measurement in line 3, we know that $o \in \operatorname{Supp} f_{p, i}\left(x_{i}\right)$. By NTCF property 2 a , for $i \in\{0,1\}$ :

$$
\begin{equation*}
x_{i}=\operatorname{INV}_{\mathcal{F}}(v, i, o) . \tag{5}
\end{equation*}
$$

Consider the case $b=0$. In this case, the output of $S_{\mathcal{Z}}$ is $a \equiv\left(i, x_{i}\right)$, where, by Eq. (5), $x_{i}=\operatorname{INV}_{\mathcal{F}}(v, i, o)$. Therefore, $V_{\mathcal{Z}}$ will return 1 in line 6 . In the case of $b=1$, before line 6 in $S_{\mathcal{Z}}$ the state is $\frac{1}{\sqrt{2}}\left(|0\rangle\left|x_{0}\right\rangle+|1\rangle\left|x_{1}\right\rangle\right)$, after the evaluation of $J$ on the second register the state is $\frac{1}{\sqrt{2}} \sum_{j \in\{0,1\}}|j\rangle\left|J\left(x_{j}\right)\right\rangle$, and after the Hadamard on both registers (which consist of $w+1$ qubits), the state is

$$
\begin{aligned}
& \frac{1}{\sqrt{2^{w+2}}} \sum_{i \in\{0,1\}, d \in\{0,1\}^{w}}\left(\sum_{j \in\{0,1\}}(-1)^{i j+d \cdot J\left(x_{j}\right)}\right)|i\rangle|d\rangle \\
= & \frac{1}{\sqrt{2^{w}}} \sum_{d \in\{0,1\}^{w}}(-1)^{d \cdot J\left(x_{0}\right)}\left|d \cdot\left(J\left(x_{0}\right) \oplus J\left(x_{1}\right)\right)\right\rangle|d\rangle .
\end{aligned}
$$

Therefore, the outcome of the measurement in line 7 will provide a random $d \in\{0,1\}^{w}$ and an $i \in\{0,1\}$ that satisfy $i=d \cdot\left(J\left(x_{0}\right) \oplus J\left(x_{1}\right)\right)$. Since $d$ is random, property 4 a guarantees that the first condition in line 12 of $V_{\mathcal{Z}}$ will be met (up to a negligible probability), and the analysis in the previous sentence guarantees that the second condition will be met. Overall, the probability that $V_{\mathcal{Z}}$ outputs 1 is $1-\operatorname{negl}(\lambda)$, for some negligible function negl, as required.

Soundness: We need to show that Eq. (4) holds for every QPT $\mathcal{T}$. In Algorithm 3, we show a reduction that maps a 2 -of- 2 solver $\mathcal{T}$ for the 1-of- 2 puzzle as in Eq. (4) to an NTCF adversary $\mathcal{A}$ as in Eq. (9).

```
Algorithm 3 The Adversary \(\mathcal{A}\)
    procedure \(\mathcal{A}_{\mathcal{F}}(k)\)
        \(\left(o, a_{0}, a_{1}\right) \leftarrow \mathcal{T}(k)\)
        Interpret \(a_{0}\) as \(i, x\) and \(a_{1}\) as \(i^{\prime}, d\).
        return \(\left(i, x, d, i^{\prime}\right)\)
    end procedure
```

If $\mathcal{T}$ succeeds with probability $\frac{1}{2}+\epsilon(\lambda)$ (where $\epsilon(\lambda)$ is not necessarily negligible), then the l.h.s. in Eq. (9) is lower-bounded by $2 \epsilon(\lambda)$ with respect to $\mathcal{A}$. Plugging the definition of $V_{2}$ (see Eq. (3)) and the acceptance criteria of $V_{\mathcal{Z}}$ into lines 6 and 12 , we see that the 2 -of- 2 solver
$\mathcal{T}$ needs to find $o, i, x, d, i^{\prime}$ such that $d \in G_{p, 0, x_{0}} \cap G_{p, 1, x_{1}}$ and $x=x_{i}$, where $x_{0}=\operatorname{INV}_{\mathcal{F}}(v, 0, o), x_{1}=\operatorname{INV}_{\mathcal{F}}(v, 1, o)$ and $i^{\prime}=d \cdot\left(J\left(x_{0}\right) \oplus J\left(x_{1}\right)\right)$. This implies the membership of ( $i, x, d, i^{\prime}$ ) in $H_{k}$ (see Eq. (88). Therefore, $\operatorname{Pr}_{\left(k, t_{k}\right) \leftarrow \operatorname{key}_{-\operatorname{gen}_{\mathcal{F}}\left(1^{\lambda}\right)}\left[\mathcal{A}(k) \in H_{k}\right] \geq \frac{1}{2}+\epsilon(\lambda) \text {. Since } H_{k} \text { and } \vec{H}_{k} \text { are disjoint, }}$ $\operatorname{Pr}_{\left(k, t_{k}\right) \leftarrow \text { key-gen }}^{\mathcal{F}\left(1^{\lambda}\right)}\left[\mathcal{A}(k) \in \bar{H}_{k}\right] \leq \frac{1}{2}-\epsilon(\lambda)$, and

$$
\begin{aligned}
& \mid \operatorname{Pr}_{\left(k, t_{k}\right) \leftarrow \text { key-gen }}^{\mathcal{F}\left(1^{\lambda}\right)} \\
& {\left[\mathcal{A}(k) \in H_{k}\right]-\operatorname{Pr}_{\left(k, t_{k}\right) \leftarrow \text { key-gen }}^{\mathcal{F}\left(1^{\lambda}\right)} } \\
& \geq 2 \epsilon(\lambda) .
\end{aligned}
$$

Since by property 2b (Definition 42) the l.h.s. of Eq. (9) is upperbounded by the negligible function $\mu(\lambda)$, we conclude that $\epsilon(\lambda)$ must be negligible, as required for $h=\frac{1}{2}$.

### 4.3 A Parallel Repetition Theorem for 1-of-2 Puzzles

Definition 14 (Parallel repetition of 1 -of-2 puzzles). Let $\mathcal{Z}$ be a 1-of-2 puzzle system, and let $n \in \mathbb{N}$. We denote by $G^{n}$ the algorithm that, on security parameter $\lambda$, runs $G\left(1^{\lambda}\right)$ for $n(\lambda)$ times and outputs all the $n$ puzzles with their verification keys:

$$
\begin{equation*}
\left.\left(\left(p_{1}, \ldots, p_{n}\right)\right),\left(v_{1}, \ldots, v_{n}\right)\right) \leftarrow G^{n}\left(1^{\lambda}\right) \tag{6}
\end{equation*}
$$

where $\left(p_{i}, v_{i}\right) \leftarrow G\left(1^{\lambda}\right)$. A similar approach is used for all other algorithms in $\mathcal{Z}$ :

$$
\begin{equation*}
\left.\left(\left(o_{1}, \ldots, o_{n}\right)\right),\left(\rho_{1} \otimes \cdots \otimes \rho_{n}\right)\right) \leftarrow O^{n}\left(p_{1}, \ldots, p_{n}\right) \tag{7}
\end{equation*}
$$

where $\left(o_{i}, \rho_{i}\right) \leftarrow O\left(p_{i}\right)$.

$$
\left(a_{1}, \ldots, a_{n}\right) \leftarrow S^{n}\left(\left(p_{1}, \ldots, p_{n}\right),\left(o_{1}, \ldots, o_{n}\right), \rho_{1} \otimes \cdots \otimes \rho_{n}, b\right)
$$

where $a_{i} \leftarrow S\left(p_{i}, o_{i}, \rho_{i}, b\right)$. The algorithm

$$
V^{n}\left(\left(p_{1}, \ldots, p_{n}\right),\left(v_{1}, \ldots, v_{n}\right),\left(o_{1}, \ldots, o_{n}\right), b,\left(a_{1}, \ldots, a_{n}\right)\right)
$$

outputs 1 iff for all $i \in[n], V\left(p_{i}, v_{i}, o_{i}, b, a_{i}\right)=1$.
The $n$-fold parallel repetition of $\mathcal{Z}$ is the 1 -of-2 puzzle

$$
\mathcal{Z}^{n}=\left(G^{n}, O^{n}, S^{n}, V^{n}\right) .
$$

We emphasize that $\mathcal{Z}^{n}$ is a 1 -of- 2 puzzle (and not a 1 -of $-2^{n}$ puzzle), which explains why the algorithm contains a single challenge bit $b$ rather than $n$ bits. The reason for that should be made clear later - see Fact 19

It is natural to ask what is the hardness parameter of $\mathcal{Z}^{n}$, relative to the hardness of $\mathcal{Z}$. In the settings of 2-prover games: on the one hand, the
soundness parameter decreases exponentially [Raz98; yet the exponential decrease is not as fast as one would expect Raz11. We say the parallel repetition is perfect if the fact that $\mathcal{Z}$ is $1-h$-hard means that $\mathcal{Z}^{n}$ is $1-h^{n}$ hard. Note that a perfect parallel repetition means that the adversary can do no better than solving each 1 -of- 2 puzzle independently (up to a negligible correction). This is indeed the case for 1 -of- 2 puzzles:

Theorem 15 (Perfect parallel repetition of 1-of-2 puzzles). Let $\mathcal{Z}$ be a 1-of-2 puzzle with completeness $\eta$ and hardness parameter $h$. For a function $n(\lambda)$ that satisfies $n(\lambda)=\operatorname{poly}(\lambda)$, the 1-of-2 puzzle $\mathcal{Z}^{n}$ has completeness $\eta^{n}$ and hardness parameter $h^{n}$.

Proof. First we prove the completeness property (see Eq. (22). For ease of notation, we write $n$, negl, $\eta$,etc., instead of $n(\lambda)$, negl $(\lambda), \eta(\lambda)$. Suppose that the success probability of $\mathcal{Z}$ is $\eta$ - negl for some negligible function negl. Since the repeated game $\mathcal{Z}^{n}$ is an independent repetition of $\mathcal{Z}$, its success probability is $(\eta-\text { negl })^{n}$. We show that for the negligible function negl' $=n^{2} \operatorname{negl}(\lambda)$, indeed $(\eta-\text { negl })^{n} \geq \eta^{n}-\operatorname{negl}^{\prime}:$

$$
\begin{aligned}
(\eta-\text { negl })^{n} & =\eta^{n}+\sum_{k=1}^{n}(-1)^{k}\binom{n}{k} \eta^{n-k} \text { neg }^{k} \\
& \geq \eta^{n}-\sum_{k=1}^{n} n^{k} \text { neg }^{k} \\
& \geq \eta^{n}-\sum_{k=1}^{n} n \cdot \text { negl }=\eta^{n}-\text { negl }^{\prime},
\end{aligned}
$$

where the last inequality holds for all $\lambda \geq \lambda_{0}$ (where $n \cdot$ negl $\leq 1$ ).
We are now ready to prove the soundness. Our main tool is the notion of a weakly verifiable puzzle system defined by Canetti, Halevi and Steiner:

Definition 16 (A weakly verifiable puzzle, adapted from [CHS05]). A system for weakly verifiable puzzles consists of a pair of efficient classical algorithms $\hat{\mathcal{Z}}=(G, V)$ such that

1. The puzzle generator $G$ outputs, on security parameter $\lambda$, a random puzzle $p$ along with some verification information $v:(p, v) \leftarrow G\left(1^{\lambda}\right)$.
2. The puzzle verifier $V$ is a deterministic efficient classical algorithm that, on input of a puzzle $p$, verification key $v$, and answer $a$, outputs either zero or one: $V(p, v, a) \in\{0,1\}$.

The hardness of a weakly verifiable puzzle is defined as follows:

Definition 17 (Hardness of a weakly verifiable puzzle, adapted from CHS05). Let $h: \mathbb{N} \mapsto[0,1]$ be an arbitrary function. A weakly verifiable puzzle $\hat{\mathcal{Z}}$ is said to be $1-h$-hard if, for any QPT ${ }^{13}$ algorithm $S$, there exists a negligible function $\operatorname{negl}(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{SOLVE}_{\mathcal{S}, \hat{\mathcal{Z}}}(\lambda)\right] \leq h(\lambda)+\operatorname{negl}(\lambda)
$$

The event $\operatorname{SOLVE}_{\mathcal{S}, \hat{\mathcal{E}}}(\lambda)$ is defined by the following security game:

1. The puzzle giver runs $(p, v) \leftarrow G\left(1^{\lambda}\right)$.
2. The solver $\mathcal{S}$ is given input $p$ and outputs an answer $a$.
3. The puzzle giver runs $r \leftarrow V(p, v, a)$. The event $\operatorname{SOLVE}_{\mathcal{S}, \hat{\mathcal{Z}}}(\lambda)$ is when $r=1$.

To avoid confusion, we always use $\mathcal{Z}$ to denote a 1 -of- 2 puzzle and $\hat{\mathcal{Z}}$ to denote a weakly verifiable puzzle.

Definition 18 (Repetition of weakly verifiable puzzles, from CHS05). Let $\hat{\mathcal{Z}}=(G, V)$ be a weakly verifiable puzzle system, and let $n: \mathbb{N} \mapsto \mathbb{N}$ be some function. We denote by $G^{n}$ the algorithm that, on security parameter $\lambda$, runs $G\left(1^{\lambda}\right)$ for $n(\lambda)$ times and outputs all the $n$ puzzles with their respective verification keys:

$$
\left.\left(\left(p_{1}, \ldots, p_{n}\right)\right),\left(v_{1}, \ldots, v_{n}\right)\right) \leftarrow G^{n}\left(1^{\lambda}\right)
$$

where $\left(p_{i}, v_{i}\right) \leftarrow G^{n}\left(1^{\lambda}\right)$. $V^{n}$ receives $n$ inputs and accepts if and only if all $n$ runs of $V$ accept:

$$
V^{n}\left(\left(p_{1}, \ldots, p_{n}\right),\left(v_{1}, \ldots, v_{n}\right),\left(a_{1}, \ldots, a_{n}\right)\right) \equiv \prod_{i=1}^{n(\lambda)} V\left(p_{i}, v_{i}, a_{i}\right)
$$

There is a tight relation between the hardness of a 1 -of- 2 puzzle and the hardness of a weakly verifiable puzzle. Given a 1 -of- 2 puzzle $\mathcal{Z}=$ ( $G, O, S, V$ ), we define the weakly verifiable puzzle

$$
\hat{\mathcal{Z}}=\left(G, V_{2}\right)
$$

(where $V_{2}$ is defined in Eq. (3)).
Fact 19. For every polynomially bounded function $n: \mathbb{N} \mapsto \mathbb{N}$, the 1-of2 puzzle $\mathcal{Z}^{n}$ is $1-h$-hard if and only if the weakly verifiable puzzle $\hat{\mathcal{Z}}^{n}$ is 1 - h-hard.

[^8]This fact follows from the observation that the hardness property of the 1-of-2 puzzle $\mathcal{Z}$ is equivalent to the hardness of the weakly verifiable puzzle $\hat{\mathcal{Z}}$ (see Definitions 16 and 12 . Furthermore, the hardness of $\mathcal{Z}^{n}$ is equivalent to the hardness of $\mathcal{Z}^{n}$ (see Definitions 14 and 18 ).

Canetti, Halevi and Steiner proved a parallel repetition theorem for weakly verifiable puzzles.

Theorem 20 ([CHS05, Theorem 1]). Let $h: \mathbb{N} \mapsto[0,1]$ be an efficiently computable function, let $n: \mathbb{N} \mapsto \mathbb{N}$ be efficiently computable and polynomially bounded, and let $\hat{\mathcal{Z}}=(G, V)$ be a weakly verifiable puzzle system. If $\hat{\mathcal{Z}}$ is $1-h$-hard, then $\hat{\mathcal{Z}}^{n}$, the $n$-fold repetition of $\hat{\mathcal{Z}}$, is $1-h^{n}$-hard.

Although the original proof of Canetti, Halevi and Steiner assumed that the hardness is with respect to a classical solver, the result holds also when we consider our definition, in which the solvers are quantum. The reason is as follows. Their proof maps an efficient solver of the n-fold repetition of a puzzle, which succeeds with probability which is non-negligibly higher than $1-h^{n}$, to an efficient solver that succeeds with probability non-negligibly higher than $1-h$ for a single puzzle, which is of course a contradiction. This reduction is black-box, and in particular there is no rewinding (which, of course, could cause an issue in the quantum setting) ${ }^{14}$. For completeness, a sketch of their proof is provided in Appendix F.

We use Theorem 20 to prove the soundness of the 1 -of- 2 puzzle $\mathcal{Z}^{n}$. We assume $\mathcal{Z}=(G, O, S, V)$ is $1-h$-hard. We define the weakly verifiable puzzle $\hat{\mathcal{Z}}=\left(G, V_{2}\right)$. By the equivalence in Fact 19 , we know that $\hat{\mathcal{Z}}$ is also $1-h$-hard. By Theorem 20, we know that $\hat{\mathcal{Z}}^{n}$ is $1-h^{n}$-hard. Using the equivalence in Fact 19 again, we conclude that $\mathcal{Z}^{n}$ is $1-h^{n}$-hard, which completes the proof.

Corollary 21. A weak 1-of-2 puzzle implies a strong 1-of-2 puzzle.
Note that we define a weak 1-of- 2 puzzle to have completeness $\eta=1$. We refrain from answering the question whether any puzzle in which $\eta(\lambda)-h(\lambda)$ is noticeable, implies a strong puzzle.

Proof. By using Theorem 15 with $n(\lambda)=\frac{\log ^{2}(\lambda)}{\log \left(\frac{1}{h}\right)}$ repetitions $\sqrt{15}$ of the weak $h$-hard 1-of-2 puzzle, we construct a 1 -complete ${ }^{16} 1-h^{n}=1-\frac{1}{\lambda^{\log (\lambda)}}=$ $1-\operatorname{negl}(\lambda)$-hard 1-of-2 puzzle. Note that a $1-\operatorname{neg}(\lambda)$-hard 1-of-2 puzzle is equivalent to a 1 -hard 1 -of-2 puzzle, which completes the proof.

[^9]
## 5 Strong 1-of-2 Puzzles Imply Semi-Quantum Money

In this section, we show a construction of a private semi-quantum money scheme using strong 1-of-2 puzzles.

In Section 5.1, we define interactive private quantum money. We define three degrees of security. Full scheme security means that every QPT counterfeiter cannot pass $t+1$ verifications given $t$ quantum money states. We define mini-scheme security as a weaker variant of full security, which is secure only when the adversary is given a single banknote. Finally, we define 2 -of- 2 mini-scheme security as an even weaker variant wherein the adversary does not have a banknote verification oracle. We also formally define semi-quantum money.

In Section 5.2, we show the construction of a 2 -of- 2 mini-scheme, and show that our 2-of-2 mini-scheme is in fact a mini scheme (see Definition 25).

In Section 5.3, we show that any (interactive private quantum money) mini scheme can be elevated to a full (interactive private quantum money) scheme - see Definition 24,

### 5.1 Definitions of Private Semi-Quantum Money

The following definitions are slight variations of the definitions in Section 3.1.
Definition 22 (Interactive memoryless private quantum money). An interactive memoryless private quantum money scheme consists of a classical PPT key generation algorithm key-gen and two-party interactive memoryless QPT protocols mint and verify. key-gen $\left(1^{\lambda}\right)$ outputs a key $k$. Both the minting protocol and the verification protocol are two-party quantum protocols involving an Acquirer $A$ and a Bank B. During both protocols, the bank receives the key $k$ as input, and the user does not. At the end of the honest run of mint, the user holds a quantum money state that, in general, could be a mixed state. In this work, the protocols will end with a pure state, usually denoted $|\$\rangle$. In the following sections, for the sake of clarity, we work with the pure-state formalism. The banknote the user chooses to verify is denoted in this work as the input of the verify protocol. At the end of the verification protocol, the bank outputs a bit b that states whether the money is valid or not.

Correctness. The scheme is correct if there exists a negligible function negl $(\lambda)$ such that:

$$
\begin{aligned}
& \operatorname{Pr}\left(k \leftarrow \text { key-gen }\left(1^{\lambda}\right) ;|\$\rangle \leftarrow \operatorname{mint}_{k}\left(1^{\lambda}\right) ; b \leftarrow \operatorname{verify}_{k}(|\$\rangle):\right. \\
& \quad b=1)=1-\operatorname{neg}(\lambda)
\end{aligned}
$$

Definition 23. We say that the protocol has classical minting (verification) if $B$ is classical in mint (verify). To emphasize that the verification is classical, we use cverify to denote the (classical) verification algorithm. We define private semi-quantum money as any secure memoryless interactive private quantum money protocol that has classical minting and classical verification.

In the quantum setting, there are a number of possible verifications with different qualities; a notable quality is whether the verification "destroys" the banknote (i.e., whether the banknote can be used again after verification). This distinction can be thought of as the difference between verifying - proving that a legal money state exists - and spending - proving a legal money state does not exist - and it becomes more interesting when considering the public setting. There, a banknote can be spent with the bank in the same manner as in the private setting, but it can also be verified with other users. In such a case it is important that the banknote is preserved, so it could be transferred. Another distinction is added by the introduction of classically verified money: whether the verification is a classical or quantum protocol. Moreover, a classical verification must be a challenge-response protocol - otherwise the same proof can be passed twice, effectively spending the same banknote twice. In our scheme, verification is classical and does not preserve the banknote, proving both that it existed and that it does not exist anymore.

In this definition, we emphasize that the protocols mint and verify are memoryless: i.e., all outgoing messages depend solely on the key and the input from the user. In other words, the bank does not maintain a variable state that changes between different runs of the protocols - each run is independent. Constructing a stateful scheme is trivial even in the classical setting, as discussed in Appendix E. In addition, it is interesting to note that our protocols are composed of a fixed number of messages, independent of the security parameter: verify has 2 messages (a single round) and mint has 3 messages.

Definition 24. We say that an interactive private quantum money scheme $\$$ is secure if for every QPT counterfeiter $\mathcal{A}$ there exists a negligible function negl $(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{COUNTERFEIT}{ }_{\mathcal{A}, \$}^{\text {full }}(\lambda)=1\right] \leq \operatorname{negl}(\lambda)
$$

The money counterfeiting game COUNTERFEIT $\mathcal{A}_{\mathcal{A}, \$}^{\text {full }}(\lambda)$ :

1. The bank generates a key $k \leftarrow$ key-gen $\left(1^{\lambda}\right)$.
2. The bank and the counterfeiter interact. The counterfeiter can ask the bank to run $\operatorname{mint}_{k}(\cdot)$ and verify ${ }_{k}(\cdot)$ polynomially many times, in
any order the counterfeiter wishes. The counterfeiter is not bound to following his side of the protocols honestly. The counterfeiter can keep ancillary registers from earlier runs of these protocols and use them in later steps. Let $w$ be the number of successful verifications, $\ell$ the number of times that mint was called by the counterfeiter and $v$ the number of times that verify was called by the counterfeiter.
3. The bank outputs $(w, \ell, v)$.

The value of the game is 1 iff $w>\ell$. In this case we sometimes simply say that the counterfeiter wins.

We state here again that this security definition is stronger than the one previously used - this is further explained in Remark 4 .

Following previous works AC13, GK15, we define a private quantum money mini-scheme, with a slight deviation. Additionally, we define a 2 -of- 2 mini-scheme, which is a weaker variant of the mini-scheme.

Definition 25 (quantum money mini-scheme and 2-of-2 mini-scheme). We define mini-scheme security as we defined full scheme security but with regard to COUNTERFEIT $\mathcal{B}, \$_{\operatorname{mini}}(\lambda)$, wherein the counterfeiter $\mathcal{B}$ wins iff $w>$ $\ell \wedge \ell=1$.

We define 2-of-2 mini-scheme security as we did above but with regard to COUNTERFEIT $_{\mathcal{C}, \$}^{2-o f-2}(\lambda)$, where the counterfeiter $\mathcal{C}$ wins iff $w>\ell \wedge \ell=$ $1 \wedge v=2$.

Note that the definitions in this sections could be naturally extended to the public settings.

### 5.2 Construction of a Mini-Scheme

In this section, we show the construction of a scheme that we then prove to be a 2 -of- 2 semi-quantum mini-scheme. Later we prove that our construction in fact achieves a stronger security notion - a semi-quantum mini-scheme.

We now give an informal description of our construction, which is defined formally in Algorithm 4. The construction uses a strong 1-of-2 puzzle and a post-quantum existentially unforgeable under an adaptive chosen-message attack (PQ-EU-CMA) MAC (see Definitions 12, 34 and 35). In key-gen, the bank generates a MAC signing key and $n$ pairs of strong 1-of-2 puzzles and their respective verification keys. The minting process is done as follows. The bank sends these $n$ puzzles to the user, who then runs the obligation protocol $\mathcal{Z} . O$ on all the $n$ puzzles. The user keeps the quantum output of $O$ and sends the classical outputs (called the obligations) to the bank. The bank signs these obligations using the classical MAC scheme and sends these tags back to the user. The verification starts with the bank sending random challenges to the user. The user then has to present a set of signed
obligations (which the user should have from the mint protocol) together with a set of solutions to the challenges of these puzzles. The bank verifies the solution to each puzzle with its respective verification key (the set of verification keys is part of the key). Due to the fact that this verification is classical, it is denoted cverify. We show that a counterfeiter cannot doublespend a banknote without breaking the soundness of a strong 1-of-2 puzzle (or the security of the MAC).

Intuitively, an adversary could try to double-spend the banknote using the solutions he received from the first verification, while hoping to be given the same challenges. However, assuming a sufficiently large number of puzzles (say, $n=\log ^{2}(\lambda)$ ), the probability of encountering the exact same set of challenges more than once is negligible. Passing two verifications of any banknote in which the challenges were not the same both times essentially requires one to pass the SOLVE -2 security game for the 1 -of- 2 puzzle. Insofar as this is considered a strong 1-of-2 puzzle, the probability that it can occur is therefore negligible.

For ease of notation, we write:

- $\bar{p}^{n}:=\left(p_{1}, \ldots, p_{n}\right)$
- $\bar{v}^{n}:=\left(v_{1}, \ldots, v_{n}\right)$
- $\bar{o}^{n}:=\left(o_{1}, \ldots, o_{n}\right)$
- $\bar{\psi}^{n}:=\left|\psi_{1}\right\rangle \otimes \cdots \otimes\left|\psi_{n}\right\rangle$
- $\bar{b}^{n}:=\left(b_{1}, \ldots, b_{n}\right)$
- $\bar{a}^{n}:=\left(a_{1}, \ldots, a_{n}\right)$

Proposition 26 (Correctness of $\$ \mathcal{Z}$ ). Assuming MAC has perfect completeness and $\mathcal{Z}$ is a 1-of-2 puzzle with completeness $\eta=1, \$_{\mathcal{Z}}$ (Algorithm 4) is a semi-quantum money scheme that satisfies the correctness property (see Definition 22).

Proof. Clearly, the communication and the bank's operation in mint and cverify are classical - therefore, the scheme is semi-quantum.

From the perfect completeness property of MAC (see Definition 34) we get:

$$
\operatorname{Pr}\left[M A C \cdot \operatorname{verify}{ }_{k}\left(\bar{o}^{n}, M A C \cdot \operatorname{mac}_{k}\left(\bar{o}^{n}\right)\right)=1\right]=1
$$

meaning $\operatorname{Pr}\left[r_{M A C}=1\right]=1$.
From the completeness $\eta=1$ of $\mathcal{Z}$ we get:

$$
\begin{aligned}
\operatorname{Pr}[(p, v) & \leftarrow \mathcal{Z} \cdot G(\lambda) ;(o,|\psi\rangle) \leftarrow \mathcal{Z} . O(p) ; b \leftarrow\{0,1\} \\
& a \leftarrow \mathcal{Z} \cdot S(p, o,|\psi\rangle, b): \\
& \mathcal{Z} \cdot v(p, v, o, b, a)=1] \\
& \geq 1-\operatorname{negl}(\lambda)
\end{aligned}
$$

```
Algorithm 4 The Interactive Private Money Scheme \(\$_{\mathcal{Z}}\)
\(\$_{\mathcal{Z}} \cdot\) key-gen \(\left(1^{\lambda}\right)\)
    \(n \leftarrow \log ^{2}(\lambda)\)
    foreach \(i \in[n]\) :
        \(\left(p_{i}, v_{i}\right) \leftarrow \mathcal{Z} . G\left(1^{\lambda}\right)\)
    \(\bar{p}^{n} \leftarrow\left(p_{1}, \ldots, p_{n}\right), \bar{v}^{n} \leftarrow\left(v_{1}, \ldots, v_{n}\right)\)
    \(k \leftarrow M A C\).key-gen \(\left(1^{\lambda}\right)\)
    \(k_{\$} \leftarrow\left(\bar{p}^{n}, \bar{v}^{n}, k\right)\)
    return \(k_{\$}\)
\(\$_{\mathcal{Z}} \cdot\) mint \(_{k_{\$}}\)
Acquirer
                                    Bank
\begin{tabular}{lcc}
\hline \(1:\) & \(\longleftrightarrow \bar{p}^{n}\) \\
\(2:\) & foreach \(i \in[n]:\) & \\
\(3:\) & \(\left(o_{i}, \psi_{i}\right) \leftarrow \mathcal{Z} \cdot O\left(p_{i}\right)\) & \\
\(4:\) & \(\bar{o}^{n} \leftarrow\left(o_{1}, \ldots, o_{n}\right), \bar{\psi}^{n} \leftarrow\left(\psi_{1}, \ldots, \psi_{n}\right)\) & \\
\(5:\) & & \(\bar{o}^{n}\) \\
\(6:\) & & \(t_{o} \leftarrow M A C \cdot \operatorname{mac}_{k}\left(\bar{o}^{n}\right)\) \\
\(7:\) & & \\
& & \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|}
\hline \multicolumn{3}{|l|}{\(\$_{\mathcal{Z} .}\) Cverify \({ }_{k_{\Phi}}\left(\bar{o}^{n}, t_{o}, \bar{\psi}^{n}\right)\)} \\
\hline Acquirer & & Bank \\
\hline 1: & \(\bar{b}^{n} \in_{R}\{0,1\}^{n}\) & \\
\hline 2: foreach \(i \in[n]\) : & & \\
\hline 3: \(\quad a_{i} \leftarrow \mathcal{Z} . S\left(p_{i}, o_{i},\left|\psi_{i}\right\rangle, b_{i}\right)\) & & \\
\hline 4: & \(\bar{a}^{n}, \bar{o}^{n}, t_{o}\) & \\
\hline \(5:\) & &  \\
\hline 6 : & & foreach \(i \in[n]\) : \\
\hline 7: & & \[
r_{i} \leftarrow \mathcal{Z} . V\left(p_{i}, v_{i}, o_{i}, b_{i}, a_{i}\right)
\] \\
\hline 8: & & \[
r \leftarrow r_{M A C} \cdot \prod_{i=1}^{n} r_{i}
\] \\
\hline 9 : & & return \(r\) \\
\hline
\end{tabular}
```

Let $b_{i}$ be the event of failing verification on the $i^{\text {th }}$ puzzle. From the previous equation, $\operatorname{Pr}\left[b_{i}\right] \leq \operatorname{negl}(\lambda)$ for some negligible function negl $(\lambda)$. Let $\operatorname{negl}^{\prime}(\lambda):=n \cdot \operatorname{negl}(\lambda)=\log ^{2}(\lambda) \cdot \operatorname{negl}(\lambda)$. Using the union bound:

$$
\operatorname{Pr}\left[\cup_{i=1}^{n} b_{i}\right] \leq \sum_{i=1}^{n} \operatorname{Pr}\left[b_{i}\right]=\log ^{2}(\lambda) \cdot \operatorname{negl}(\lambda)=\operatorname{negl}^{\prime}(\lambda)
$$

meaning $\operatorname{Pr}\left[\left(\prod_{i=1}^{n} r_{i}\right)=1\right] \geq 1-\operatorname{negl}^{\prime}(\lambda)$. Thus:

$$
\begin{aligned}
\operatorname{Pr} & {\left[k_{\$} \leftarrow \$_{\mathcal{Z}} \cdot \operatorname{key}-\operatorname{gen}\left(1^{\lambda}\right) ;\left(\bar{p}^{n}, \bar{o}^{n}, t_{o}, \bar{\psi}^{n}\right) \leftarrow \$_{\mathcal{Z}} \cdot \operatorname{mint}_{k_{\S}}() ;\right.} \\
& \left.\${ }_{\mathcal{Z}} \cdot \operatorname{cverify}_{k_{\S}}\left(\bar{p}^{n}, \bar{o}^{n}, t_{o}, \bar{\psi}^{n}\right)=1\right] \\
& =\operatorname{Pr}\left[r_{M A C}=1 \bigcap\left(\prod_{i=1}^{n} r_{i}\right)=1\right] \\
& \geq 1-\operatorname{negl}^{\prime}(\lambda) .
\end{aligned}
$$

Proposition 27 ( $\$_{\mathcal{Z}}$ is a 2-of-2 mini-scheme). Assuming $\mathcal{Z}$ is a strong 1-of-2 puzzle and MAC is a PQ-EU-CMA MAC, the scheme $\$_{\mathcal{Z}}$ (Algorithm (4) is a 2-of-2 mini-scheme (see Definition 25).

Proof. We show that the probability of a QPT counterfeiter to win the 2 -of-2 mini-scheme security game against $\$_{\mathcal{Z}}$ (Algorithm (4) is bound by the negligible probability to solve both challenges of the strong 1 -of- 2 puzzle $\mathcal{Z}$. Intuitively, double-spending a banknote entails solving both challenges for at least one of its $n$ puzzles, which is intractable. For this proof, as well as the following security proofs of our money scheme (Proposition 28 and Theorem 30), we use a sequence-of-games based technique adapted from Sho04. The following sequence of games binds the success probability of any QPT 2-of-2 mini-scheme counterfeiter to that of a QPT 2-of-2 puzzle solver (see Eq. (4)):

Game 0. Let $\mathcal{C}$ be a QPT 2-of-2 mini-scheme counterfeiter. We assume w.l.o.g. that $\mathcal{C}$ performs exactly two verifications and one mint (i.e., $\ell=1$ and $v=2$ ) - an adversary which does not comply with this assumption will necessarily fail (see Definition 25). We define Game 0 to be COUNTERFEIT $\mathcal{C}_{\mathcal{C}, \$_{\mathcal{Z}}}^{2-o f-2}(\lambda)$.

Let $S_{0}$ be the event where $w>1$ (see Definition 25) in Game 0 (this is the original win condition for $\mathcal{C}$, since we assume $\ell=1 \wedge v=2$ ).

Game 1. We now transform Game 0 into Game 1 , simply by changing the win condition: Game 1 is identical to Game 0 , but we define the following event: let ${\overline{b^{1}}}^{n},{\overline{b^{2}}}^{n}$ be the random bit strings that were generated in line 1
of $\$_{\mathcal{Z}}$.cverify the first and second times $\mathcal{C}$ asked for verification, respectively. Let $S_{1}$ be the event where $w>1 \wedge{\overline{b^{1}}}^{n} \neq{\overline{b^{2}}}^{n}$ in Game 1.

Let $F$ be the event where ${\overline{b^{1}}}^{n}={\overline{b^{2}}}^{n}$ in Game 1, and $F^{\prime}$ the event where $w>1 \wedge \overline{b^{1}}=\overline{b^{2}}$ in Game 1. Since ${\overline{b^{1}}}^{n}$ and ${\overline{b^{2}}}^{n}$ are generated uniformly and independently, $\operatorname{Pr}[F]=\frac{1}{2^{n}} \leq \operatorname{negl}(\lambda)$ for some negligible function negl $(\lambda)$. Therefore: $\operatorname{Pr}\left[S_{0}\right]=\operatorname{Pr}\left[S_{1} \cup F^{\prime}\right] \leq \operatorname{Pr}\left[S_{1} \cup F\right] \leq \operatorname{Pr}\left[S_{1}\right]+\operatorname{Pr}[F] \leq \operatorname{Pr}\left[S_{1}\right]+$ negl $(\lambda)$. So $\operatorname{Pr}\left[S_{0}\right] \leq \operatorname{Pr}\left[S_{1}\right]+\operatorname{negl}(\lambda)$.

Game 2. We now add a small change to the game above: at the start of the game, a uniform $i^{\prime} \in_{R}[n]$ is chosen by the bank. Let $j$ be the first index such that $b_{j}^{1} \neq b_{j}^{2}\left(j=\infty\right.$ if $\left.b^{1}=b^{2}\right)$.

Let $S_{2}$ be the event where $w>1 \wedge b^{1} \neq b^{2} \wedge i^{\prime}=j$ in Game 2.
$S_{1} \Rightarrow b^{1} \neq b^{2}$, so since $i^{\prime}$ was chosen uniformly and independently of $w$, $b^{1}, b^{2}$ and $j$, we get that $\operatorname{Pr}\left[S_{2} \mid S_{1}\right]=\frac{1}{n}$. Moreover, it is easy to see that $\operatorname{Pr}\left[S_{2} \mid \neg S_{1}\right]=0$. So $\operatorname{Pr}\left[S_{2}\right]=\frac{1}{n} \cdot \operatorname{Pr}\left[S_{1}\right]$, meaning $\operatorname{Pr}\left[S_{1}\right]$ is a polynomial multiplicative factor of $\operatorname{Pr}\left[S_{2}\right]$.

Game 3. Game 3 is identical to Game 2, but we now add an additional constraint to the win condition. Let $\bar{o}^{n}$ be the set of obligations $\mathcal{C}$ sent in line 5 of $\$_{\mathcal{Z}}$.mint, and let ${\overline{o^{1}}}^{n},{\overline{o^{2}}}^{n}$ be the sets of obligations sent by $\mathcal{C}$ during line 4 of $\$_{\mathcal{Z}}$.cverify the first and second times $\mathcal{C}$ asks for verification, respectively.

Let $S_{3}$ be the event where $w>1 \wedge b^{1} \neq b^{2} \wedge i^{\prime}=j \wedge{\overline{o^{1}}}^{n}={\overline{o^{2}}}^{n}=\bar{o}^{n}$ in Game 3.

Let $F$ be the event where $\mathcal{C}$ passes one or more verifications such that ${\overline{o^{1}}}^{n} \neq \bar{o}^{n}$ or ${\overline{o^{2}}}^{n} \neq \bar{o}^{n}$. It is easy to see that $S_{2} \wedge \neg F \Longleftrightarrow S_{3} \wedge \neg F$. Therefore, from the Difference Lemma (Lemma 38) we get

$$
\left|\operatorname{Pr}\left[S_{3}\right]-\operatorname{Pr}\left[S_{2}\right]\right| \leq \operatorname{Pr}[F]
$$

From the unforgeability of MAC (see Definition 35 ), $\operatorname{Pr}[F]$ is negligible ${ }^{17}$ Therefore, $\operatorname{Pr}\left[S_{2}\right] \leq \operatorname{Pr}\left[S_{3}\right]+\operatorname{negl}(\lambda)$.

Game 4. We now change the behavior of verifications. Let ${\overline{a^{1}}}^{n},{\overline{a^{2}}}^{n}$ be the sets of answers sent by $\mathcal{C}$ in line 4 of $\$_{\mathcal{Z}}$.cverify the first and second times $\mathcal{C}$ asks for verification, respectively ${ }^{18}$. Instead of performing verifications both

[^10]times, the bank now performs both verifications only on the second time: the first time $\$_{\mathcal{Z}}$.cverify is called, after line 4 the bank returns 1 and stops. The second time $\$_{\mathcal{Z}}$.cverify is called, the bank performs both verifications: i.e., on the second verification we replace everything from line 7 with:
\[

$$
\begin{array}{ll}
6: & \text { foreach } i \in[n]: \\
7: & r_{i} \leftarrow \mathcal{Z} . V\left(p_{i}, v_{i}, o_{i}, b_{i}^{1}, a_{i}^{1}\right) \\
8: & r_{i}^{\prime} \leftarrow \mathcal{Z} \cdot V\left(p_{i}, v_{i}, o_{i}, b_{i}^{2}, a_{i}^{2}\right) \\
9: & r \leftarrow r_{M A C} \cdot \prod_{i=1}^{n} r_{i} \cdot_{i}^{\prime} \\
10: & \text { return } r
\end{array}
$$
\]

Let $S_{4}$ be the event where $w>1 \wedge b^{1} \neq b^{2} \wedge i^{\prime}=j \wedge{\overline{o^{1}}}^{n}={\overline{o^{2}}}^{n}=\bar{o}^{n}$ in Game 4.

Verifying both inputs on the second request is equivalent to verifying them individually: $S_{3} \Rightarrow S_{4}$ since if both verifications pass in Game 3, then both pass in Game 4 (the first one always passes, the second one runs both verifications that passed in $S_{3}$ ), and $\neg S_{3} \Rightarrow \neg S_{4}$ since that means one of the verifications in Game 3 fail, which means the second verification in Game 4 fails. So $\operatorname{Pr}\left[S_{3}\right]=\operatorname{Pr}\left[S_{4}\right]$.

Game 5. We now change the second verification: on the $i^{\prime t h}$ pair of puzzles, if $b_{i^{\prime}}^{1} \neq b_{i^{\prime}}^{2}$ (we note that this always holds when $i^{\prime}=j$ ), we perform $V_{2}$ instead of normal verification - i.e., we replace everything from line 7 forward in $\$_{\mathcal{Z}}$.cverify in the second verification with:

$$
\begin{array}{ll}
6: & \text { foreach } i \in[n]: \\
7: & \text { if } i=i^{\prime} \wedge b_{i^{\prime}}^{1} \neq b_{i^{\prime}}^{2}: \\
8: & \hat{\text { if } b_{i}^{1}=0: \hat{a_{0}} \leftarrow a_{i}^{1}, \hat{a_{1}} \leftarrow a_{i}^{2}} \\
9: & \hat{\text { else }: \hat{a_{0}} \leftarrow a_{i}^{2}, \hat{a_{1}} \leftarrow a_{i}^{1}} \\
10: & \hat{r_{i}, r_{i}^{\prime} \leftarrow V_{2}\left(p_{i}, v_{i}, o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)} \\
11: & \text { else }: \\
12: & r_{i} \leftarrow \mathcal{Z} . V\left(p_{i}, v_{i}, o_{i}, b_{i}^{1}, a_{i}^{1}\right) \\
13: & r_{i}^{\prime} \leftarrow \mathcal{Z} . V\left(p_{i}, v_{i}, o_{i}, b_{i}^{2}, a_{i}^{2}\right) \\
14: & \text { endif } \\
15: & r \leftarrow r_{M A C} \cdot \prod_{i=1}^{n} r_{i} \cdot r_{i}^{\prime} \\
&
\end{array}
$$

16: return $r$
Let $S_{5}$ be the event where $w>1 \wedge b^{1} \neq b^{2} \wedge i^{\prime}=j \wedge{\overline{o^{1}}}^{n}={\overline{o^{2}}}^{n}=$ $\bar{o}^{n} \wedge V_{2}\left(p_{i}, v_{i}, o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)=1$ in Game 5.

In the case where $i=i^{\prime} \wedge b_{i}^{1} \neq b_{i}^{2}$, running $V_{2}\left(p_{i}, v_{i}, o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)$ is equivalent
to running $\mathcal{Z} . v$ twice, since we assign $\hat{a_{0}}$ and $\hat{a_{1}}$ respective to $b_{i}^{1}$ and $b_{i}^{2}$. So $S_{4} \Longleftrightarrow S_{5}$, meaning $\operatorname{Pr}\left[S_{4}\right]=\operatorname{Pr}\left[S_{5}\right]$.

Game 6. We now simply relax the win condition: Game 6 goes exactly the same as Game 5 , but we define the following event: let $S_{6}$ be the event where $V_{2}\left(p_{i}, v_{i}, o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)=1$ in Game 6. Since this is a relaxation of the conditions of $S_{5}$, we get $\operatorname{Pr}\left[S_{5}\right] \leq \operatorname{Pr}\left[S_{6}\right]$.

Bound on success probability. We show a reduction mapping a 2 -of- 2 mini-scheme counterfeiter to a 2 -of-2 solver (see Definition 17):

Let $\mathcal{C}$ be a QPT 2 -of- 2 mini-scheme counterfeiter. We construct a QPT 2 -of-2 solver $\mathcal{T}$ in the following manner:

Let $(p, v)$ be the output of $G\left(1^{\lambda}\right)$ at step 1 of the solving game. On step $2, \mathcal{T}$ simulates a Game 6 bank (by honestly running mints and verifications as defined in Game 5, as well as choosing $i^{\prime}$ uniformly) with two changes:

1. The $i^{\prime t h}$ puzzle is replaced with $p$.
2. In line 10 of the second verification, $\mathcal{T}$ outputs $\left(o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)$ to the puzzle giver instead of running $V_{2}$. The honest puzzle giver runs $V_{2}\left(p, v, o_{i}, \hat{a_{0}}, \hat{a_{1}}\right)$ and returns the result, which $\mathcal{T}$ uses as $r_{i}$ and $r_{i^{\prime}}$.
We can see that for any $\mathcal{C}, \operatorname{Pr}\left[S_{6}\right]$ is not affected by the above changes: in the first change we replace a random puzzle with another random puzzle, which has no affect on $\operatorname{Pr}\left[S_{6}\right]$. In the second change, the honest puzzle giver runs $V_{2}$ with exactly the same input as the bank in the original Game 6 should, and returns the result - this also does not affect $\operatorname{Pr}\left[S_{6}\right]$.
$\mathcal{T}$ runs $\mathcal{C}$ against Game $6 . S_{6}$ is exactly the win condition of the 2 -of- 2 solving game, which means $\mathcal{T}$ wins the 2 -of- 2 solving game with probability $\operatorname{Pr}\left[S_{6}\right]$. Since $\mathcal{Z}$ is a strong 1 -of- 2 puzzle, the success probability of any QPT 2-of-2 solver is negligible - meaning $\operatorname{Pr}\left[S_{6}\right]$ is negligible for any QPT counterfeiter.

For each pair of consecutive games $i$ and $i+1$, we have shown that $\operatorname{Pr}\left[S_{i}\right] \leq \operatorname{poly}(\lambda) \cdot \operatorname{Pr}\left[S_{i+1}\right]+\operatorname{negl}(\lambda)$ for some poly $(\lambda)$, negl $(\lambda)$. Finally, we have shown that $\operatorname{Pr}\left[S_{6}\right]$ is negligible in $\lambda$, so we can conclude that $\operatorname{Pr}\left[S_{0}\right]$ is negligible in $\lambda$. Since Game 0 is defined as the 2 -of- 2 mini-scheme counterfeiting game, and $S_{0}$ is defined as its win condition, no QPT 2-of-2 mini-scheme counterfeiter can win the game with more than negligible probability.

We now prove that $\$_{\mathcal{Z}}$ (Algorithm 4) is, in fact, a mini-scheme (see Definition 25). Unlike the others, this proof is not modular - not every 2-of2 mini-scheme is a mini-scheme. For example, consider a scheme wherein the bank shares with the counterfeiter a single bit of the key on each verification. This scheme could have 2 -of- 2 mini-scheme security, but obviously, it would not be secure for a counterfeiter with a verification oracle, which could easily discern the key.

Proposition 28 ( $\$_{\mathcal{Z}}$ is a mini-scheme). Assuming $\$_{\mathcal{Z}}$ is a 2-of-2 minischeme (where $\$_{\mathcal{Z}}$ is given in Algorithm 4, and a 2-of-2 mini-scheme is defined in Definition 25), $\$_{\mathcal{Z}}$ is a mini-scheme (see Definition 25).

Proof. We use an idea very similar to that used in Ref. [PYJ ${ }^{+}$12, Theorm 5] (a slightly different variation also appeared in Ref. BS16a, Appendix C]); we show that if a counterfeiter with access to a verification oracle can ask for $v$ verifications and have two of them succeed, a 2 -of- 2 counterfeiter could guess the two success indices randomly and apply the same strategy, thus breaking the security of the 2 -of- 2 mini-scheme. The following sequence of games binds the success probability of any QPT mini-scheme counterfeiter to that of a QPT 2-of-2 mini-scheme counterfeiter against $\$_{\mathcal{Z}}$ :

Game 0. Let $\mathcal{B}$ be a QPT mini-scheme counterfeiter. We assume w.l.o.g. that $\mathcal{B}$ asks for mint only once (i.e., $\ell=1$ ), and for verification $v$ times such that $v$ is polynomial in $\lambda$ - an adversary which does not comply with this assumption necessarily fails (see Definition 25. We define the first game to be COUNTERFEIT $\mathcal{B}_{\mathcal{B}, \$_{\mathcal{Z}}}^{\operatorname{mini}}(\lambda)$ (see Definition 25 ).

Let $S_{0}$ be the event where $w>1$ (see Definition 25) in Game 0 (this is the original win condition for $\mathcal{B}$ since we assume $\ell=1$ and $v$ is polynomial in $\lambda$ ).

Game 1. We now make one small change to Game 0, namely, that the game stops after $\mathcal{B}$ receives two successful verifications (i.e., the counterfeiter is not allowed to make additional verifications after receiving two successful ones. We model this by defining additional verification attempts as failures).

Let $S_{1}$ be the event where $w=2$ in Game 1 .
It is obvious why $S_{1} \Rightarrow S_{0}$. In addition, $S_{0} \Rightarrow S_{1}$, since any run of Game 0 with more than two successful verifications is equivalent to a run of Game 1 in which all verifications beyond the second successful one are ignored. So $\operatorname{Pr}\left[S_{0}\right]=\operatorname{Pr}\left[S_{1}\right]$.

Game 2. We model a run of $v$ verifications using a string $r \in\{0,1\}^{v}$, such that $r_{i}=1$ if and only if the $i^{t h}$ time $\mathcal{B}$ asked for verification was successfu ${ }^{19}$. At the beginning of Game 2, a uniform binary string $r^{\prime} \in_{R}$ $\{0,1\}^{v}$ is generated such that $\sum_{i=1}^{v} r_{i}^{\prime}=2$.

Let $S_{2}$ be the event where $w=2 \wedge r^{\prime}=r$ in Game 2.
Given $S_{1}$, we know that the string $r$ representing the verifications in Game 1, like $r^{\prime}$, also holds $\sum_{i=1}^{v} r_{i}=2$. There are $\binom{v}{2}$ such strings, so since $r^{\prime}$ was chosen uniformly and independently of $r$, there is a $\frac{1}{\binom{v}{2}}$ probability that $r^{\prime}=r$. So $\operatorname{Pr}\left[S_{2}\right]=\frac{1}{\binom{v}{2}} \cdot \operatorname{Pr}\left[S_{1}\right]$, meaning $\operatorname{Pr}\left[S_{1}\right]=\binom{v}{2} \cdot \operatorname{Pr}\left[S_{2}\right]$.

[^11]Game 3. We transform Game 2 into Game 3 by changing the following: for each $i \in[v]$, for the $i^{t h}$ time $\mathcal{B}$ runs a verification protocol with the bank, instead of receiving the actual result of the MAC and puzzle verifications $(r)$, it receives $r_{i}^{\prime}$; i.e., we change line 9 with return $r_{i}^{\prime}$.

Let $S_{3}$ be the event where $w=2 \wedge r^{\prime}=r$ in Game 3.
Given $S_{2}$, since $r^{\prime}=r$ in both Game 2 and Game 3, the fact that $\mathcal{B}$ receives $r_{i}^{\prime}$ instead of $r_{i}$ changes nothing. So $\operatorname{Pr}\left[S_{3} \mid S_{2}\right]=1$. Trivially, $\operatorname{Pr}\left[S_{3} \mid \neg S_{2}\right]=0$. So $\operatorname{Pr}\left[S_{2}\right]=\operatorname{Pr}\left[S_{3}\right]$.

Game 4. Let $k, h$ be the two indices such that $r_{k}^{\prime}=r_{h}^{\prime}=1, k \neq h$ (by construction there are exactly two such indices). In Game 4, for every verification other than the $k^{t h}$ and the $h^{t h}$, the MAC verification and puzzle verifications are not called at all - $b_{i}$ is generated and $r_{i}^{\prime}$ is returned; i.e., lines 5 to 8 are removed.

Let $S_{4}$ be the event where $w=2 \wedge r^{\prime}=r$ in Game 4.
It is easy to see that $\operatorname{Pr}\left[S_{3}\right]=\operatorname{Pr}\left[S_{4}\right]$, since for every verification but the $k^{t h}$ and the $h^{t h}$, the bank did nothing with the result of the MAC or puzzle verifications, so whether we run them at all changes nothing.

Bound on success probability. We show a reduction mapping a minischeme counterfeiter to a 2 -of-2 mini-scheme counterfeiter (see Definition 25):

Let $\mathcal{B}$ be a QPT mini-scheme counterfeiter. We construct a a QPT 2-of-2 mini-scheme counterfeiter $\mathcal{C}$ in the following manner:
$\mathcal{C}$ simulates a Game 4 bank with the following difference: when asked to run $\$_{\mathcal{Z}}$.mint, it, in turn, asks the real bank to run $\$_{\mathcal{Z}}$.mint and returns the result, and on the $k^{t h}$ an $h^{t h}$ verifications, it asks the real bank to run $\$_{\mathcal{Z}}$.cverify and returns the result. We note that for any other verification, $\mathcal{C}$ can simulate the bank since MAC and puzzle verifications are not performed; all it needs to do is choose a uniform $b$ and return $r_{i}^{\prime}$. $\mathcal{C}$ runs $\mathcal{B}$ against the simulated Game 4 bank.

So $\operatorname{Pr}\left[S_{4}\right]=\operatorname{Pr}\left[\operatorname{COUNTERFEIT} \mathcal{C}_{\mathcal{C}, \$_{\mathcal{Z}}}^{2-2}(\lambda)=1\right] \leq$ negl $(\lambda)$ for some negligible function negl $(\lambda)$. Therefore, by construction, we get that $\operatorname{Pr}\left[S_{0}\right] \leq$ $\operatorname{poly}(\lambda) \cdot \operatorname{Pr}\left[S_{4}\right]$ for some poly $(\lambda)$ and therefore is also negligible for any QPT counterfeiter. Game 0 is defined as the original mini-scheme security game, and $S_{0}$ is defined as its original win condition; therefore, $\$_{\mathcal{Z}}$ (Algorithm 4) is a mini-scheme.

### 5.3 A Mini-Scheme Implies a Full Blown Scheme

We show how a mini-scheme $\$$ can be used to construct a full blown scheme $\hat{\$}$. The construction is based on a very similar idea to those in Refs. BS16a, Appendix C] and [AC13, Section 3.3].

Here we provide an informal description of our full scheme $\hat{\$}$. The construction is defined formally in Algorithm5. Our full scheme is constructed
by minting mini-scheme banknotes, and including the key of the mini-scheme in each one. To that end, a MAC and a private-key encryption scheme are used: on minting, the bank mints a mini-scheme banknote, encrypts the mini-scheme key that was generated in the process, signs it in its encrypted form, and hands it to the user together with the mini-scheme banknote. The secure nature of the encryption scheme prevents the user from exploiting the mini-scheme key to break the mini-scheme's underlying security. On verification, the bank uses the MAC scheme to verify that the note was indeed minted by a bank, after which it decrypts the mini-scheme key to verify the mini-scheme banknote itself.

In both BS16a and AC13, the core idea of the construction is the same, with minor differences: in BS16a algorithms are used instead of interactive protocols, and AC13 is in the public setting, so a digital signature scheme is used instead of MAC, and an encryption scheme is not necessary.

We prove the security of the full-blown scheme by showing a reduction mapping a full-blown scheme counterfeiter to a mini-scheme counterfeiter, such that the mini-scheme counterfeiter generates fake bank notes for the full-blown counterfeiter.

Proposition 29 (Correctness of $\hat{\$}$ ). Assuming $\$$ is a correct mini-scheme (see Definition 22) and that both MAC and ENC have perfect completeness, \$ (Algorithm 5) is correct (see Definition 22).

Proof. From the perfect completeness of MAC (see Definition 34), we get that $\operatorname{Pr}\left[S_{m}\right]=1$, where

$$
S_{m}:=M A C \cdot \text { verify } k_{k_{m}}\left(c, M A C \cdot \operatorname{mac}_{k_{m}}(c)\right)=1 .
$$

Therefore, when the acquirer is honest, we know that he will send $t=$ MAC.sign $k_{k_{m}}(c)$ (that he received during the run of $\hat{\$}$. mint) to the bank on line 1 of $\hat{\$} . c v e r i f y$. Thus, the MAC verification on line 2 will succeed.

From the perfect completeness of ENC (see Definition 32), we get that $\operatorname{Pr}\left[S_{e}\right]=1$, where

$$
S_{e}:=E N C . \operatorname{decrypt}_{k_{e}}\left(E N C . \operatorname{encrypt}_{k_{e}}\left(k_{\S}\right)\right)=k_{\$} .
$$

Therefore, when the acquirer is honest, we know that he will send $c=$ $E N C$. encrypt $_{k_{e}}\left(k_{\$}\right)$ (that he received during the run of $\hat{\$}$.mint) to the bank on line 1 of $\hat{\$}$.cverify. Thus, the decryption in line 3 will succeed.

From the above, we conclude that for an honest acquirer both the decryption and MAC verification in $\hat{\$}$.cverify always succeeds. As such, the verification can only fail in $\$$. cverify $_{k_{8}}$. We know that the result of the decryption is $k_{\$}$ as it was generated in $\hat{\$}$.mint, and that this $k_{\$}$ was generated by running $\$$.key-gen. Thus, from the correctness of the mini-scheme $\$$ (see

| Algorithm 5 The Interactive Private Money Scheme $\hat{\$}$ |  |
| :---: | :---: |
| $\hat{\$}$. key $-\operatorname{gen}\left(1^{\lambda}\right)$ <br> 1: $k_{m} \leftarrow M A C$.key-gen <br> 2: $k_{e} \leftarrow E N C$.key-gen <br> 3: return $\left(k_{m}, k_{e}\right)$ |  |
| $\hat{\$} \cdot \operatorname{mint}_{\left(k_{m}, k_{e}\right)}$ |  |
| Acquirer | Bank |
| 1: | $k_{\$} \leftarrow$ \$.key-gen $\left(1^{\lambda}\right)$ |
| 2 : | $c \leftarrow E N C$. encrypt $_{k_{e}}\left(k_{\$}\right)$ |
| 3 : | $t \leftarrow M A C \cdot \mathrm{mac}_{k_{m}}(c)$ |
| $\|\$\rangle \leftarrow \$ . \operatorname{mint}_{k_{\Phi}}()$ |  |
| 4: $\quad c^{2} t$ |  |
| $\hat{\$ . c v e r i f y}_{\left(k_{m}, k_{e}\right)}(c, t,\|\$\rangle)$ |  |
| Acquirer | Bank |
| $1: \quad c, t$ |  |
| 2 : | $r_{m} \leftarrow M A C$.verify ${ }_{k_{m}}(c, t)$ |
| 3 : | $k_{\$} \leftarrow E N C$. decrypt $_{k_{e}}(c)$ |
| $r_{v} \leftarrow \$$. cverify $_{k_{s}}(\|\$\rangle)$ |  |
| 4: | return $r_{v}$ |
| $5:$ | return 0 |

Definition 22), we get that $\operatorname{Pr}\left[S_{\S}\right] \geq 1-\operatorname{negl}(\lambda)$ for some negligible function $\operatorname{negl}(\lambda)$, where

$$
\begin{aligned}
& S_{\$}:=k_{\$} \leftarrow \$ . \operatorname{key}-\operatorname{gen}\left(1^{\lambda}\right) ;|\$\rangle \leftarrow \$ \text { mint }_{k_{\S}}() ; \\
& \left.\quad \$ . \text { cverify }_{k_{\S}}(\$ \$\rangle\right)=1 .
\end{aligned}
$$

$\hat{\$}$. cverify passes when MAC.verify, ENC.decrypt and $\$$.cverify all pass, so for an honest acquirer:

$$
\begin{aligned}
\operatorname{Pr}\left[\left(k_{m}, k_{e}\right)\right. & \leftarrow \hat{\$} \cdot \operatorname{key}^{-g e n}\left(1^{\lambda}\right) ;(c, t,|\hat{\$}\rangle) \leftarrow \hat{\$} \cdot \operatorname{mint}_{\left(k_{m}, k_{e}\right)}() ; \\
& \left.\hat{\$} . \operatorname{cverify}_{\left(k_{m}, k_{e}\right)}(c, t,|\hat{\$}\rangle)=1\right] \\
& =1-\operatorname{Pr}\left[\neg S_{m} \cup \neg S_{e} \cup \neg S_{\$}\right] \\
& \geq 1-\operatorname{negl}(\lambda) .
\end{aligned}
$$

Theorem 30 ( $\hat{\$}$ is a secure interactive private quantum money scheme). Assuming $\$$ is an interactive private quantum money mini-scheme, MAC is a PQ-EU-CMA MAC (see Definition 35) and ENC has PQ-IND-CPA (see Definition 33), $\hat{\$}$ (Algorithm 5) is a secure interactive private quantum money scheme (see Definition 24). Moreover, if \$ is semi-quantum, $\hat{\$}$ is also semi-quantum.

Proof. The proof idea is very similar to that used in BS16a, Appendix C]: we show that the success probability of any full-scheme counterfeiter able to verify more banknotes than he received is upper-bounded by the success probability of a mini-scheme counterfeiter; a mini-scheme counterfeiter could guess which banknote the full-scheme counterfeiter will double-spend, generate fake banknotes, and with non-negligible probability double-spend the single mini-scheme banknote. The following sequence of games binds the success probability of any QPT full-scheme counterfeiter to that of a QPT mini-scheme counterfeiter:

Game 0. Let $\mathcal{A}$ be a QPT full scheme counterfeiter. We assume that the amount of mints and verifications requested by $\mathcal{A}$ is polynomial in $\lambda$ (i.e., $\ell$ and $v$ are polynomial in $\lambda$ ) since an adversary that does not comply with this assumption is not QPT. We define the first game to be the original interactive private quantum money security game, $\operatorname{COUNTERFEIT}_{\mathcal{A}, \hat{\mathbb{S}}}^{\text {full }}(\lambda)$ (see Definition 24).

Let $S_{0}$ be the event where $w>\ell$ (see Definition 24) in Game 0 (this is the original win condition of the interactive private quantum money security game, since we assume $\ell$ and $v$ are polynomial in $\lambda$ ).

Game 1. We change Game 0 slightly by adding the condition that a specific banknote is double-spent: recall that $\ell$ and $v$ are the numbers of times $\hat{\$}$.mint and $\hat{\$}$.cverify are run during Game 0 , respectively. In the start of Game 1 a uniform $i \in_{R}[\ell]$ is chosen by the bank. Let $\left(c_{j}, t_{j},\left|\$_{j}\right\rangle\right)$ be the result of the $j^{\text {th }}$ mint, and let $w_{j}$ be the amount of verifications such that $\$ . \operatorname{cverify}\left(c_{j}, t,|\$\rangle\right)=1$ for some $t,|\$\rangle$. Let $\hat{j}$ be the smallest $j$ such that $w_{j} \geq 2\left(\hat{j}=\infty\right.$ if for all $\left.j \in[\ell]: w_{j}<2\right)$.

Let $S_{1}$ be the event where $w>\ell \wedge i=\hat{j}$ in Game 1.
Assume $S_{0}$ occurred. Due to the unforgeability of MAC (see Definition 35), we know that in every successful verification, $\mathcal{A}$ presented $\left(c_{j}, t,|\$\rangle\right)$ for some $j \in[\ell], t, \mid \${ }^{20}$. Therefore, since $\mathcal{A}$ was given only $\ell$ pairs $\left(c_{j}, t_{j}\right)$ (from the $\ell$ times that $\widehat{\$}$.mint was run), and there were $w>\ell$ successful verifications from the assumption that $S_{0}$ occurred, then from the pigeonhole principle we conclude that $w_{j} \geq 2$ for some $j$, meaning $1 \leq \hat{j} \leq \ell$. Since $i \in[\ell]$ was chosen randomly and independently to $\hat{j}$, given $S_{0}$, there is a $\frac{1}{\ell}$ probability that $i=\hat{j}$; in which case $S_{1}$ occurs - therefore $\operatorname{Pr}\left[S_{1} \mid S_{0}\right]=$ $\frac{1}{\ell} \cdot \operatorname{Pr}\left[S_{0}\right]$.

Assume $S_{0}$ did not occur: then we know $w \leq \ell$, meaning $S_{1}$ also did not occur - namely, $\operatorname{Pr}\left[S_{1} \mid \neg S_{0}\right]=0$. So $\operatorname{Pr}\left[S_{1}\right]=\frac{1}{\ell} \cdot \operatorname{Pr}\left[S_{0}\right]$, meaning $\operatorname{Pr}\left[S_{0}\right]$ is $\operatorname{Pr}\left[S_{1}\right]$ times some polynomial in $\lambda$.

Game 2. We now change the above game such that now, on the $i^{\text {th }} \operatorname{mint}{ }^{21}$ instead of encrypting and signing the mini-scheme key from line 1 (the one later used in $\$$.mint), the bank encrypts and signs $0_{\$}$ (where $0_{\$}$ is a string of 0 's the length of a mini-scheme key); i.e., on the $i^{\text {th }}$ mint we replace lines 2 and 3 with:

$$
\begin{aligned}
& c \leftarrow E N C . \text { encrypt }_{k_{e}}\left(0_{\$}\right) \\
& t \leftarrow M A C \cdot \operatorname{mac}_{k_{m}}(c) .
\end{aligned}
$$

On $\hat{\$}$.cverify $(c, t,|\$\rangle)$, if $E N C$.decrypt ${ }_{k_{e}}(c)=0_{\S}$, then the bank runs $\$$.cverify with the original mini-scheme key that was used in the $i^{\text {th }}$ mint (the one originally generated in line 1 of the $i^{\text {th }}$ mint) rather than with $0_{\$}$.

Let $S_{2}$ be the event where $w>\ell \wedge i=\hat{j}$ in Game 2 .
Game 2 is different from Game 1 only in the $i^{\text {th }}$ mint, and the sole difference in the $i^{\text {th }}$ mint is that $\mathcal{A}$ receives an encrypted and signed $0_{\$}$

[^12]rather than the key that was used in $\$$.mint. Similarly, in a verification for $(c, t,|\$\rangle)$ such that $c=E N C$. encrypt $_{k_{e}}\left(0_{\$}\right)$ for some $t,|\$\rangle$, the mini-scheme bank verifies $|\$\rangle$ with the mini-scheme key that was generated in the $i^{\text {th }}$ mint, that in Game 1 is sent instead of $0_{\$}$.

That means that the only difference between Game 2 and Game 1 is in what $\mathcal{A}$ receives on the $i^{\text {th }}$ mint; on the $i^{\text {th }}$ mint, $\mathcal{A}$ receives a signed encryption of a random key rather than the key used to sign the mini-scheme banknote he received, but the same key will be used to verify it, just like on a normal verification. So, due to the indistinguishability of ENC, replacing an encryption of one message with the encryption of another message of the same length ${ }^{22}$ cannot change the behavior of $\mathcal{A}$, i.e., $\left|\operatorname{Pr}\left[S_{2}\right]-\operatorname{Pr}\left[S_{1}\right]\right| \leq$ $\operatorname{neg}(\lambda){ }^{23}$.

Bound on success probability. We show a reduction mapping a Game 2 counterfeiter to a mini-scheme counterfeiter:

Let $\mathcal{A}$ be a QPT full-scheme counterfeiter. We construct a mini-scheme counterfeiter $\mathcal{B}$ in the following manner:
$\mathcal{B}$ simulates the bank of Game 2, with one exception: on the $i^{\text {th }}$ mint, instead of generating the actual mini-scheme key and banknote, $\mathcal{B}$ asks the actual mini-scheme bank to run $\$$. mint. Similarly, when performing a verification for $0_{\$}, \mathcal{B}$ asks the actual mini-scheme bank for verification. $\mathcal{B}$ runs $\mathcal{A}$ against the altered version of Game 2. The only difference from the original Game 2 is that on the $i^{\text {th }}$ mint $\mathcal{B}$ asks the bank to generate the banknote, and when he receives $0_{\$}$ he asks the bank to verify that same note. The honest mini-scheme bank runs minting and verification on that banknote in the exact same way as the bank in Game 2 should, meaning that $\operatorname{Pr}\left[S_{2}\right]$ is unchanged for any $\mathcal{A}$ by the simulated Game 2 . In the case that $S_{2}$ occurred, $\mathcal{B}$ passed at least two verifications with $0_{\Phi}$, meaning he passed two verifications with the actual bank, while only asking mint once. So $\mathcal{B}$ has a probability of $\operatorname{Pr}\left[S_{2}\right]$ to pass win the mini-scheme counterfeiting game, and

[^13]we showed that the success probability of any QPT counterfeiter to do so is negligible - meaning $\operatorname{Pr}\left[S_{2}\right]$ must be negligible.

From construction, $\operatorname{Pr}\left[S_{0}\right]$ is also negligible. Game 0 was defined to be the original full-scheme security game, and $S_{0}$ was defined as its win condition; so $\hat{\$}$ is secure (see Definition 24 ).

## 6 LWE Implies Semi-Quantum Money

For convenience, we restate the main theorem of our private semi-quantum money result:

Theorem 2 (Private Semi-Quantum Money). Assuming that the Learning With Errors (LWE) problem with certain sets of parameters is hard for BQP, then a secure private semi-quantum money scheme exists (Definition 23).

Proof. From Theorem 43 we get that the hardness of LWE with certain parameters implies that an NTCF family exists. From Theorem 13 we get that an NTCF implies $\frac{1}{2}$-hard 1-of- 2 puzzles, and from Corollary 21 we get that weak 1-of-2 puzzles (and in particular, $\frac{1}{2}$-hard 1 -of- 2 puzzles) imply strong 1-of-2 puzzles.

From BZ13, GHS16] we get that the hardness of LWE with certain parameters ${ }^{24}$ (that are different to those used for NTCF) implies that a PQ-EU-CMA MAC and a PQ-IND-CPA encryption exist. By combining these with Propositions 26, 27, 28 and 29 and Theorem 30 (based on the constructions of Algorithm 4 and Algorithm 5), we get secure semi-quantum private money from 1-of-2 puzzles.

## 7 Semi-Quantum Money Requires Computational Assumptions

Both our public and private semi-quantum money constructions rely on computational assumptions. For our public construction we know this is necessary directly from the fact that public quantum money schemes cannot be secure against computationally unbounded adversaries AC13. However, it is known that private quantum money schemes could be informationtheoretically secure ${ }^{25}$. In this section we show that for semi-quantum money schemes, and in fact for any quantum money scheme with classical minting, computational assumptions are necessary. We do so by showing an adversary

[^14]that, after running a single minting protocol with the bank, can create 2 states which pass verification with non-negligible probability:

Theorem 31. Let $\$_{C}$ be a quantum money scheme with classical minting. Then there is a computationally unbounded adversary $\mathcal{A}$ such that:

$$
\begin{aligned}
& \operatorname{Pr}\left[k \leftarrow \$_{C} \cdot \text { key-gen }\left(1^{\lambda}\right) ;|\$\rangle \leftarrow \$_{C} \cdot \operatorname{mint}_{k}\left(1^{\lambda}\right) ;\left(\left|\$_{1}\right\rangle,\left|\$_{2}\right\rangle\right) \leftarrow \mathcal{A}\left(1^{\lambda},|\$\rangle\right) ;\right. \\
& \left.\quad \$_{C} \cdot \text { verify }\left(\left|\$_{1}\right\rangle\right)=\$_{C} \cdot \text { verify }\left(\left|\$_{2}\right\rangle\right)=1\right]=p
\end{aligned}
$$

where $p$ is non-negligible.
Proof. The adversary's strategy is simple - she simply recreates the exact run of the protocol on her own by aborting and rerunning whenever a measurement result differs from the original run with the bank, and uses the messages from the original run to substitute for the bank. Eventually she will be able to rerun the protocol with the same measurement results (this would occur only after an exponential number of tries, but this is fine since the adversary is unbounded), resulting in another quantum money state identical to that minted with the bank.

Note that Theorem 31 does not hold for a memory-dependant scheme (see Definition 6) where the database can be modified during the verification protocol, since the bank could keep track of money that has been verified. However, a privat ${ }^{[26}$ scheme where this is allowed is not interesting, since if such modifications are considered we can even construct a simple secure classical scheme (see Appendix E). Therefore, when considering memorydependant private quantum money schemes, the database can be modified only during the minting protocol.

It is known that only private quantum money schemes which are memorydependant can be information-theoretically secure - Aaronson Aar18, Theorem 8 ] showed that any semi-quantum money scheme with a fixed key (i.e., any scheme that is not memory-dependant) is vulnerable to counterfeiting by an unbounded adversary. Obviously, this also holds for schemes with classical minting that are not memory-dependant. However, Theorem 31 also holds for schemes with classical minting which are memory-dependant, thus getting a stronger result - that no quantum money scheme with classical minting can be information-theoretically secure.

## 8 Discussion

The main question that is raised in this work is the following. There are many multi-party quantum cryptographic protocols which require that both parties have quantum resources. This work elicits an important question: is

[^15]there a way (preferably, as general as possible) to convert some of these protocols to ones in which at least one of the parties does not need a quantum computer? A weaker open question can be posed from the perspective of device-independent cryptography: can at least one party use an untrusted quantum computer in unison with a trusted classical computer? We emphasize that device independent protocols (see VV19, FRV19 and references therein), such as DI quantum key distribution, DI randomness expansion ${ }^{27}$ and randomness amplification, use unconditional (information theoretic) security notions, while our protocols are only computationally secure.

Additional questions are raised when considering Remote State Preparation (RSP). RSP is a protocol that allows a classical client to create quantum states on an untrusted quantum server, with different levels of security; a protocol can promise blindness DK16 (i.e., the server learns nothing about the state) and verifiability GV19, CCKW19] (i.e., the client can verify the right states were created). This raises the following question: is there a general way to turn a classically-verifiable quantum money scheme into a semi-quantum money scheme using remote state preparation for the userside minting? We note that Pastawski et al. $\mathrm{PYJ}^{+} 12$ proved that a simple variant of Wiesner's scheme is classically verifiable. Their scheme also tolerates constant level of noise. Their construction only requires preparations of the states $|0\rangle,|1\rangle,|+\rangle$ and $|-\rangle$ which the RSP constructions above support. A recent work by Badertscher et al. $\mathrm{BCC}^{+} 20$ shows that an RSP protocol cannot be composable in the abstract cryptography framework this result is somewhat discouraging, but there might still be a way to use RSP to create semi-quantum money schemes.

Moreover, can the randomness generation protocol from $\mathrm{BCM}^{+} 18$ be amplified by our parallel repetition result? Currently the protocol has $N$ rounds; could this number be made constant using parallel repetition?

## Acknowledgments

We wish to thank Zvika Brakerski, Urmila Mahadev and Thomas Vidick for fruitful discussions. We also wish to thank an anonymous referee for insightful comments. O.S. and R.R. are supported by the Israel Science Foundation (ISF) grant No. 682/18 and 2137/19, and by the the Cyber Security Research Center at Ben-Gurion University.

## References

R. Amiri and J. M. Arrazola. Quantum money with nearly optimal error tolerance. Physical Review A, 95(6), June 2017.

[^16][Aar09] S. Aaronson. Quantum Copy-Protection and Quantum Money, In Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009, pages 229-242. IEEE Computer Society, 2009, arXiv: 1110.5353
[Aar18] S. Aaronson. Shadow Tomography of Quantum States, In I. Diakonikolas, D. Kempe, and M. Henzinger, editors, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018, pages 325-338. ACM, 2018, arXiv: 1711.01053 .
[AC13] S. Aaronson and P. Christiano. Quantum Money from Hidden Subspaces. Theory of Computing, 9:349-401, 2013, arXiv: 1203.4740 ,
[ACGH19] G. Alagic, A. M. Childs, A. B. Grilo, and S.-H. Hung. Noninteractive Classical Verification of Quantum Computation, 2019, arXiv: 1911.08101.
[AGKZ20] R. Amos, M. Georgiou, A. Kiayias, and M. Zhandry. Oneshot signatures and applications to hybrid quantum/classical authentication. In K. Makarychev, Y. Makarychev, M. Tulsiani, G. Kamath, and J. Chuzhoy, editors, Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020, pages 255268. ACM, 2020.
[AMR20] G. Alagic, C. Majenz, and A. Russell. Efficient Simulation of Random States and Random Unitaries. In A. Canteaut and Y. Ishai, editors, Advances in Cryptology - EUROCRYPT 2020 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 1014, 2020, Proceedings, Part III, volume 12107 of Lecture Notes in Computer Science, pages 759-787. Springer, 2020, arXiv: 1910.05729
$\left[\mathrm{BCC}^{+} 20\right]$ C. Badertscher, A. Cojocaru, L. Colisson, E. Kashefi, D. Leichtle, A. Mantri, and P. Wallden. Security Limitations of Classical-Client Delegated Quantum Computing, 2020, arXiv: 2007.01668.
$\left[\mathrm{BCM}^{+} 18\right]$ Z. Brakerski, P. Christiano, U. Mahadev, U. V. Vazirani, and T. Vidick. A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device. In M. Thorup, editor, 59th IEEE Annual Symposium on Foundations of

Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, pages 320-331. IEEE Computer Society, 2018, arXiv: 1804.00640
[BIN97] M. Bellare, R. Impagliazzo, and M. Naor. Does Parallel Repetition Lower the Error in Computationally Sound Protocols? In 38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997, pages 374-383. IEEE Computer Society, 1997.
[BPR12] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom Functions and Lattices. In D. Pointcheval and T. Johansson, editors, Advances in Cryptology - EUROCRYPT 2012-31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings, volume 7237 of Lecture Notes in Computer Science, pages 719-737. Springer, 2012.
[BS16a] S. Ben-David and O. Sattath. Quantum Tokens for Digital Signatures, 2016, arXiv: 1609.09047.
[BS16b] A. Broadbent and C. Schaffner. Quantum cryptography beyond quantum key distribution. Des. Codes Cryptography, 78(1):351382, 2016, arXiv: 1510.06120
[BZ13] D. Boneh and M. Zhandry. Quantum-Secure Message Authentication Codes, In T. Johansson and P. Q. Nguyen, editors, Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings, volume 7881 of Lecture Notes in Computer Science, pages 592-608. Springer, 2013.
[CCKW19] A. Cojocaru, L. Colisson, E. Kashefi, and P. Wallden. QFactory: Classically-Instructed Remote Secret Qubits Preparation. In S. D. Galbraith and S. Moriai, editors, Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I, volume 11921 of Lecture Notes in Computer Science, pages 615-645. Springer, 2019, arXiv: 1904.06303.
[CHS05] R. Canetti, S. Halevi, and M. Steiner. Hardness Amplification of Weakly Verifiable Puzzles. In J. Kilian, editor, Theory of Cryptography, Second Theory of Cryptography Conference,

TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings, volume 3378 of Lecture Notes in Computer Science, pages 17-33. Springer, 2005.
[Col19] A. Coladangelo. Smart Contracts Meet Quantum Cryptography, 2019, arXiv: 1902.05214 .
[CS20] A. Coladangelo and O. Sattath. A Quantum Money Solution to the Blockchain Scalability Problem. Quantum, 4:297, July 2020.
[Die82] D. Dieks. Communication by EPR Devices. Physics Letters A, 92(6):271-272, 1982.
[DK16] V. Dunjko and E. Kashefi. Blind quantum computing with two almost identical states, 2016, arXiv: 1604.01586.
$\left[\mathrm{FGH}^{+} 12\right]$ E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. W. Shor. Quantum Money From Knots. In S. Goldwasser, editor, Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012, pages 276-289. ACM, 2012, arXiv: 1004.5127 .
[FRV19] R. A. Friedman, R. Renner, and T. Vidick. Simple and Tight Device-Independent Security Proofs. SIAM J. Comput., 48(1):181-225, 2019, arXiv: 1607.01797.
[Gav12] D. Gavinsky. Quantum Money with Classical Verification. In Proceedings of the 27th Conference on Computational Complexity, CCC 2012, Porto, Portugal, June 26-29, 2012, pages 42-52. IEEE Computer Society, 2012, arXiv: 1109.0372.
[GHS16] T. Gagliardoni, A. Hülsing, and C. Schaffner. Semantic Security and Indistinguishability in the Quantum World. In M. Robshaw and J. Katz, editors, Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III, volume 9816 of Lecture Notes in Computer Science, pages 60-89. Springer, 2016, arXiv: 1504.05255.
[GK15] M. Georgiou and I. Kerenidis. New Constructions for Quantum Money. In 10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015, May 20-22, 2015, Brussels, Belgium, pages 92-110. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
[Gol01] O. Goldreich. The Foundations of Cryptography - Volume 1, Basic Techniques. Cambridge University Press, 2001.
[Gol04] O. Goldreich. The Foundations of Cryptography - Vol. 2, Basic Applications. Cambridge University Press, 2004.
[GV19] A. Gheorghiu and T. Vidick. Computationally-Secure and Composable Remote State Preparation. CoRR, abs/1904.06320, 2019, arXiv: 1904.06320.
[HS20] K. Horodecki and M. Stankiewicz. Semi-device-independent quantum money. New Journal of Physics, 22(2):023007, feb 2020, arXiv: 1811.10552.
[JLS18] Z. Ji, Y. Liu, and F. Song. Pseudorandom Quantum States. In H. Shacham and A. Boldyreva, editors, Advances in Cryptology - CRYPTO 2018-38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III, volume 10993 of Lecture Notes in Computer Science, pages 126-152. Springer, 2018, arXiv: 1711.00385 .
[KL14] J. Katz and Y. Lindell. Introduction to Modern Cryptography, Second Edition. CRC Press, 2014.
$\left[L^{2}{ }^{+} 10\right]$ A. Lutomirski, S. Aaronson, E. Farhi, D. Gosset, J. A. Kelner, A. Hassidim, and P. W. Shor. Breaking and Making Quantum Money: Toward a New Quantum Cryptographic Protocol. In A. C. Yao, editor, Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings, pages 20-31. Tsinghua University Press, 2010, arXiv: 0912.3825 .
[Lut10] A. Lutomirski. An Online Attack Against Wiesner's Quantum Money, 2010, arXiv: 1010.0256.
[Lut11] A. Lutomirski. Component Mixers and a Hardness Result for Counterfeiting Quantum Money, 2011, arXiv: 1107.0321.
[Mah18a] U. Mahadev. Classical Homomorphic Encryption for Quantum Circuits. In M. Thorup, editor, 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, pages 332-338. IEEE Computer Society, 2018, arXiv: 1708.02130
[Mah18b] U. Mahadev. Classical Verification of Quantum Computations. In M. Thorup, editor, 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, pages 259-267. IEEE Computer Society, 2018, arXiv: 1804.01082 .
[MS10] M. Mosca and D. Stebila. Quantum Coins, volume 523 of Contemp. Math., pages 35-47. Amer. Math. Soc., 2010, arXiv: 0911.1295.
[MVW13] A. Molina, T. Vidick, and J. Watrous. Optimal Counterfeiting Attacks and Generalizations for Wiesner's Quantum Money. In Theory of Quantum Computation, Communication, and Cryptography, pages 45-64. Springer, 2013, arXiv: 1202.4010 .
[NC11] M. A. Nielsen and I. L. Chuang. Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
[NSBU16] D. Nagaj, O. Sattath, A. Brodutch, and D. Unruh. An adaptive attack on Wiesner's quantum money, Quantum Inf. Comput., 16(11\&12):1048-1070, 2016, arXiv: 1404.1507.
[Par70] J. L. Park. The Concept of Transition in Quantum Mechanics Foundations of Physics, 1(1):23-33, Mar 1970.
$\left[\mathrm{PDF}^{+} 19\right]$ M. C. Pena, R. D. Díaz, J. Faugère, L. H. Encinas, and L. Perret. Non-quantum Cryptanalysis of the Noisy Version of AaronsonChristiano's Quantum Money Scheme. IET Information Security, 13(4):362-366, 2019.
$\left[\mathrm{PYJ}^{+} 12\right] \quad$ F. Pastawski, N. Y. Yao, L. Jiang, M. D. Lukin, and J. I. Cirac. Unforgeable Noise-Tolerant Quantum Tokens. Proceedings of the National Academy of Sciences, 109(40):16079-16082, 2012, arXiv: 1112.5456.
[Raz98] R. Raz. A Parallel Repetition Theorem. SIAM J. Comput., 27(3):763-803, 1998.
[Raz11] R. Raz. A Counterexample to Strong Parallel Repetition. SIAM J. Comput., 40(3):771-777, 2011.
[Rob19] B. Roberts. Toward Secure Quantum Money. PhD thesis, Princeton University, 2019.
[Sho04] V. Shoup. Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, 2004.
[TOI03] Y. Tokunaga, T. Okamoto, and N. Imoto. Anonymous Quantum Cash, 2003.
[VV19] U. V. Vazirani and T. Vidick. Fully Device Independent Quantum Key Distribution. Commun. ACM, 62(4):133, 2019, arXiv: 1210.1810.
[Wie83] S. Wiesner. Conjugate Coding, ACM Sigact News, 15(1):78-88, 1983.
[WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned Nature, 299(5886):802-803, October 1982.
[Zha12] M. Zhandry. How to Construct Quantum Random Functions. In 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, pages 679-687. IEEE Computer Society, 2012.
[Zha19] M. Zhandry. Quantum Lightning Never Strikes the Same State Twice. In Y. Ishai and V. Rijmen, editors, Advances in Cryptology - EUROCRYPT 2019 - Germany, May 19-23, 2019, Proceedings, Part III, volume 11478 of Lecture Notes in Computer Science, pages 408-438. Springer, 2019, arXiv: 1711.02276 .

## A Nomenclature

$\mathcal{A}$ Counterfeiter of a full blown private quantum money scheme (also NTCF adversary in Section 44, page 33

B A bank, page 13
$\mathcal{B} \quad$ Counterfeiter of a private quantum money mini-scheme, page 34
$\mathcal{C} \quad$ Counterfeiter of a private quantum money 2-of-2 mini-scheme (also quantum lightning certificate forger), page 34
$\mathcal{D} \quad$ An encryption distinguisher - i.e., an adversary whose goal is to distinguish between an actual encryption and a random string, page 58
$\$_{C} \quad$ A quantum money scheme with classical minting, page 49
$\$_{P} \quad$ Our public semi-quantum money scheme based on quantum lightning with bolt-to-certificate, page 21
$\$_{\mathcal{Z}}$ Our private semi-quantum money construction based on 1-of-2 puzzles, page 36

DS A digital signature scheme, page 19
ENC A symmetric encryption scheme, page 43
$\mathcal{F} \quad$ A MAC or digital signature forger - i.e., an adversary whose goal is to break the security of a MAC or digital signature scheme, page 59
$\in_{R} \quad$ For a finite set $S$ we denote $s \in_{R} S$ to be the process in which $s$ is sampled uniformly from $S$, page 36
$\mathcal{L} \quad$ A quantum lightning adversary, i.e., an adversary whose goal is to pass verification for two bolts with the same serial number, page 60

LWE Learning With Errors, page 7
MAC Message Authentication Code, page 58
[ $n$ ] We denote $[n] \equiv\{1, \ldots, n\}$, page 28
NTCF Noisy Trapdoor Claw-free Function, page 7
$P \quad$ A payer in a transaction, page 13
PQ-EU-CMA Post-quantum existentially unforgeable under an adaptive chosen-message attack, page 59

PQ-IND-CPA Post-quantum indistinguishable encryptions under a chosenplaintext attack, page 58
QL A quantum lightning scheme, page 19
$R \quad$ A receiver in a transaction, page 13
$\mathcal{T} \quad$ A QPT 2-of-2 solver - an adversary that attempts to solve both challenges of a 1 -of- 2 puzzle, page 25
$V_{2}$ An algorithm that verifies solutions for both challenges of a puzzle, page 25
$\mathcal{Z} \quad$ A 1-of-2 puzzle, page 25
$\hat{\mathcal{Z}} \quad$ A weakly verifiable puzzle, page 29

## B Preliminaries

This appendix contains mainly the standard definitions of private-key encryption and message authentication codes (MAC), and can be safely skipped by readers already familiar with these notions.

We use the standard definitions for negligible, non-negligible and noticeable functions - see, e.g., Gol01.
Definition 32 (Private-key encryption system, KL14, Definition 3.7]). $A$ private-key encryption scheme consists of three PPT algorithms key-gen, encrypt and decrypt such that:

1. The randomized key-generation algorithm key-gen takes as input $1^{\lambda}$ and outputs a key $k \leftarrow$ key-gen $\left(1^{\lambda}\right)$.
2. The (possibly randomized) encryption algorithm encrypt takes as input $a$ key $k$ and a plaintext message $m \in\{0,1\}^{*}$, and outputs a ciphertext $c \leftarrow \operatorname{encrypt}_{k}(m)$.
3. The deterministic decryption algorithm decrypt takes as input a key $k$ and a ciphertext $c$, and outputs a message $m:=\operatorname{Dec}_{k}(c)$.

A private-key encryption system is required to have perfect completeness, meaning that for every $\lambda$, every $k$ output by key-gen $\left(1^{\lambda}\right)$, and every $m \in$ $\{0,1\}^{*}$, it holds that decrypt $\left(\operatorname{encrypt}_{k}(m)\right)=m$.
Definition 33 (PQ-IND-CPA, adapted from [KL14, Definition 3.22]). A private-key encryption scheme $\Pi$ has post-quantum indistinguishable encryptions under a chosen-plaintext attack ( $P Q-I N D-C P A$ ) if for every $Q P T$ distinguisher $\mathcal{D}$ there is a negligible function $\operatorname{negl}(\lambda)$ such that, for all $\lambda$ :

$$
\operatorname{Pr}\left[\operatorname{IND}-\operatorname{CPA}_{\mathcal{D}, \Pi}(\lambda)=1\right] \leq \frac{1}{2}+\operatorname{neg}(\lambda)
$$

The indistinguishability game $\operatorname{IND}-\mathrm{CPA}_{\mathcal{D}, \Pi}(\lambda)$ :

1. A key $k$ is generated by running key-gen $\left(1^{\lambda}\right)$.
2. The distinguisher $\mathcal{D}$ is given input $1^{\lambda}$ and classical oracle access to encrypt $_{k}(\cdot)$, and outputs a pair of messages $m_{0}, m_{1}$ of the same length.
3. A uniform bit $b \in_{R}\{0,1\}$ is chosen, and then a ciphertext $c \leftarrow$ encrypt $_{k}\left(m_{b}\right)$ is computed and given to $\mathcal{D}$.
4. $\mathcal{D}$ continues to have oracle access to $\operatorname{encrypt}_{k}(\cdot)$ and outputs a bit $b^{\prime}$.
5. The output of the game is defined to be 1 if $b^{\prime}=b$, and 0 otherwise. In the former case, we say that $\mathcal{D}$ succeeds.

Definition 34 (Message authentication code [KL14, Definition 4.1]). A message authentication code (MAC) consists of 3 PPT algorithms key-gen, mac and verify satisfying:

1. key-gen takes as input the security parameter $1^{\lambda}$ and outputs a key $k$
2. mac takes as input a key $k$ and a message $m \in\{0,1\}^{*}$ and outputs a $\operatorname{tag} t \leftarrow \operatorname{mac}_{k}(m)$.
3. verify takes as input a key $k$, a message $m$, and a tag $t$. It outputs a bit $b:=\operatorname{verify}_{k}(m, t)$, with $b=1$ meaning valid and $b=0$ meaning invalid.

A MAC is required to have perfect completeness, i.e., for every $\lambda$, every key $k \leftarrow \operatorname{key}-\operatorname{gen}\left(1^{\lambda}\right)$ and every $m \in\{0,1\}^{*}$, it holds that verify $\left(m, \operatorname{mac}_{k}(m)\right)=$ 1.

Definition 35 (PQ-EU-CMA MAC, adapted from [KL14, Definition 4.2]). A message authentication code $\Pi$ is Post-Quantum Existentially Unforgeable under an adaptive Chosen-Message Attack ( $P Q-E U-C M A$ ) if for every QPT forger $\mathcal{F}$, there exists a negligible function negl $(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{MAC}-\operatorname{FORGE} \mathcal{F}_{\mathcal{F}, \Pi}(\lambda)=1\right] \leq \operatorname{neg}(\lambda) .
$$

The CMA message authentication game $\mathrm{MAC}^{-\operatorname{FORGE}_{\mathcal{F}, \Pi}}(\lambda)$ :

1. A key $k$ is generated by running key-gen $\left(1^{\lambda}\right)$.
2. The forger $\mathcal{F}$ is given input $1^{\lambda}$, classical oracle access to $\operatorname{mac}_{k}(\cdot)$ and classical oracle access to verify $k(\cdot)$ (note that the forger cannot query the oracles in superposition). The forger eventually outputs ( $m, t$ ). Let $\mathcal{Q}$ denote the set of all queries that $\mathcal{F}$ asked its signing oracle.
3. $\mathcal{F}$ succeeds if and only if (1) verify ${ }_{k}(m, t)=1$ and (2) $m \notin \mathcal{Q}$. In that case the output of the game is defined to be 1 .

Definition 36 (Digital signature scheme KL14, Definition 12.1]). A digital signature scheme consists of three PPT algorithms key-gen, sign and verify such that:

1. The key-generation algorithm key-gen takes as input a security parameter $1^{\lambda}$ and outputs a pair of keys $(p k, s k)$. These are called the public key and the private key, respectively. We assume that pk and sk each has length of at least $\lambda$, and that $\lambda$ can be determined from either.
2. The signing algorithm sign takes as input a private key sk and a message $m$. It outputs a signature $\sigma \leftarrow \operatorname{sign}_{s k}(m)$.
3. The deterministic verification algorithm verify takes as input a public key $p k$, a message $m$ and a signature $\sigma$. It outputs a bit $b \leftarrow$ verify $_{s k}(m, \sigma)$, with $b=1$ meaning valid and $b=0$ meaning invalid.

A digital signature scheme is required to have perfect completeness, meaning that except with negligible probability over ( $p k, s k$ ) output by key-gen $\left(1^{\lambda}\right)$, it holds that verify ${ }_{p k}\left(m, \operatorname{sign}_{s k}(m)\right)=1$ for every legal message $m$.

Definition 37 (PQ-EU-CMA digital signature scheme, adapted from KL14 Definition 12.2]). A digital signature scheme $\Pi$ is Post-Quantum Existentially Unforgeable under an adaptive Chosen Message Attack ( $P Q-E U-C M A$ ) if for every QPT forger $\mathcal{F}$, there exists a negligible function negl $(\lambda)$ such that:

$$
\operatorname{Pr}\left[\operatorname{SIG}-\operatorname{FORGE}_{\mathcal{F}, \Pi}(\lambda)=1\right] \leq \operatorname{negl}(\lambda) .
$$

The signature experiment $\operatorname{SIG-FORGE} \mathcal{F}_{\mathcal{F}, \Pi}(\lambda)$ :

1. key-gen is run to generate to obtain keys $(p k, s k)$.
2. Forger $\mathcal{F}$ is given $p k$ and access to a signing oracle $\operatorname{sign}_{s k}(\cdot)$. The forger than outputs $(m, \sigma)$. Let $Q$ denote the set of all queries that $\mathcal{F}$ asked its oracle.
3. $\mathcal{F}$ succeeds iff verify ${ }_{p k}(m, \sigma)=1$ and $m \notin Q$. In this case the output of the experiment is defined to be 1 (and otherwise 0).

Lemma 38 (Difference Lemma [Sho04, Lemma 1]). Let $A, B, F$ be events defined in some probability distribution, and suppose that $A \wedge \neg F \Longleftrightarrow$ $B \wedge \neg F$. Then $|\operatorname{Pr}[A]-\operatorname{Pr}[B]| \leq \operatorname{Pr}[F]$.

## C Quantum Lightning with Bolt-to-Certificate

The following definitions are taken almost verbatim from [CS20. The definitions originate in Zha19] and Col19], but in this work we use the notations of the superseding work CS20].

Definition 39 (Quantum Lightning [Zha19]). A quantum lightning scheme consists of a PPT algorithm QL.setup $\left(1^{\lambda}\right)$ (where $\lambda$ is a security parameter) which samples a pair of QPT algorithms (gen-bolt, verify-bolt). gen-bolt outputs a pair of the form $|\psi\rangle \in \mathcal{H}_{\$}, s \in\{0,1\}^{\lambda}$. We refer to $|\psi\rangle$ as a "bolt" and to $s$ as a "serial number". verify-bolt takes as input a pair of the same form, and outputs either "accept" (1) or "reject" (0). They satisfy the following:

$$
\begin{aligned}
\operatorname{Pr}[(\text { gen-bolt, verify-bolt }) & \leftarrow \mathrm{QL} \cdot \operatorname{setup}\left(1^{\lambda}\right) ;(|\psi\rangle, s) \leftarrow \text { gen-bolt }(): \\
& \text { verify-bolt }(|\psi\rangle, s)=1] \\
& =1-\operatorname{negl}(\lambda) .
\end{aligned}
$$

- For all $s^{\prime} \in\{0,1\}^{\lambda}$ :

$$
\begin{aligned}
\operatorname{Pr}[(\text { gen-bolt, verify-bolt }) & \leftarrow \mathrm{QL} . \operatorname{setup}\left(1^{\lambda}\right) ;(|\psi\rangle, s) \leftarrow \text { gen-bolt }(): \\
& \left.s \neq s^{\prime} \wedge \text { verify-bolt }\left(|\psi\rangle, s^{\prime}\right)=1\right] \\
& =\operatorname{negl}(\lambda)
\end{aligned}
$$

Definition 40 (Security [Zha19]). A quantum lightning scheme QL is secure if, for all QPT bolt forgers $\mathcal{L}$ :

$$
\operatorname{Pr}\left[\operatorname{FORGE}-\operatorname{BOLT}_{\mathcal{L}, \mathrm{QL}}(\lambda)=1\right]=\operatorname{negl}(\lambda)
$$

The bolt forging game $\mathrm{FORGE}^{-\mathrm{BOLT}_{\mathcal{L}, \mathrm{QL}}}(\lambda)$ :

1. The challenger runs (gen-bolt, verify-bolt) $\leftarrow \mathrm{QL} . \operatorname{setup}\left(1^{\lambda}\right)$ and sends (gen-bolt, verify-bolt) to $\mathcal{L}$.
2. $\mathcal{L}$ produces a pair $\left|\Psi_{12}\right\rangle \in \mathcal{H}_{\$}^{\otimes 2}, s \in\{0,1\}^{\lambda}$.
3. The challenger runs verify-bolt(•,s) on each half of $\left|\Psi_{12}\right\rangle$. The output of the game is 1 if both outcomes are "accept" (and otherwise 0).

Definition 41 (Bolt-to-certificate). For a quantum lightning scheme QL to have bolt-to-certificate capability, we change the procedure QL.setup $\left(1^{\lambda}\right)$ slightly, so that it outputs a quadruple

> (gen-bolt, verify-bolt, gen-certificate, verify-certificate),
where gen-certificate is a QPT algorithm that takes as input a quantum money state and a serial number and outputs a classical string of some fixed length $l(\lambda)$ for some polynomially bounded function $l$, to which we refer as a certificate, and verify-certificate is a PPT algorithm that takes as input a serial number and a certificate, and outputs "accept" (1) or "reject" (0).

Let $\lambda \in \mathbb{N}$. We say that a quantum lightning scheme QL has bolt-tocertificate capability if:
$\operatorname{Pr}\left[(\right.$ gen-bolt, verify-bolt, gen-certificate, verify-certificate $) \leftarrow$ QL.setup $\left(1^{\lambda}\right)$;

$$
\begin{aligned}
& (|\psi\rangle, s) \leftarrow \operatorname{gen}-\text { bolt }() ; c \leftarrow \text { gen-certificate }(|\psi\rangle, s): \\
& \text { verify-certificate }(s, c)=1] \\
& =1-\operatorname{neg} \mid(\lambda) .
\end{aligned}
$$

- For all QPT algorithms $\mathcal{C}$ :

$$
\operatorname{Pr}\left[\operatorname{FORGE}-\operatorname{CERTIFICATE} \mathcal{E}_{\mathcal{C}, \mathrm{QL}}(\lambda)=1\right]=\operatorname{negl}(\lambda) .
$$

The certificate forging game $\operatorname{FORGE}^{-C E R T I F I C A T E} \mathrm{E}_{\mathcal{C}, \mathrm{QL}}(\lambda)$ :

1. The challenger runs
(gen-bolt, verify-bolt, gen-certificate, verify-certificate) $\leftarrow Q L . \operatorname{setup}\left(1^{\lambda}\right)$,
and sends the quadruple to $\mathcal{C}$.
2. $\mathcal{C}$ returns $c \in\{0,1\}^{l(\lambda)}$ and $(|\psi\rangle, s)$.
3. The challenger runs verify-certificate $(s, c)$ and verify-bolt $(|\psi\rangle, s)$, and outputs 1 if they both accept (otherwise outputs 0).

## D Trapdoor Claw-Free Families

Most of this section is taken verbatim from Brakerski et al. $\mathrm{BCM}^{+} 18$. Let $\lambda$ be a security parameter, and let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets (depending on $\lambda$ ). For our purposes, an ideal family of functions $\mathcal{F}$ would have the following properties. For each public key $k$, there are two functions $\left\{f_{k, b}: \mathcal{X} \rightarrow\right.$ $\mathcal{Y}\}_{b \in\{0,1\}}$ that are both injective, that have the same range (equivalently, $(b, x) \mapsto f_{k, b}(x)$ is 2-to-1), and that are invertible given a suitable trapdoor $t_{k}$ (i.e., $t_{k}$ can be used to compute $x$ given $b$ and $\left.y=f_{k, b}(x)\right)$. Furthermore, the pair of functions should be claw-free: it must be hard for an attacker to find two pre-images $x_{0}, x_{1} \in \mathcal{X}$ such that $f_{k, 0}\left(x_{0}\right)=f_{k, 1}\left(x_{1}\right)$. Finally, the functions should satisfy an adaptive hardcore bit property, which is a stronger form of the claw-free property: assuming for convenience that $\mathcal{X}=\{0,1\}^{w}$, we want it to be computationally infeasible to simultaneously generate $\left(b, x_{b}\right) \in\{0,1\} \times \mathcal{X}$ and a non-zero string $d \in\{0,1\}^{w}$ such that with a non-negligible advantage over $\frac{1}{2}$ the equation $d \cdot\left(x_{0} \oplus x_{1}\right)=0$ holds, where $x_{1-b}$ is defined as the unique element such that $f_{k, 1-b}\left(x_{1-b}\right)=f_{k, b}\left(x_{b}\right)$.

Unfortunately, we (as well as Brakerski et al.) do not know how to construct a function family that exactly satisfies all these requirements under standard cryptographic assumptions. Instead, Brakerski et al. construct a family that satisfies slightly relaxed requirements based on the hardness of the learning with errors (LWE) problem, and we will show that these are still adequate for our purposes. The requirements are relaxed as follows. First, the range of the functions is no longer a set $\mathcal{Y}$; instead, it is $\mathcal{D}_{\mathcal{Y}}$, the set of probability densities over $\mathcal{Y}$. That is, each function returns a density, rather than a point. The trapdoor injective pair property is then described in terms of the support of the output densities: these supports should either be identical for a colliding pair or be disjoint in all other cases.

The consideration of functions that return densities elicits an additional requirement of efficiency: there should exist a quantum polynomial-time procedure that efficiently prepares a superposition over the range of the function, i.e., for any key $k$ and $b \in\{0,1\}$, the procedure can prepare a state that is close (up to a negligible trace distance) to the state

$$
\frac{1}{\sqrt{\mathcal{X}}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{f_{k, b}(x)(y)}|x\rangle|y\rangle
$$

We modify the adaptive hardcore bit requirement slightly. Since the set $\mathcal{X}$ may not be a subset of binary strings, we first assume the existence of an injective, efficiently invertible map $J: \mathcal{X} \rightarrow\{0,1\}^{w}$. Next, we only require the adaptive hardcore bit property to hold for a subset of all nonzero strings rather than for the set $\{0,1\}^{w} \backslash\left\{0^{w}\right\}$. Finally, membership in the appropriate set should be efficiently checkable, given access to the trapdoor.

Definition 42 (NTCF family). Let $\lambda$ be a security parameter. Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets. Let $\mathcal{K}_{\mathcal{F}}$ be a finite set of keys. A family of functions

$$
\mathcal{F}=\left\{f_{k, b}: \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}\right\}_{k \in \mathcal{K}_{\mathcal{F}}, b \in\{0,1\}}
$$

is called a noisy trapdoor claw free (NTCF) family if the following conditions hold:

1. Efficient Function Generation. There exists an efficient probabilistic algorithm key-gen $\mathcal{F}^{\text {which generates a key } k \in \mathcal{K}_{\mathcal{F}} \text { together }}$ with a trapdoor $t_{k}$ :

$$
\left(k, t_{k}\right) \leftarrow \operatorname{key}-\operatorname{gen}_{\mathcal{F}}\left(1^{\lambda}\right)
$$

2. Trapdoor Injective Pair. For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold.
(a) Trapdoor: For all $b \in\{0,1\}$ and $x \neq x^{\prime} \in \mathcal{X}, \operatorname{Supp}\left(f_{k, b}(x)\right) \cap$ $\operatorname{SUPP}\left(f_{k, b}\left(x^{\prime}\right)\right)=\emptyset$. Moreover, there exists an efficient deterministic algorithm $I N V_{\mathcal{F}}$ such that for all $b \in\{0,1\}, x \in \mathcal{X}$ and $y \in \operatorname{SUPP}\left(f_{k, b}(x)\right), I N V_{\mathcal{F}}\left(t_{k}, b, y\right)=x$.
(b) Injective pair: There exists a perfect matching $\mathcal{R}_{k} \subseteq \mathcal{X} \times \mathcal{X}$ such that $f_{k, 0}\left(x_{0}\right)=f_{k, 1}\left(x_{1}\right)$ if and only if $\left(x_{0}, x_{1}\right) \in \mathcal{R}_{k}$.
3. Efficient Range Superposition 28 There exists an efficient procedure $S A M P_{\mathcal{F}}$ that on input $k$ and $b \in\{0,1\}$ prepares a state $\left|\psi^{\prime}\right\rangle$ which has a negligible trace distance to the state

$$
|\psi\rangle=\frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{\left(f_{k, b}(x)\right)(y)}|x\rangle|y\rangle
$$

4. Adaptive Hardcore Bit. For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold, for some integer $w$ that is a polynomially bounded function of $\lambda$.
(a) For all $b \in\{0,1\}$ and $x \in \mathcal{X}$, there exists a set $G_{k, b, x} \subseteq\{0,1\}^{w}$ such that $\operatorname{Pr}_{d \leftarrow_{U}\{0,1\}^{w}}\left[d \notin G_{k, b, x}\right]$ is negligible, and moreover there exists an efficient algorithm that checks for membership in $G_{k, b, x}$ given $k, b, x$ and the trapdoor $t_{k}$.

[^17](b) There is an efficiently computable injection $J: \mathcal{X} \rightarrow\{0,1\}^{w}$, such that $J$ can be inverted efficiently on its range, and such that the following holds. If
\[

$$
\begin{array}{cr}
H_{k}= & \left\{\left(b, x_{b}, d, d \cdot\left(J\left(x_{0}\right) \oplus J\left(x_{1}\right)\right)\right) \mid\right. \\
\bar{H}_{k}= & b \in\{0,1\},\left(x_{0}, x_{1}\right) \in \mathcal{R}_{k}, d \in G_{\left.k, 0, x_{0} \cap G_{k, 1, x_{1}}\right\}, 29}^{29}  \tag{8}\\
\left\{\left(b, x_{b}, d, c\right) \mid(b, x, d, c \oplus 1) \in H_{k}\right\},
\end{array}
$$
\]

then for any quantum polynomial-time procedure $\mathcal{A}$ there exists a negligible function $\mu(\cdot)$ such that

$$
\begin{align*}
&\left|\operatorname{Pr}_{\left(k, t_{k}\right) \leftarrow \text { key }-\operatorname{gen}_{\mathcal{F}}\left(1^{\lambda}\right)}\left[\mathcal{A}(k) \in H_{k}\right]-\underset{\left(k, t_{k}\right) \leftarrow \text { key }-\operatorname{gen}_{\mathcal{F}}\left(1^{\lambda}\right)}{ }\left[\mathcal{A}(k) \in \bar{H}_{k}\right]\right| \\
& \leq \mu(\lambda) . \tag{9}
\end{align*}
$$

Theorem 43 (Informal). Under the assumption that the Learning With Errors (LWE) problem with certain parameters is hard for BQP, an NTCF family exists.

The hardness definition of LWE and the exact parameters required for the theorem above are given in $\left[\mathrm{BCM}^{+} 18\right.$, Theorem 26].

## E The Advantage of Memoryless Money

When discussing any form of quantum money, we must consider the motivation, i.e., the benefits over classical constructions - for example, we could construct a rudimentary private classical money scheme in the following way: upon minting, the bank would produce a random serial number significantly long for some security parameter $\lambda$ and sign it using a MAC. The bank would maintain a database of all banknotes that have already been spent, and upon verification, after verifying the MAC tag of the banknote, the bank would search for its serial number within the database if it is not there, the verification succeeds and the serial number is added to the database, and if it is there the bank would know the money was already spent and thus verification will fail (of course, the bank would have to mint a new banknote for the user after a successful verification). Gavinsky [Gav12] discusses a similar notion.

This scheme is counterfeit-resistant according to our security definitions. However, it is memory-dependent (also known as state-based); i.e., the bank has to maintain a database to represent an ongoing state, remembering the banknotes that were spent. On its own, a memory-dependent protocol is

[^18]not a terrible problem; many services maintain a database. This, however, becomes a liability when considering multiple branches of the same bank: a central database must maintain the shared state and synchronize the access to it (otherwise information would have to propagate between the branches, causing potential security breaches during the propagation time); this has a toll in terms of response time and communication.

Constructing a memory-dependent (classical) private money scheme is trivial - the scheme above is an extremely simple example - so such a construction is not particularly interesting. The case is different, however, in the public setting; constructing even a memory-dependent quantum money scheme that is publicly secure is challenging (and impossible to achieve classically), and thus such a construction is an interesting result.

## F Parallel Repetition of Weakly Verifiable Puzzles

As seen in Section 4.3, our main tool for proving parallel repetition for 1-of2 puzzles is the notion of weakly verifiable puzzles introduced by CHS05. This section provides a brief informal overview of their parallel repetition proof for weakly verifiable puzzles (Theorem 20).

The goal is to show a reduction from an algorithm $A$ (which in our case may be quantum) that solves $n$ weakly verifiable puzzles in parallel with probability at least $h^{n}+\operatorname{negl}(\lambda)$ for some non-negligible $h$ to an algorithm $A^{\prime}$ solving a single puzzle with probability at least $h+\operatorname{negl}(\lambda)$. This shows that the best strategy for solving $n$ puzzles has the same probability (up to a negligible difference) of solving each puzzle separately.

Denote the event where $A$ solves the puzzle $p_{i}$ correctly by $S_{i}$, and denote the event where $A$ solves puzzles $p_{k}, \ldots, p_{n}$ correctly by $R_{k}$. We call a puzzle coordinate $i \in[n]$ good if it holds that $\underset{p_{1}, \ldots, p_{n}}{\operatorname{Pr}}\left[S_{i} \mid R_{i+1}\right] \geq h+\operatorname{negl}(\lambda)$, i.e., if the probability that $A$ solved $p_{i}$ correctly conditioned on the probability that $A$ solved $p_{i+1}, \ldots, p_{n}$ correctly is at least $h$, where the probability is taken over the randomness of $A$ as well as over the random choices of puzzles $p_{1}, \ldots, p_{n}$. We make the following statistical observation:

$$
\operatorname{Pr}\left[R_{1}\right]=\operatorname{Pr}\left[S_{1} \wedge R_{2}\right]=\operatorname{Pr}\left[S_{1} \mid R_{2}\right] \cdot \operatorname{Pr}\left[R_{2}\right] \geq h^{n}+\operatorname{negl}(\lambda)
$$

The last inequality holds because we know the probability of $A$ to solve all $n$ puzzles is at least $h^{n}$. Therefore, we conclude that at least one of the following must hold: either $\operatorname{Pr}_{p_{1}, \ldots, p_{n}}\left[S_{1} \mid R_{2}\right] \geq h$, or $\underset{p_{1}, \ldots, p_{n}}{\operatorname{Pr}}\left[R_{2}\right] \geq h^{n-1}$. Meaning either 1 is a good coordinate, or the probability that $A$ solves the last $n-1$ puzzles of its input is at least $h^{n-1}$.

If 1 is a good coordinate, then $A^{\prime}$ solves its original puzzle in the following way: $A^{\prime}$ runs $A$ with $\left(p, p_{2}, \ldots, p_{n}\right)$ - where $p$ is the original challenge puzzle and $p_{2}, \ldots, p_{n}$ are randomly generated by $A^{\prime}$ - until puzzles $p_{2}, \ldots, p_{n}$
were solved correctly ${ }^{30}$ Since 1 is a good coordinate, we know that if $A$ solved puzzles $p_{2}, \ldots p_{n}$ correctly, the first puzzle is also solved correctly with probability at least $h$, as needed.

If 1 is not a good coordinate, than we know the probability of $A$ to solve $n-1$ puzzles is at least $h^{n-1}$. This provides us with a solver for $n-1$ puzzles with probability at least $h^{n-1}+\operatorname{negl}(\lambda)$, and the reduction continues by induction. Eventually, $A^{\prime}$ will either reach a good coordinate $i$ and solve it with probability at least $h$ as described above, or arrive at a solver that solves a single puzzle - the last puzzle - with probability at least $h$.

To find out whether $i$ is a good coordinate, $A$ estimates $\underset{p_{1}, \ldots, p_{n}}{\operatorname{Pr}}\left[S_{i} \mid R_{i+1}\right]$. This is done by generating $n$ puzzles, and running $A$ multiple times, all the while checking whether puzzles $p_{i+1}, \ldots, p_{n}$ were solved correctly, and if so checking whether puzzle $p_{i}$ was also solved correctly. Since $\operatorname{Pr}\left[R_{i+1}\right] \geq$ $\operatorname{Pr}\left[R_{1}\right] \geq h^{n}+\operatorname{negl}(\lambda)$, this estimation can be done efficiently with high probability.

The proof itself in [CHS05] is more complex than the outline given here. For example, we cannot know for sure if a coordinate is a good coordinate, since we can only estimate the probabilities within some polynomial bounds - the original proof deals with these issues.

Note that this parallel repetition is "perfect" - i.e., it shows that the best strategy for solving multiple puzzles has the same probability as solving them separately (up to a negligible difference). This would not be the case if we had more than 2 rounds - for 3 rounds the repetition would not be perfect, and for 4 rounds the soundness error does not decrease exponentially at all BIN97. We can see that the proof above indeed would not work for more than 2 rounds: in our reduction, when we find a good coordinate $i$, we run the multiple-puzzle adversary multiple times until the puzzles $p_{i+1}, \ldots, p_{n}$ are solved correctly, and then our original puzzle is solved correctly with probability $h$. But in the case of multiple rounds, the single-puzzle adversary is required to provide the multiple-puzzle adversary with messages from the challenger for all the puzzles, including the original one. The answers of the multiple-puzzles adversary for each puzzle can therefore depend on messages for all the puzzles, but we can run the protocol for the original puzzle with the challenger only once - i.e., we cannot provide the challenger's answers for the original puzzles more than once. Therefore, we cannot condition on the adversary solving $p_{i+1}, \ldots, p_{n}$ correctly, meaning that if we can run the adversary only once our success probability will not be $h$ as needed.

[^19]
## G Transaction Figures

This section contains figures showing transactions can be made with each flavor of quantum money and the kind of communication they require (quantum or classical). Fat arrows $\Rightarrow$ indicate quantum communication and thin arrows $\rightarrow$ indicate classical communication. A two-sided arrow indicates two-sided communication.


Figure 2：Regular（non－quantum）money direct transaction．The communi－ cation required is one－way，though usually two－way communication is used to send confirmations for actions．


Figure 3：Regular（non－quantum）money bank transaction．Though it is possible for communication between payer and bank to be one－way，usually the authentication process is an interactive protocol，and the payer receives confirmation from the bank，making communication two－way．


Figure 4：Private standard direct transaction．


Figure 5: Standard bank transaction.


Figure 6: Private classically verifiable direct transaction. In step (b), receiver acts as a relay between payer and bank and thus can be sure the banknote was valid.


Figure 7：Classically verifiable bank transaction．


Figure 8：Private classically verifiable direct transaction for a scheme which allows multiple verifications for the same banknote．Note that after a finite number of verifications the banknote is destroyed and communication with the bank is required，like in Fig．6．


Figure 9: Private classically verifiable bank transaction for a scheme which allows multiple verifications for the same banknote. For some such schemes it could be possible to have a classical verification that ensures the banknote was destroyed, in which case a transaction through the bank could be executed like in Fig. 11


Figure 10：Private classical minting direct transaction．


Figure 11: Classical minting bank transaction.


Figure 12: Private semi-quantum direct transaction. In step (b), receiver acts as a relay between payer and bank and thus can be sure the banknote was valid.


Figure 13: Semi-quantum bank transaction.


Figure 14：Public quantum money direct transaction．


Figure 15：One－shot signature direct transaction．


[^0]:    ${ }^{1}$ This is not problematic; the bank would simply mint a new banknote and send it, through a quantum channel, to the receiver. The concept of non-reusable money is not a new one; in fact, an otherwise secure quantum money scheme could sometimes be broken if banknotes are reused (see NSBU16, Lut10, Aar09).

[^1]:    ${ }^{2}$ We refer to such a scheme as "memory-dependent", and explore its consequences in Appendix E

[^2]:    ${ }^{3}$ In our public semi-quantum money construction, the classical verification can only be done with the bank; users still require quantum communication to transfer banknotes without the bank. We leave it as an open question whether public semi-quantum money with entirely classical communication can be made. A recent work of Amos et al. AGKZ20] exhibits completely classical communication between users, but in the oracle model.

[^3]:    ${ }^{4}$ We deviate here slightly from the formal definitions; Since the bank does not have a "database", verification should only use the key. This is handled by using a message authentication code (MAC) and by returning to the user a tag for these values, and then verifying that tag during the verification. For the sake of clarity, we omit this part in the discussion - refer to Algorithm 4 to see how we work around this issue.

[^4]:    ${ }^{5}$ Some classically verifiable schemes Gav12 allow to verify the same banknote multiple times before it is destroyed. Schemes with this property behave a little differently.
    ${ }^{6}$ The authors are not aware of an instantiation of a private quantum money scheme with classical minting that is not semi-quantum, but it is a valid setting to consider.
    ${ }^{7}$ The authors are not aware of any instantiation of a classically-verifiable public quantum money scheme (that is not also semi-quantum), but it is a valid setting to consider.
    ${ }^{8}$ Quantum lightning Zha19 can be seen as quantum money with classical minting which is not necessarily classically verifiable - in Section 3 we construct public semiquantum quantum money from quantum lightning with bolt-to-certificate, but we can see that classical minting is straightforward even without this quality.
    ${ }^{9}$ This is a construction from AGKZ20 for public semi-quantum money which allows

[^5]:    to transfer banknotes between users using classical communication. However, the only construction available requires an oracle.

[^6]:    ${ }^{10}$ If there is non-negligible probability that a generated bolt would have the same serial number as one already in $D$, assuming $D$ contains an amount of serial numbers polynomial in $\lambda$, we could construct a bolt forger $\mathcal{L}$ that would generate $|D|+1$ bolts and with non-negligible probability end up with two bolts with the same serial number that pass verify-bolt, winning the bolt forging security game (Definition 40 ).
    ${ }^{11}$ We arbitrarily number the runs of $\$_{P}$.mint according to the order they were initiated.

[^7]:    ${ }^{12}$ Suppose $\mathcal{A}$ passes with non-negligible probability a verification of some $(|\psi\rangle, s, \sigma)$ on querify or a verification of $s, \sigma, c$ on cverify such that $s \neq s_{j} \forall j \in[\ell]$. In that case DS.verify $p_{p k_{\sigma}}(s, \sigma)=1$ with non-negligible probability (either on line 2 of querify or on line 4 of ceverify). We could use $\mathcal{A}$ to construct a digital signature forger $\mathcal{F}$ with non-negligible success probability: $\mathcal{F}$ simulates a bank, but instead of signing with $s k_{\sigma}$ and verifying with $p k_{\sigma}$ that were generated in key-gen he uses his signing oracle and the real $p k$. He then runs $\mathcal{A}$ against the simulated bank, and will be able to present $(s, \sigma)$ which pass DS verification with non-negligible probability, while $s \notin Q$ since $s \neq s_{j} \forall j \in[\ell]$.

[^8]:    ${ }^{13}$ In Ref. CHS05] this notion is defined for any PPT algorithm.

[^9]:    ${ }^{14}$ The weakly verifiable puzzle $\hat{Z}$ constructed from a 1 -of- 2 puzzle is classical — each attempt to solve the puzzle creates a new quantum state, so even when running a solver for a puzzle multiple times there is indeed no rewinding.
    ${ }^{15}$ Note that $n(\lambda)$ is indeed polynomial in $\lambda$ - since a weak 1 -of- 2 puzzle holds that $1-h$ is noticeable (see Definition 12, by using the inequality $\ln (1-\varepsilon) \leq-\varepsilon$ we get that $\log (1 / h)$ is noticeable.
    ${ }^{16}$ Recall that a weak 1-of-2 puzzle has completeness $\eta=1$ (see Definition 12 ).

[^10]:    ${ }^{17}$ Otherwise, we could construct a MAC forger $\mathcal{F}$ with non-negligible success probability. Assume towards a contradiction that with non-negligible probability, $\mathcal{C}$ passes verification by sending in line $4{\overline{o^{\prime}}}^{n}, t_{o}^{\prime}$ such that ${\overline{o^{\prime}}}^{n} \neq \bar{o}^{n}$. That means that the MAC verification in line 5 passed. So $\mathcal{F}$ could simulate a bank, but instead of signing and verifying with $k$ generated in $\$_{\mathcal{Z}}$.key-gen, $\mathcal{F}$ uses the signing and verification oracles. $\mathcal{F}$ runs $\mathcal{C}$ against the simulated bank, and present ${\overline{o^{\prime}}}^{n},{\tilde{o^{\prime}}}^{n}$. With non-negligible probability, MAC verification passes, and since ${\overline{o^{\prime}}}^{n} \neq \bar{o}^{n}$, and no other signings are requested (mint was run only once), $\mathcal{F}$ wins MAC-FORGE $\mathcal{F}, M A C$ ( $\lambda$.
    ${ }^{18} \mathcal{C}$ can, of course, run both verification protocols simultaneously. We number the verifications according to the one that got to line 4 of the protocol first.

[^11]:    ${ }^{19} \mathcal{B}$ can, of course, run several verification protocols simultaneously. We number the verifications according to the order in which they were initiated.

[^12]:    ${ }^{20}$ Suppose $\mathcal{A}$ passes with non-negligible probability a verification of $(c, t,|\$\rangle)$ such that $c \neq c_{j} \forall j \in[\ell]$. In that case $M A C$.verify $(c, t)=1$ with non-negligible probability. We could use $\mathcal{A}$ to construct a forger $\mathcal{F}$ with non-negligible success probability: $\mathcal{F}$ simulates a bank, but instead of signing and verifying with the MAC key generated in $\hat{\$}$.key-gen, he uses the signing and verification oracles. He then runs $\mathcal{A}$ against the simulated bank, and will be able to present $c, t$ which pass MAC verification with non-negligible probability, while he did not ask for a tag of $c$ before since $c \neq c_{j} \forall j \in[\ell]$.
    ${ }^{21} \mathcal{A}$ can, of course, run several mint protocols simultaneously. We number them according to the order they were initiated.

[^13]:    ${ }^{22}$ Indistinguishability works for messages of the same length. Here we assume, without loss of generality, that key-gen $\left(1^{\lambda}\right)$ always outputs keys of the same length.
    ${ }^{23}$ Assume $\left|\operatorname{Pr}\left[S_{2}\right]-\operatorname{Pr}\left[S_{1}\right]\right|$ is non-negligible. Assume without loss of generality that $\operatorname{Pr}\left[S_{2}\right] \leq \operatorname{Pr}\left[S_{1}\right]$. In that case, we could construct a distinguisher $\mathcal{D}$ with non-negligible success probability: $\mathcal{D}$ will simulate a bank, but instead of encrypting with the ENC key generated in $\hat{\$}$. key-gen, he will use the encryption oracle, and instead of decrypting he will "remember" each encryption he made and thus could match each encryption to the relevant key (any unrecognized encryption would not have passed the real bank verification either, because the encryptions are MAC signed). On the $i^{t h}$ mint, he will present a random key and the actual mini-scheme key used in that mint as $m_{0}$ and $m_{1}$ (the chosen messages whose encryptions he needs to recognize in the CPA game) respectively, and proceed with the encryption he received to finish the game. $\mathcal{D}$ returns $b^{\prime}=1$ (guessing the encryption he received was of the real key) if and only if he wins the counterfeiting game (since he wins with higher probability when he receives encryption of the real key). $\mathcal{D}$ has a $\frac{1}{2}+\frac{\left|\operatorname{Pr}\left[S_{2}\right]-\operatorname{Pr}\left[S_{1}\right]\right|}{2}$ probability to win, which is non-negligible, in contradiction to the security of ENC.

[^14]:    ${ }^{24}$ Both BZ13] and GHS16] rely on Quantum Pseudorandom Functions (QPRF). From Banerjee et al. BPR12 and Zhandry Zha12 we get that QPRFs can be constructed from LWE with certain parameters.
    ${ }^{25}$ Wiesner's original quantum money scheme Wie83 was proven to be informationtheoretically secure MVW13, and Pastawski et al. $\mathrm{PYJ}^{+} 12$ even showed a slightly modified version of the scheme that is also noise-tolerant.

[^15]:    ${ }^{26}$ This discussion is not relevant to the public setting.

[^16]:    ${ }^{27}$ Also known as certified randomness.

[^17]:    ${ }^{28}$ Here we use a slightly weaker (and simpler) definition compared to Brakerski et al. Our definition follows from theirs by using Lemma 2 in $\mathrm{BCM}^{+} 18$, which relates the Hellinger distance to the trace distance.

[^18]:    ${ }^{29}$ Note that although both $x_{0}$ and $x_{1}$ are referred to to define the set $H_{k}$, only one of them, $x_{b}$, is explicitly specified in any 4-tuple that lies in $H_{k}$.

[^19]:    ${ }^{30}$ This can be verified since $A^{\prime}$ generated those puzzles along with their verification information and therefore can verify them efficiently. Moreover, $A$ solves $n-1$ puzzles with non-negligible probability, meaning $A^{\prime}$ can get correct solutions in polynomial time with high probability.

