Quantum money with nearly optimal error tolerance

Ryan Amiri
SUPA, Institute of Photonics and Quantum Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom

Juan Miguel Arrazola
Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543
(Received 28 March 2017; published 26 June 2017)

We present a family of quantum money schemes with classical verification which display a number of benefits over previous proposals. Our schemes are based on hidden matching quantum retrieval games and they tolerate noise up to 23%, which we conjecture reaches 25% asymptotically as the dimension of the underlying hidden matching states is increased. Furthermore, we prove that 25% is the maximum tolerable noise for a wide class of quantum money schemes with classical verification, meaning our schemes are almost optimally noise tolerant. We use methods in semidefinite programming to prove security in a substantially different manner to previous proposals, leading to two main advantages: first, coin verification involves only a constant number of states (with respect to coin size), thereby allowing for smaller coins; second, the reusability of coins within our scheme grows linearly with the size of the coin, which is known to be optimal. Last, we suggest methods by which the coins in our protocol could be implemented using weak coherent states and verified using existing experimental techniques, even in the presence of detector inefficiencies.

DOI: 10.1103/PhysRevA.95.062334

I. INTRODUCTION

Quantum cryptography has traditionally been associated exclusively with quantum key distribution [1], but it encompasses a much larger class of tasks and protocols [2]. Notable examples are quantum signature schemes [3–5], two-party quantum cryptography [6–8], delegated quantum computation [9,10], covert quantum communication and steganography [11–14], quantum random number generation [15–17], quantum fingerprinting [18–21], and quantum money [22–24]. Historically, many of these protocols have been extremely challenging to implement with available technologies, but we are currently approaching a point where both theoretical and experimental developments have made it possible for the first experimental demonstrations to emerge. We are thus entering an exciting stage where practical quantum cryptography has begun to expand rapidly beyond the realms of quantum key distribution.

Quantum money, which was suggested by Weisner in 1970 [22] as a means to create money that is physically impossible to counterfeit, is one of the first examples of quantum cryptography. The basic aim of any quantum money scheme is to enable a trusted authority, the bank, to provide untrusted users with finitely reusable, verifiable coins that cannot be forged. Verifiability ensures that honest users can prove the money they hold is genuine, while unforgeability restricts the ability of an adversary to dishonestly fabricate additional coins. Potential drawbacks of Weisner's original scheme were that verification required quantum communication between the holder and the bank, and moreover security of the scheme had not been proved rigorously. Indeed, it was shown in Refs. [25–27] that many variants of the scheme were vulnerable to so-called "adaptive attacks"—attacks in which the adversary is allowed a number of auxiliary interactions with the bank before trying to forge a coin.

In 2012, Gavinsky [23] addressed both issues and presented a fully secure quantum money scheme in which coins are verified using three rounds of *classical* communication between the holder of the coin and the bank. The scheme was based on hidden matching quantum retrieval games (QRGs), introduced in Ref. [28]. Nevertheless, the scheme could not be considered *practical*, as the security analysis did not include the effects of noise. This issue was addressed by Pastawski *et al.* [29], in which a noise tolerant quantum money scheme with classical verification was proposed that remains secure as long as the overall transmission fidelity is greater than $\frac{1}{2} + \frac{1}{\sqrt{8}} \approx 85.4\%$. The scheme requires only two rounds of communication for verification and is secure even against adaptive attacks. Following this, Ref. [24] presented a simpler protocol, again based on hidden matching QRGs, in which the verification procedure contained only a single round of communication and displayed an increased noise tolerance of up to 12.5%, where noise is defined as the probability of a single honest verifier measurement returning an incorrect outcome.

Beyond the secret-key quantum money schemes discussed above, there has also been significant interest in public-key quantum money schemes, proposed in [25], offering computational security against quantum adversaries. Since then, Farhi *et al.* [30] introduced the concepts of quantum state restoration and single-copy tomography to further rule out a large class of seemingly promising schemes. Following this result, Farhi *et al.* [31] suggested a scheme based on knot theory and conjectured that it is secure against computationally bounded adversaries. However, whether a secure public-key quantum money scheme exists without the use of oracles is an open question and, so far, the majority of schemes that were proposed have subsequently been broken [32].

In this work, we focus on secret-key quantum money schemes with classical verification and propose a scheme based on hidden matching QRGs. Utilising semidefinite programming, we provide a full security proof of our scheme, and show that by increasing the dimension of the underlying

states, we can increase the error tolerance to as much as 23.03% for states of dimension $n = 14$, while also proving that the maximum noise tolerance in that case is 23.3%. Thus, the error tolerance of our protocols is nearly optimal. We conjecture that for large dimension, the error tolerance of our protocols approaches 25% asymptotically, and we further prove that 25% is the maximum possible error tolerance for a wide range of quantum money protocols, including all those based on hidden matching QRGs. Increasing the error tolerance has a twofold benefit: as well as allowing the protocol to be performed in regions of higher noise than was previously possible, it also increases protocol efficiency since we show that security relies on the size of the gap between the expected error rate and the maximum tolerable error rate of the scheme, thereby allowing smaller coins.

We use methods in semidefinite programming to prove security in a substantially simpler manner compared to previous proposals such as Refs. [23,24]. Besides an increase in noise tolerance, our scheme also has two additional advantages compared to previous proposals: coin verification involves only a constant number of states with respect to coin size, thereby allowing for smaller coins, and the reusability of coins within our scheme grows linearly with the size of the coin, which is known to be optimal. Finally, we discuss how our schemes can be implemented in practice using a coherent state encoding, while also showing that they remain secure even in the presence of limited detection efficiency.

Definitions and previous results

In this section we state various definitions that are needed to introduce our quantum money schemes. We consider the case of quantum money "minischemes" in which the bank creates only a single quantum coin and the adversary attempts to use this coin to forge another copy. It has been shown in Ref. [33] that by adding a classical serial number to each coin, a secure full quantum money scheme can be created directly from the secure minischeme, and so the two are essentially equivalent.

Definition 1. A quantum money minischeme with classical verification consists of an algorithm, *Bank*, which creates a quantum coin \$ and a verification protocol Ver, which is a classical protocol run between a holder H of \$ and the bank B, designed to verify the authenticity of the coin. The final output of this protocol is a bit $b \in \{0,1\}$ sent by the bank, which corresponds to whether the coin is valid or not. Denote by $\mathrm{Ver}_H^B(\$)$ this final bit. The scheme must satisfy two properties to be secure:

(i) Correctness: The scheme is ϵ-correct if for every honest holder, we have

$$\Pr\left[\mathrm{Ver}_H^B(\$) = 1\right] \geqslant 1 - \epsilon.$$

(ii) Unforgeability: Coins in the scheme are ϵ-unforgeable if for any quantum adversary who has interacted a finite and bounded number of times with the bank and holds a valid coin \$, the probability that she can produce two coins $\$_1$ and $\$_2$ that are verified by an honest user satisfies

$$\Pr\left[\mathrm{Ver}_H^B(\$_1) = 1 \wedge \mathrm{Ver}_H^B(\$_2) = 1\right] \leqslant \epsilon,$$

where H is any honest holder.

The first property guarantees that all honest participants can prove the coins they own are valid, while the second property guarantees that a dishonest adversary cannot forge the coins. The definition covers adaptive attacks by allowing the adversary to interact with the bank (via the verification procedure) a finite number of times before attempting to forge the coin.

The schemes presented in this paper are based on quantum retrieval games (QRGs), which we have mentioned but not formally introduced. A QRG is a protocol performed between two parties, Alice and Bob, and can be seen as a generalization of state discrimination. Alice holds an n-bit string x, selected at random according to a probability distribution $p(x)$, which she encodes into a quantum state ρ_x. She sends the state to Bob, whose goal is to provide a correct answer to a given question about x. Mathematically, a question is modelled as a relation: if X is the set of possible values x can take, and if A is the set of possible answers, the relation σ is a subset of $X \times A$. If $(x,a) \in \sigma$, this means that, given x, the answer a is a correct answer to the "question" σ. Formally, a quantum retrieval game is defined as follows.

Definition 2. Let X and A be the sets of inputs and answers respectively. Let $\sigma \subset X \times A$ be a relation and $\{p(x),\rho_x\}$ an ensemble of states and their a priori probabilities. Then the tuple $G = (X,A,\{p(x),\rho_x\},\sigma)$ is called a quantum retrieval game. If Bob may choose to find an answer to one of a finite number of distinct relations $\sigma_1,\ldots,\sigma_k$, then we write the game as $G = (X,A,\{p(x),\rho_x\},\sigma_1,\ldots,\sigma_k)$.

A particularly useful class of QRGs is the *hidden matching* QRGs [23,24,34], in which the relations are defined by matchings. A matching M on the set $[n] := \{1,2,\ldots,n\}$, where n is an even number, is a partitioning of the set into $n/2$ disjoint pairs of numbers [35]. A matching can be visualized as a graph with n nodes, where edges define the elements in the matching, as illustrated in Fig. 1. In general, there are $1\times3\times\ldots\times(n-1) = (n-1)!!$ distinct matchings of any set containing n elements. For our purposes, we focus on sets of matchings where no two matchings in the set contain a common element. We call such sets *pairwise disjoint*. The maximum number of pairwise disjoint matchings is $n-1$, since if we consider the element $1 \in [n]$, it must be paired in each matching with a distinct integer less than or equal to n.

Definition 3. A maximal pairwise disjoint set of matchings $\mathcal{R}$ is a set of pairwise disjoint matchings on $[n]$ such that $|\mathcal{R}| := n-1$.

A matching on the set $[n]$ can be equivalently represented as a graph with n nodes, with each element (i,j) of the matching identified with an edge in the graph. Maximal pairwise disjoint sets of matchings for $n = 4$, 6, and 8 are illustrated in Fig. 1.

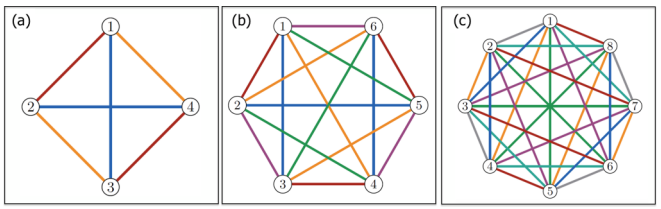

FIG. 1. Maximal pairwise disjoint set of matchings for (a) $n = 4$, (b) $n = 6$, and (c) $n = 8$. Color is used to represent each matching within the maximal pairwise disjoint set.

In hidden matching QRGs the set of possible inputs is the set of all n-bit strings, each chosen with equal probability, where n is an even number. Alice encodes her input into the n-dimensional pure state

$$|\phi_x\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} (-1)^{x_i} |i\rangle, \tag{1}$$

where x_i is the ith bit of the string x. Note that this state corresponds to a $O(\log_2 n)$ qubit state, so that the number of qubits needed in the scheme scales favorably with n.

The relations in this game are defined by the matchings: given a matching, the correct answers are the ones which correctly identify the parity of the bits connected by an edge in the matching. For example, if (1,2) is an element of the matching, the measurement should output $x_1 \oplus x_2$. Formally, given a perfect matching M_1, the set of answers is given by

$$A = \{(i,j,b) : i,j \in \{1,\ldots,n\}, b \in \{0,1\}\}$$

and the corresponding relation is

$$\sigma_1 = \{(x,i,j,b) : x_i \oplus x_j = b \text{ and } (i,j) \in M_1\}.$$

Bob is able to find a correct answer to any matching of his choice with certainty simply by measuring in the basis

$$\mathcal{B} = \left\{ \frac{1}{\sqrt{2}} (|i\rangle \pm |j\rangle) \right\}, \quad \text{with} \quad (i,j) \in M. \tag{2}$$

This is because the outcome $\frac{1}{\sqrt{2}}(|i\rangle + |j\rangle)$ can only occur if $x_i \oplus x_j = 0$, and similarly $\frac{1}{\sqrt{2}}(|i\rangle - |j\rangle)$ can only occur if $x_i \oplus x_j = 1$.

Previous quantum money schemes based on hidden matching QRGs have used only two matchings for verification. In the following section, we generalize these schemes to the case of an arbitrary number of matchings and show that this allows us to significantly increase the noise tolerance of the resulting schemes.

II. QUANTUM MONEY SCHEME

Here we present a quantum money scheme which is secure even in the presence of up to 23% noise. As in Ref. [24], the verification protocol requires only one round of classical communication.

In this scheme, the bank randomly chooses a number of n-bit classical strings and encodes each of them into the hidden matching states, given by Eq. (1). Essentially, the coin is a collection of these independent quantum states, and each of the quantum states can be thought of as an instance of a QRG. We assume that there is a maximal pairwise disjoint set of matchings on $[n]$, known to all participants, which we call $\mathcal{R}$. This set specifies the $n-1$ possible relations defined within each QRG, and each state in the coin represents a QRG. To verify a coin, the holder will pick a small selection of the states from the coin and randomly choose a relation for each. The holder will perform the appropriate measurement [defined by Eq. (2)] to get an answer for each QRG under each chosen relation. The holder then sends these answers to the bank which returns whether or not more than a specified fraction of the answers are correct. If they are, the coin is accepted as valid; otherwise, it is rejected. The scheme is formally defined below and illustrated in Figs. 2 and 3.

Bank algorithm

(1) The bank independently and randomly chooses q n-bit strings which we will call $x^1,\ldots,x^q$.

(2) For $i \in [q]$, the bank creates $\phi_{x^i} := |\phi_{x^i}\rangle\langle\phi_{x^i}|$, where

$$|\phi_{x^i}\rangle := \frac{1}{\sqrt{n}} \sum_{j=1}^{n} (-1)^{x^i_j} |j\rangle.$$

For each i we define the QRG $G_i = (S_i, A_i, \{\phi_{x^i}\}_{x^i}, \sigma_1,\ldots,\sigma_{n-1})$, where $\mathcal{R} = \{\sigma_1,\ldots,\sigma_{n-1}\}$ is a maximal pairwise disjoint set of matchings known to all participants in the scheme.

(3) The bank creates the classical binary register r and initializes it to 0^q.

(4) The bank creates the counter variable s and initializes it to 0.

(5) The pair $(\$,r) = (\bigotimes_{i=1}^{q} \phi_{x^i}, r)$ is the coin for the minischeme. The bank keeps the counter s in order to keep track of the number of verification attempts.

Ver algorithm

(1) The holder of the coin randomly chooses a subset of indices, $L \subset [q]$ such that $r_i = 0$ for each $i \in L$. The indices $i \in L$ specify the selection of games G_i which will be used as tests in the verification procedure. For each $i \in L$, the holder sets the corresponding bit of r to be 1 so that this game cannot be used in future verifications.

(2) For each $i \in L$, the holder picks a relation σ'_i at random from $\mathcal{R}$ and applies the appropriate measurement to obtain outcome d_i.

(3) The holder sends all triplets (i,σ'_i,d_i) to the bank.

(4) The bank checks that $s < T$, where T is the predefined maximum number of allowed verifications for the coin. If $s = T$, the bank declares the coin as invalid.

(5) For each i, the bank checks whether the answer is correct by comparing (i,σ'_i,d_i) to the secret x^i values. The bank accepts the coin as valid if and only if more than $l(c-\delta)$ of the answers are correct, where c is a correctness parameter of the protocol, $l = |L|$, and δ is a small positive constant.

(6) The bank updates s to $s+1$.

We say that an instance of the verification algorithm has been passed or failed if the final output by the bank is "valid"

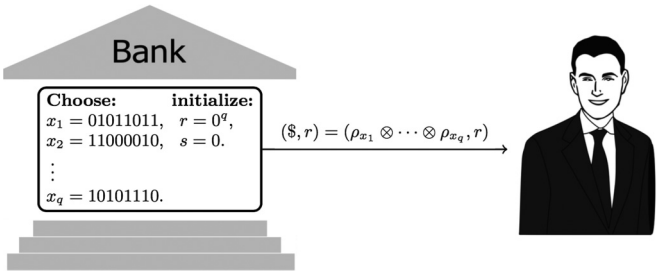

FIG. 2. Schematic illustration of the Bank algorithm for $n = 8$. The bank selects q eight-bit strings and initializes the q-bit register r to the zero string. The bank creates the corresponding hidden matching states and sends these, together with r, to the holder of the coin.

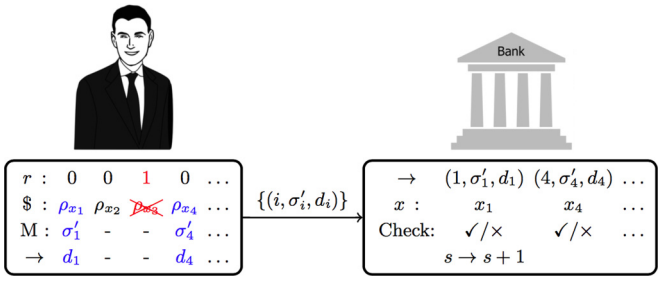

FIG. 3. Schematic showing the verification algorithm. The verifier selects a sample $\{\rho_{x_1}, \rho_{x_4}, \ldots\}$ of the states contained within the coin which have an r value of 0. He randomly chooses matching measurements and applies them to get classical measurement outcomes which he sends to the bank, together with the index of the state and the matching chosen. The bank checks these against its secret strings, as well as checking $s < T$. Finally, the bank declares an output based on the number of incorrect outcomes.

or "invalid" respectively. Coins can be verified at most T times until the Hamming weight of r is greater than Tl, at which point the coin is returned to the bank to be refreshed. We choose T to be small but linear in q. Any such choice would be acceptable but, for the sake of definiteness, in what follows we set $T := q/(1000l)$. We note that having T scale linearly with q is optimal for any quantum money scheme [23] and that this is an improvement over previous protocols (for example those in Refs. [23,24]).

The noise of the protocol is defined as the probability that an honest verifier obtains an incorrect outcome when making the honest measurement on a single QRG state (i.e., in step 2 of the verification procedure). In the ideal setting we can set $c = 1$, since an honest participant in possession of a correct state will always get a correct answer to a relation. Of course, in the practice system imperfections inevitably lead to errors so that even when all participants are honest, it is not certain that the holder's measurement will return a correct answer. Thus, in the presence of errors, we must have $c < 1$, and the smallest value of c for which we can retain security determines the noise tolerance of the protocol.

We note that this scheme requires the bank to maintain a small classical database to record the number of times the verification protocol has been run—i.e., the bank's database is "nonstatic," and must be updated after each run of verification. Although this requirement demands more from the bank than completely static database models, we believe the requirement is both minimal and realistic, and allows significant simplifications to the security analysis.

Nevertheless, in some cases it may be desirable for the bank to have a completely static database—for example in applications in which the bank consists of many small decentralized branches. In such a scenario, attacks targeting multiple branch locations may be able to compromise security by gaining additional verification attempts. To provide safeguards against these types of attack, our scheme could be modified in two different ways.

The simplest method would be to assume that all bank branches have access to a single common database, thereby preventing verifiers from performing too many verification attempts on a single coin. Alternatively, we could add an additional round of classical communication to the verification protocol, similarly to Ref. [23], in which the bank selects the states to be used in the verification protocol. The effect would be to transform our scheme into one which uses a fully static database, but still retains the same level of noise tolerance. Security of this modified scheme can be proved by directly applying the arguments in Ref. [23] to show that the additional verification attempts do not (significantly) help the adversary [36].

Security

In this section we prove that the scheme defined above is secure according to Definition 1.

1. *Correctness*

Correctness of the scheme follows simply from the Hoeffding bound [37]. In the honest case, if the holder of a coin has probability c of getting a correct answer for each of the l QRGs selected in the verification protocol, then his probability of getting fewer than $(c - \delta)l$ correct answers overall is bounded by

$$\mathbb{P}(\text{honest fail}) \leqslant e^{-2l\delta^2}. \tag{3}$$

Based on the security analysis in the following section, we choose δ to be half of the gap between the error rate an honest participant expects and the minimum error rate the adversary can achieve. I.e., we set $\delta := (e_{\min} - \beta)/2$, where $e_{\min}$ is the minimum error rate achievable by the adversary [derived below in Eq. (27)], and $\beta := 1 - c$ is the error rate expected in an honest run of the protocol.

2. *Unforgeability*

We assume the adversary is in possession of a valid coin and first address a simple forging strategy available to the adversary based on manipulating the r register attached to the coin. The adversary is allowed to set at most $q/1000$ of the r register entries to 1. She creates $(\$_1, r_1)$ and $(\$_2, r_2)$ to send to the two honest verifiers, Ver_1 and Ver_2 respectively. If she sets $r_1(i) = 1$ and $r_2(i) = 0$, she can be certain that Ver_1 will not select the ith state to test, and so can forward the perfect state to Ver_2. In this way, $q/1000$ of the states in the coins sent to each verifier will be perfect, and will not cause errors. The remaining positions must have r register values of 0 for both verifiers. Similarly, the adversary is able to use the auxiliary verification attempts to her advantage. We make a worst-case assumption and assume that the adversary gets full knowledge of every state used in an auxiliary verification attempt. Since there are at most T attempts allowed, each of which involve l states, the adversary knows the identity of at most $q/1000$ of the states. Since the states are prepared independently, this knowledge does not provide any information on the remaining states.

The combined effect of the above two strategies is that the adversary is able to exactly replicate $q/500$ of the states in the coin, as shown in Fig. 4. To prove coins are unforgeable, we consider the remaining $997q/1000$ states for which the r register is zero for both verifiers, and for which the adversary has no auxiliary information. In reference to Fig. 4, we refer to these states as the white states, and start by considering a single such state, $\phi_{x^i} := |\phi_{x^i}\rangle\langle\phi_{x^i}|$, contained in the coin. For

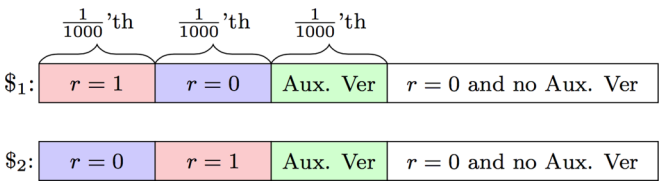

FIG. 4. Representation of the states within the quantum coins sent to the verifiers. The first block on the far left represents all states for which the adversary set $r = 1$ for Ver_1, and $r = 0$ for Ver_2. The adversary knows that Ver_1 cannot select these states for testing, and so is able to forward on the perfect states to Ver_2. The second block of states represents the same, but with the roles of the verifiers reversed. The "Aux. Ver" states in the diagram are the ones that we assume are known to the adversary via auxiliary verifications. The remaining states in white are the ones we consider below—those states for which the r register is zero for both verifiers, and which have not been used in auxiliary verifications.

simplicity, we drop the superscript on the n-bit strings x^i in all that follows.

The idea behind the proof is to relate the probability that the forger can use a single white state to create two states that pass the verification test of the two honest verifiers, to the average fidelity of these two states with the original state $|\phi_x\rangle$. The maximization of this average fidelity corresponds to the optimal attack, which can be cast as a semidefinite program. By focusing on the dual program, we can upper bound the value of the semidefinite program and therefore bound the forging probability of the adversary. Last, we show that coherent attacks on multiple states cannot help the adversary to forge.

Since the adversary has a valid coin, she holds the unknown state

$$|\phi_x\rangle = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}(-1)^{x_i}|i\rangle. \tag{4}$$

From this state, the adversary wishes to create two states, η_x and τ_x, which, when measured by the honest verifiers, will give the correct answer to a randomly chosen relation in $\mathcal{R}$. At this stage we ignore any auxiliary verification attempts available to her. Consider the normalized state sent to Ver_1,

$$\eta_x = \sum_{i,j=1}^{n} a_{ij}|i\rangle\langle j|. \tag{5}$$

Suppose the verifier chooses to measure using the matching $M_\alpha = \{(i_1, j_1), \ldots, (i_{n/2}, j_{n/2})\}$, where $\alpha \in \{1, 2, \ldots, n-1\}$. To find a correct answer to the relation σ_α defined by this matching, an honest verifier will apply the measurement with projectors in the set $\{|+_{i_k j_k}\rangle\langle+_{i_k j_k}|, |-_{i_k j_k}\rangle\langle-_{i_k j_k}| : k = 1, \ldots, n/2\}$, where $|\pm_{i_k j_k}\rangle := \frac{1}{\sqrt{2}}(|i_k\rangle \pm |j_k\rangle)$. An incorrect result is obtained whenever the verifier finds an incorrect value for $x_{i_k} \oplus x_{j_k}$, which happens whenever the measurement outcome is one of the form

$$\frac{1}{\sqrt{2}}[|i\rangle - (-1)^{x_i \oplus x_j}|j\rangle]. \tag{6}$$

This happens with probability

$$p^{\alpha,x}_{\mathrm{Ver}_1} = \frac{1}{2}\left(1 - \sum_{k=1}^{n/2}(-1)^{x_{i_k}\oplus x_{j_k}} a_{i_k j_k} + (-1)^{x_{i_k}\oplus x_{j_k}} a_{j_k i_k}\right). \tag{7}$$

Thus, the probability of an incorrect answer to σ_α is given by a subset of the off-diagonal elements of the density matrix η_x. The off-diagonal elements occurring are exactly those with indices paired by the matching M_α. Since the set of relations form a maximal pairwise disjoint set, the off-diagonal matrix elements appearing in the error probability for different relations will all be distinct. Therefore, averaging over all possible relations that could be chosen by the verifier allows us to significantly simplify the adversary's error probability, which becomes

$$\begin{aligned} p^{x}_{\mathrm{Ver}_1} &= \frac{1}{n-1}\sum_{\alpha=1}^{n-1} p^{\alpha,x}_{\mathrm{Ver}_1} \\ &= \frac{1}{2(n-1)}\left(n - \sum_{i,j=1}^{n}(-1)^{x_i \oplus x_j} a_{ij}\right) \\ &= \frac{n}{2(n-1)}(1 - F_x), \end{aligned} \tag{8}$$

where we have defined

$$F_x := \langle\phi_x|\eta_x|\phi_x\rangle = \frac{1}{n}\sum_{i,j}(-1)^{x_i\oplus x_j} a_{ij}. \tag{9}$$

Since the adversary does not know the secret string x, rather than holding the state in Eq. (4), she instead holds a mixture over the possible x values. We define $F := \frac{1}{2^n}\sum_x F_x$ and take an average over x values to get

$$\begin{aligned} p_{\mathrm{Ver}_1} &= \frac{1}{2^n}\sum_x p^{x}_{\mathrm{Ver}_1} \\ &= \frac{1}{2^n}\sum_x \frac{n}{2(n-1)}(1 - F_x) \\ &= \frac{n}{2(n-1)}(1 - F). \end{aligned} \tag{10}$$

Essentially then, to successfully forge a coin, the adversary is trying to create two states, η_x and τ_x, which both have a high fidelity with the original state $|\phi_x\rangle$. Let us define $G_x = \langle\phi_x|\tau_x|\phi_x\rangle$, and $G := \frac{1}{2^n}\sum_x G_x$. For the purpose of forging, the adversary needs *both* Ver_1 and Ver_2 to accept the coin she sends, which requires her to make both error probabilities as small as possible. From the above result, we can relate this to maximizing the average fidelity of the states η_x and τ_x with the original state. This problem can be cast as a semidefinite program as follows.

Let $\Psi : L(\mathcal{X}) \to L(\mathcal{Y} \otimes \mathcal{Z})$ be a physical channel taking states in Hilbert space $\mathcal{X}$ to states in the Hilbert space $\mathcal{Y} \otimes \mathcal{Z}$, where both $\mathcal{Y}$ and $\mathcal{Z}$ are isomorphic to $\mathcal{X}$. We want to find the channel that maximizes

$$\overline{F} = \frac{1}{2^n}\sum_{x=1}^{2^n}\frac{\langle\phi_x|\eta_x|\phi_x\rangle + \langle\phi_x|\tau_x|\phi_x\rangle}{2}, \tag{11}$$

where $\eta_x = \mathrm{Tr}_{\mathcal{Z}}[\Psi(|\phi_x\rangle\langle\phi_x|)]$ and $\tau_x = \mathrm{Tr}_{\mathcal{Y}}[\Psi(|\phi_x\rangle\langle\phi_x|)]$. In other words, η_x is the reduced state of the channel output

representing the state held by Ver_1, and τ_x is the reduced state of the channel output representing the state held by Ver_2. This maximization is subject to Ψ being a completely positive trace preserving linear map. To express this maximization in the standard form of a semidefinite program, we express the channel as an operator using the Choi representation. We fix the preferred basis to be $\{|i\rangle\}_{i=1,\ldots,n}$, the basis used to define the hidden matching states in the ensemble. Given this choice, the Choi operator corresponding to the channel Ψ is an operator $J(\Psi)$ in $L(\mathcal{X}\otimes\mathcal{Y}\otimes\mathcal{Z})$, given by

$$J(\Psi)=\sum_{i,j=1}^{n}|i\rangle\langle j|_{\mathcal{X}}\otimes\Psi(|i\rangle\langle j|)_{\mathcal{YZ}}. \tag{12}$$

Using the facts that $\langle\phi_x|i\rangle=\langle i|\phi_x\rangle$ for all states in the ensemble, and that Ψ is a linear map, it can be shown that

$$\text{Tr}_{\mathcal{XYZ}}\left[\left(\phi_x^{\mathcal{X}}\otimes\phi_x^{\mathcal{Y}}\otimes\mathbb{1}^{\mathcal{Z}}\right)J(\Psi)\right]=\langle\phi_x|\eta_x|\phi_x\rangle_{\mathcal{Y}}, \tag{13}$$

and similarly that

$$\text{Tr}_{\mathcal{XYZ}}\left[\left(\phi_x^{\mathcal{X}}\otimes\mathbb{1}^{\mathcal{Y}}\otimes\phi_x^{\mathcal{Z}}\right)J(\Psi)\right]=\langle\phi^x|\tau_x|\phi^x\rangle_{\mathcal{Z}}, \tag{14}$$

where here, for ease of notation, we have used the superscript to denote the relevant Hilbert space. With this we can rewrite the problem in Eq. (11) as the problem of finding the operator $J(\Psi)$ which maximizes

$$\frac{1}{2^{n+1}}\sum_{x=1}^{2^n}\text{Tr}_{\mathcal{XYZ}}\times\left[\left(\left(\phi_x^{\mathcal{X}}\otimes\phi_x^{\mathcal{Y}}\otimes\mathbb{1}^{\mathcal{Z}}\right)+\left(\phi_x^{\mathcal{X}}\otimes\mathbb{1}^{\mathcal{Y}}\otimes\phi_x^{\mathcal{Z}}\right)\right)J(\Psi)\right]. \tag{15}$$

The conditions that the channel must be completely positive and trace preserving lead to the conditions that $J(\Psi)$ must be positive semidefinite and $\text{Tr}_{\mathcal{YZ}}(J(\Psi))=\mathbb{1}_{\mathcal{X}}$. Written in standard form, the semidefinite program corresponding to the maximum average fidelity is given by

$$\begin{aligned}&\text{maximize:}\quad\langle Q(n),X\rangle,\\&\text{subject to:}\quad\text{Tr}_{\mathcal{YZ}}(X)=\mathbb{1}_{\mathcal{X}},\\&\qquad\qquad X\geqslant 0,\end{aligned} \tag{16}$$

where

$$Q(n)=\frac{1}{2^{n+1}}\sum_{x=1}^{2^n}\left(\left(\phi_x^{\mathcal{X}}\otimes\phi_x^{\mathcal{Y}}\otimes\mathbb{1}^{\mathcal{Z}}\right)+\left(\phi_x^{\mathcal{X}}\otimes\mathbb{1}^{\mathcal{Y}}\otimes\phi_x^{\mathcal{Z}}\right)\right). \tag{17}$$

The dual problem is simply

$$\begin{aligned}&\text{minimize:}\quad\text{Tr}(Y),\\&\text{subject to:}\quad\mathbb{1}_{\mathcal{YZ}}\otimes Y\geqslant Q(n),\\&\qquad\qquad Y\in\text{Herm}(\mathcal{X}),\end{aligned} \tag{18}$$

since $\langle\mathbb{1}_{\mathcal{X}},Y\rangle=\text{Tr}(Y)$ and the adjoint of the partial trace is the extension by the identity. The dual problem approaches the optimal value from above, so any feasible point (i.e., any operator Y that satisfies the constraints of the dual problem) gives us an upper bound on the maximum average fidelity. A feasible point can easily be found in terms of the matrix $Q(n)$ as

$$Y=||Q(n)||_{\infty}\mathbb{1}_{\mathcal{X}} \tag{19}$$

so that we arrive at the following upper bound on the average fidelity:

$$\overline{F}\leqslant n||Q(n)||_{\infty}. \tag{20}$$

Thus, for quantum money protocols using states of dimension n and a maximal disjoint set of matchings, we can upper bound the error probability of the adversary in terms of the operator norm of $Q(n)$. Computing this norm for different values of n leads to the bound

$$\overline{F}\leqslant\frac{1}{2}+\frac{1}{n}, \tag{21}$$

which we have verified numerically for $n\leqslant 14$ and we conjecture holds for any n. From now on, we simply assume that $n\leqslant 14$. The analysis above enables us to restrict the achievable error probabilities for the two verifiers on a single game as

$$p_{\text{Ver}_1}=\frac{n}{2(n-1)}(1-F),\quad p_{\text{Ver}_2}=\frac{n}{2(n-1)}(1-G) \tag{22}$$

subject to

$$\frac{1}{2}(F+G)\leqslant\frac{1}{2}+\frac{1}{n}, \tag{23}$$

which leads to

$$p_{\text{Ver}_1}+p_{\text{Ver}_2}\geqslant\frac{1}{2}-\frac{1}{2(n-1)}. \tag{24}$$

Until now, we have considered only a single white state out of the l games used in the verification protocol. Let us now consider l such games, and let $p^{(i)}_{\text{Ver}_j}$ be the error probability for honest verifier j on the ith run of the verification protocol. We claim that when we have l independent white states (in the sense that each x^i is chosen independently), it is still the case that

$$p^{(i)}_{\text{Ver}_1}+p^{(i)}_{\text{Ver}_2}\geqslant\frac{1}{2}-\frac{1}{2(n-1)} \tag{25}$$

for all i, regardless of the outcomes of previous measurements made by the verifiers. Though intuitively reasonable, this claim is far from trivial, but can be proved using a teleportation argument due to Croke and Kent [38] (see Appendix) so that, essentially, we can imagine the adversary acts independently on each game in the verification protocol. Therefore, on each and every white state, at least one verifier must have an error probability of at least

$$\frac{1}{2}\left(p^{(i)}_{\text{Ver}_1}+p^{(i)}_{\text{Ver}_2}\right)=\frac{1}{4}-\frac{1}{4(n-1)}. \tag{26}$$

Overall, if we include the effects of r register manipulation and auxiliary verifications, at least one verifier, say Ver_1, must have an average error probability over all l games of at least

$$e_{\min}=\frac{997}{999}\left(\frac{1}{4}-\frac{1}{4(n-1)}\right)\approx\frac{1}{4}-\frac{1}{4(n-1)}. \tag{27}$$

Using Hoeffding's inequality, the probability of both verifiers accepting the coin can be bounded as

$$\begin{aligned}&\mathbb{P}(\text{both Ver}_1\text{ and Ver}_2\text{ generate outcome ``valid''})\\&\leqslant\mathbb{P}(\text{Ver}_1\text{ generates outcome ``valid''})\\&\leqslant e^{-2l\delta^2},\end{aligned} \tag{28}$$

where $\delta = (e_{\min} - \beta)/2$, as above. As long as $\beta < e_{\min}$, the Hoeffding bound can be used to show that it becomes exponentially unlikely for both verifiers to pass the verification protocol. By increasing the maximum noise tolerance of the protocol we increase the size of δ, thereby allowing smaller sample sizes in the verification protocol, which increases the reusability of coins. If we choose $n = 4$, our scheme would be able to tolerate 16.6% noise, and for $n = 14$ it can tolerate up to 23% noise. This concludes the proof of security against forging.

In the next section, we prove an upper bound on the error tolerance achievable for a general class of classical verification quantum money schemes, and show this bound limits to 25% as the dimension of the underlying states is increased. This implies that our protocols are nearly optimal in terms of error tolerance. When proving this result, we assume only that the coin is a collection of quantum states each identified with a secret classical string, and that to verify the coin the holder must declare a number of single bit values which can be checked against the classical record.

III. MAXIMUM ACHIEVABLE NOISE TOLERANCE

Suppose we have a scheme in which the coin consists of many independently chosen n-dimensional pure quantum states, $\phi_x = |\phi_x\rangle\langle\phi_x|$, with $x \in X$ and where x is a classical bit string chosen according to some probability distribution. To verify each state, the holder performs some positive operator-valued measure (POVM), $\mathcal{M}_x = \{M_x^{\text{cor}}, M_x^{\text{inc}}\}$, to ascertain one bit of information about each of the states used in the verification protocol. The bit values resulting from the measurement outcomes are checked against a classical record to verify whether or not the coin is genuine.

Lemma 1. For any quantum money scheme of the above type, the maximum tolerable noise $e_{\max}$ must be less than

$$e_{\max} \leqslant \frac{1}{2} - \frac{1}{4}\frac{n+2}{n+1}. \tag{29}$$

Proof. We prove this by explicitly illustrating a strategy available to the adversary. The adversary holds the unknown state ϕ_x, which lives in Hilbert space $\mathcal{H}$. She extends the state to $\phi_x \otimes \Phi$, where $\Phi = \frac{1}{n}\mathbb{1}_n$, and symmetrises the system. Specifically, she performs the mapping

$$\phi_x \otimes \Phi \to S_2(\phi_x \otimes \Phi)S_2, \tag{30}$$

where S_2 is the projector onto $\mathcal{H}_+^2$, the symmetric subspace of $\mathcal{H}^{\otimes 2}$, and where the state on the right-hand side is not normalized. The resulting normalized state of each clone is [39]

$$\eta_x = v\phi_x + (1-v)\Phi, \tag{31}$$

where $v := \frac{1}{2}\frac{n+2}{n+1}$. By the correctness requirement of quantum money schemes, an honest measurement on the correct state should always give a correct answer so that the coin is declared valid, i.e.,

$$\text{Tr}\left(M_x^{\text{cor}}\phi_x\right) = 1. \tag{32}$$

We further assume that, without access to the state ϕ_x, the adversary has no information on x and can do no better than to guess randomly. This means her probability of declaring a correct bit value is $1/2$, i.e., [40]

$$\text{Tr}\left(M_x^{\text{cor}}\Phi\right) = 1/2. \tag{33}$$

Both honest verifiers hold the state η_x. Using Eqs. (32) and (33), the probability that an honest verifier gets a correct measurement outcome is

$$\begin{aligned}\text{Tr}\left(M_x^{\text{cor}}\eta_x\right) &= v\text{Tr}\left(M_x^{\text{cor}}\phi_x\right) + (1-v)\text{Tr}\left(M_x^{\text{cor}}\Phi\right)\\ &= v + \frac{(1-v)}{2}.\end{aligned} \tag{34}$$

Expressing v in terms of the dimension of the system shows that this strategy (which is always available to the adversary) leads to the honest verifiers finding an error rate of

$$e_{\max} = \frac{1}{2} - \frac{1}{4}\frac{n+2}{n+1}, \tag{35}$$

and so for any such scheme to be secure an honest participant must expect an error rate less than $e_{\max}$ in an honest run of the protocol.

Our analysis shows that for any scheme with $n = 4$ the tolerable noise is at most 20%, which complements our results in Sec. II where we described a protocol with $n = 4$ which tolerated noise up to 16.6%. For $n = 14$, the bound in this section shows that any such scheme has a noise tolerance of at most 23.3%. For $n = 14$, our protocol can achieve an error tolerance of 23.03%, and so it is nearly optimal. As we increase the dimension of the quantum states used for the coins, the upper bound on the tolerable noise approaches 25% which coincides with our conjecture for the tolerable noise in our protocols above (see Fig. 5).

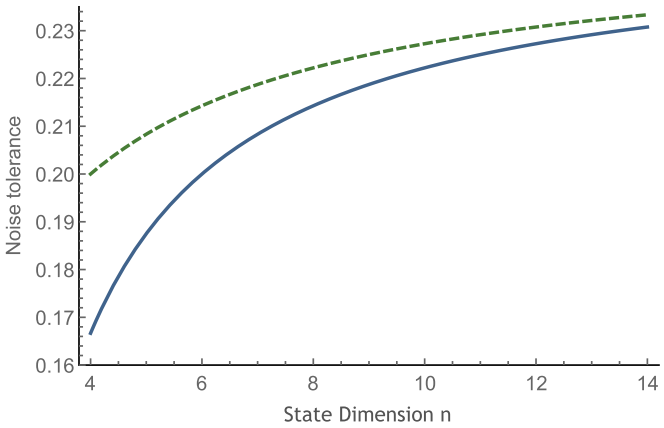

FIG. 5. Plot showing the theoretical bound on protocol noise tolerance (dotted line) and the noise tolerance achieved by the protocols in Sec. II (bold line) as the dimension of the underlying systems increase.

IV. EXPERIMENTAL IMPLEMENTATION

The protocol presented in Sec. II gives rise to three main technical challenges when one considers experimental implementations, namely, the security analysis provided does not account for losses; the bank requires a source of complex, high-dimensional states; and the protocol requires that the

coin holders have the ability to store states in quantum memory. In this section we address the first two issues so that a proof-of-principle implementation of the verification algorithm of the quantum money schemes could be performed with current technology.

A. Detector losses

Here we tackle the first of the issues, and consider an implementation in which the verifiers use imperfect detectors with efficiency η. We assume that all detector losses are random and cannot be manipulated by the adversary. In this paper we do not consider channel loss, as we assume that coin transfers occur over short distances, meaning channel losses are less relevant. Nevertheless, many of the methods presented here would remain valid in the presence of small channel loss with only minor modifications necessary. Note that detectors are employed by the holder and not the bank.

To incorporate detector loss, it is necessary to modify the verification protocol, previously stated in Sec. II, so that it becomes the following.

Ver algorithm

(1) The holder randomly chooses a subset of indices, $L \subset [q]$, with $l = |L|$, such that $r_i = 0$ for each $i \in |$. The indices $i \in L$ specify the selection of games G_i which will be used as tests for the verification procedure. For each $i \in L$, the holder then sets the corresponding bit of r to be 1 so that this game cannot be used in future verifications.

(2) For each $i \in L$, the holder picks a relation σ_i' at random from $\mathcal{R}$ and applies the appropriate measurement to get answer d_i. If there is no measurement outcome we say the measurement was unsuccessful and set $d_i = \emptyset$. We define the number of successful measurement outcomes to be l'.

(3) If $l' < l_{\min} := (\eta - \epsilon)l$, where $\epsilon > 0$ is a small security parameter, the verifier aborts the protocol.

(4) The holder sends all triplets (i,σ_i',d_i) to the bank.

(5) The bank checks that $s < T$, where T is the predefined maximum number of allowed verifications for the coin. If $s = T$, the bank declares the coin as invalid.

(6) For each i, the bank checks whether the answer is correct by comparing (i,σ_i',d_i) to the secret x^i values. The bank ignores those outcomes for which $d_i = \emptyset$, and accepts the coin as valid only if more than $l'(c - \delta)$ of the answers are correct, where $c = 1 - \beta$ is a measure of the channel correctness and δ is a small positive constant.

(7) The bank updates s to $s + 1$.

1. Correctness

Correctness of the scheme follows from Hoeffding's inequality. When all participants are honest, it is exponentially unlikely for l' to be less than $l_{\min}$, so the protocol will not abort, except with a negligible probability. If the protocol does not abort, the verifier has at least $l_{\min}$ successful measurement outcomes, each with an independent probability c of being correct. Overall, the probability of the verification failing is bounded by

$$\mathbb{P}(\text{Ver fails}) \leqslant \exp[-2l_{\min}\delta^2] + \exp[-2l\epsilon^2], \tag{36}$$

where now $\delta = (e'_{\min} - \beta)/2$, with $e'_{\min}$ derived in Eq. (40) below as the minimum average error rate achievable by the adversary.

2. Unforgeability

Since the protocol now includes detector losses, the adversary may not have to send states to each verifier for each game in the verification protocol, and she could attempt to hide losses arising from her strategy in the losses arising from detector inefficiency. As a consequence, the set of strategies available to the adversary is increased, and we must make sure our arguments in Sec. II still apply.

Let U_1 and U_2 be q-bit strings representing whether or not the adversary sent a state to Ver_1 and Ver_2 respectively, for each of the q games created by the bank. An entry of 1 means the adversary sent a state to the verifier, while an entry of 0 means the adversary did not send a state to the verifier. We want to show that, in order for the protocol not to abort, $W(U_i) \geqslant \gamma q$, where $\gamma := 1 - \frac{3\epsilon}{\eta}$ and W is the Hamming weight. Suppose $W(U_i) = \gamma q$. Then, in step 1 of the verification protocol, Ver_i takes a sample, V_i, consisting of l of the entries of U_i. Hoeffding's inequality gives

$$P\left[(V_i) \leqslant \left(\gamma + \frac{\epsilon}{\eta}\right)l\right] \geqslant 1 - \exp\left[-2\frac{\epsilon^2}{\eta^2}l\right]. \tag{37}$$

If $W(V_i) \leqslant (\gamma + \frac{\epsilon}{\eta})l$, then the probability of at least $l_{\min}$ successful measurement outcomes is given by

$$P\left[\text{at least } l_{\min} \text{ succ. meas.} \mid W(V_i) \leqslant \left(\gamma + \frac{\epsilon}{\eta}\right)l\right] \leqslant \exp[-2l\epsilon^2]. \tag{38}$$

The probability of the protocol proceeding past step 3 of verification is therefore

$$P[\text{no abort}|W(U_i) = \gamma q] \leqslant \exp\left[-2\frac{\epsilon^2}{\eta^2}l\right] + \exp[-2\epsilon^2 l]. \tag{39}$$

In what follows we assume $W(U_i) \geqslant \gamma q$, since otherwise the above shows that the verifiers will abort with near certainty. This means the adversary is able to use any strategy that leads to channel losses of at most $\frac{3\epsilon}{\eta}$ for each verifier, as these can be hidden within the normal fluctuations of detector loss. Suppose there is a strategy which gives at least $(1 - \frac{3\epsilon}{\eta})q$ states to each verifier, and which leads to an average error probability (on only the states tested) of $e'_{\min}$ for at least one of the verifiers. Then, there is a strategy which gives q states to each verifier, and leads to an average error probability for at least one of the verifiers of $(1 - \frac{3\epsilon}{\eta})e'_{\min} + \frac{3\epsilon}{2\eta}$ (the adversary simply sends the maximally mixed state to each verifier in place of the $\frac{3\epsilon}{\eta}$ losses). Since this strategy falls under the scope of the analysis in Sec. II, we know that the resulting error rate must be at least $e_{\min}$, which means

$$e'_{\min} \geqslant \frac{e_{\min} - \frac{3\epsilon}{2\eta}}{1 - \frac{3\epsilon}{\eta}}. \tag{40}$$

The parameter ϵ can be chosen to be arbitrarily small by increasing the sample size l. As such, the protocol is able

to handle arbitrarily large detector losses, and leads to noise tolerance that can be kept arbitrarily close to the noise tolerance derived for the case of perfect detectors.

Each verifier tests at least $l_{\min}$ states, and at least one verifier expects an error rate of $e'_{\min}$. The probability of this verifier passing the test is bounded as

$$P(\text{error rate} < e'_{\min} - \delta) \leqslant \exp[-2l_{\min}\delta^2]. \tag{41}$$

Combining Eqs. (39) and (41), the probability that the adversary is able to forge a coin is given by

$$P(\text{forgery}) \leqslant \exp\left[-2\frac{\epsilon^2}{\eta^2}l\right] + \exp[-2l\epsilon^2] + \exp[-2l_{\min}\delta^2]. \tag{42}$$

B. Coherent state implementation

In this section we tackle the second issue arising when considering experimental realizations of the scheme—the bank must create hidden matching states of the form in Eq. (1), which are high-dimensional states of high complexity. The implementation of hidden matching quantum retrieval games has been studied extensively in Ref. [34], where the coherent state mapping defined in Ref. [41] was used to approximate each hidden matching state by a sequence of n coherent states of the form

$$|\alpha,x\rangle = e^{-|\alpha|^2/2}\sum_{k=0}^{\infty}\frac{\alpha^k}{k!}(a_x^\dagger)^n|0\rangle = \bigotimes_{i=1}^{n}\left|(-1)^{x_i}\frac{\alpha}{\sqrt{n}}\right\rangle, \tag{43}$$

where

$$a_x^\dagger = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}(-1)^{x_i}b_i^\dagger \tag{44}$$

and $\{b_1^\dagger, b_2^\dagger, \ldots, b_n^\dagger\}$ are the creation operators of the n modes. We call each sequence of coherent states a block, so that a single block is used to approximate a hidden matching state. As outlined in Ref. [34], Bob's measurement can then be performed using linear optics circuits and single-photon detectors.

In the absence of a phase reference, the phase of each block is randomized, which implies that each block is equivalent to a classical mixture of number states [42]. More specifically, writing $\alpha = e^{i\theta}|\alpha|$, we have

$$\int_0^{2\pi}\frac{d\theta}{2\pi}|\alpha,x\rangle\langle\alpha,x| = e^{-|\alpha|^2}\sum_{k=0}^{\infty}\frac{|\alpha|^{2k}}{k!}|k\rangle\langle k|_x, \tag{45}$$

where $|k\rangle\langle k|_x$ is a state of k photons in the mode $a_x^\dagger$. Thus, the probability of obtaining a particular number of photons depends only on α, which is a free parameter within the coherent state mapping. We consider the following three cases.

1. Zero photons in the block

In this case the state emitted is simply the vacuum state. If the adversary chooses to forward a state on to the verifiers, she can do no better than to induce a 50% error rate, and it is simple to show that it is never beneficial for her to do so. This scenario can therefore be considered a "source" loss, as opposed to a channel or detector loss. Crucially, since these losses are not controllable by the adversary, they can be treated in the same manner as detector losses in Sec. IV A simply by including the source loss into the detector loss parameter η. The probability of zero photons being emitted is $p_0 = e^{-|\alpha|^2}$.

2. One photon in the block

In this case, the state emitted is equivalent to the ideal hidden matching state in Eq. (1) since

$$|1\rangle_x = a_x^\dagger|0\rangle = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}b_i^\dagger|0\rangle = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}(-1)^{x_i}|i\rangle, \tag{46}$$

where $|i\rangle$ is a single-photon state in the mode b_i. Therefore, whenever the bank's source emits a single photon, the analysis in Sec. II applies. The probability of one photon being emitted is $p_1 = |\alpha|^2e^{-|\alpha|^2}$.

3. More than one photon in the block

In this case we assume the worst case scenario: whenever the source emits more than one photon to represent a hidden matching state, the adversary can perfectly forge that state. The resulting error rate for the adversary is $e'_{\min}(\frac{p_1}{p_1+p_{2+}})$, where $p_{2+} = 1 - p_0 - p_1$. For small $|\alpha|$, $p_{2+} \approx \frac{|\alpha|^4}{2}$, while $p_1 \approx |\alpha|^2$, so that $p_{2+} \ll p_1$ and the adversary's error probability is almost unchanged by using coherent states.

V. CONCLUSION

We presented a family of unconditionally secure classical verification quantum money schemes which are tolerant to noise up to 23%, and which we conjecture tolerate noise up to 25%. We further proved that 25% is the maximum noise tolerance achievable for a wide class of quantum money schemes, including all classical verification secret-key schemes previously proposed. The security of our schemes depends on the difference between maximum tolerable noise and expected noise, meaning the increase in maximum tolerable noise increases the efficiency of our scheme, allowing for smaller, more reusable coins. The techniques we use to prove security differ considerably to previous papers, and the reusability of our coins is optimal [23] in that it scales linearly with the number of qubits in the coin. This is a significant improvement when compared to Ref. [24], in which the reusability scales as $q^{1/3}$, and Ref. [23], in which reusability scales as $q^{1/4}$, where q is the total number of qubits in the coin. With realistic assumptions on experimental equipment, we expect that, using $n = 8$, a coin containing 10^9 qubits would use $l = 18\,000$ states for each verification, and would be reusable $T = 100$ times for a security level of 10^{-6}. Last, we suggested methods of adapting our techniques to facilitate experimental implementations of the scheme. We show that the schemes can be implemented using weak coherent states even in the presence of limited detector efficiency.

ACKNOWLEDGMENTS

We would like to thank I. Kerenidis, E. Andersson, and A. Ignjatovic for helpful discussions. R.A. gratefully acknowledges EPSRC studentship funding under Grant

No. EP/L015110/1. J.M.A. recognizes funding from the Singapore Ministry of Education (partly through the Academic Research Fund Tier 3 MOE2012-T3-1-009) and the National Research Foundation of Singapore, Prime Ministers Office, under the Research Centres of Excellence programme.

APPENDIX

1. Overview of argument

In the main paper, we claim that the adversary cannot use coherent attacks on multiple states in order to beat the bound given in Eq. (24), even when conditioned on the states chosen by the bank, and on the outcomes of previous measurement results found by the verifiers. In this section we formally prove our claim using a teleportation argument similar to the one introduced by Croke and Kent in Ref. [38], so that each game can essentially be viewed as independent of all others.

In order to apply the teleportation argument, we must first introduce a modified individual setting, in which the adversary is allowed an additional ability. We show that this modification does not help the adversary to cheat. We then show that any coherent strategy can be transformed into a modified individual strategy. Therefore, any coherent strategy cannot beat the bounds proved for the unmodified individual case, as claimed.

2. Modified individual attacks

In the individual setting, the verifiers each receive a single hidden matching state and apply the verification protocol to test its authenticity. As specified by the protocol, the verifiers randomly choose to measure the state they receive using one of the matching measurements. We include this random choice of matching into the mathematical description of the measurement, and group the outcomes to be either "correct" or "incorrect." It can be shown that if the bank creates $\phi_x = |\phi_x\rangle\langle\phi_x|$, the verifiers' measurement is described by the POVM,

$$\Gamma_x = \{\Gamma^{\mathrm{cor},x}, \Gamma^{\mathrm{inc},x}\} = \frac{n}{2(n-1)}\left\{\frac{n-2}{n}\mathbb{I} + \phi_x, \mathbb{I} - \phi_x\right\}. \tag{A1}$$

Suppose now the adversary has the additional power of being able to force the verifiers to apply a correction unitary (which will be the teleportation corrections) to their measurement outcomes before they are sent to the bank. The adversary must specify the correction operation before sending the states to the verifiers, and, crucially, the correction operation is such that it is simply a permutation of the set of hidden matching states. For example, suppose the teleportation operation takes input $|\phi_x\rangle$ and outputs $|\phi_{x'}\rangle$, with correction operator C. In this case, before sending the states, the adversary will tell the verifiers that they must apply correction C to their measurement outcomes. In effect then, the verifiers will measure

$$\begin{aligned}\Gamma_{x'} &= \{\Gamma^{\mathrm{cor},x'}, \Gamma^{\mathrm{inc},x'}\} \\ &= \frac{n}{2(n-1)}\left\{\frac{n-2}{n}\mathbb{I} + \phi_{x'}, \mathbb{I} - \phi_{x'}\right\},\end{aligned} \tag{A2}$$

since the correction applied to $\Gamma^{\mathrm{inc},x'}$ is $\Gamma^{\mathrm{inc},x}$. On average, given ϕ_x, it is not possible for the adversary to create two states, η_x and τ_x, such that $\mathrm{Tr}[\Gamma^{\mathrm{inc},x'}(\eta_x + \tau_x)] < p$. If it were possible, then it would imply that the adversary can clone $\phi_{x'}$ better than what is allowed by quantum mechanics (and our arguments in the main paper). This is because if the adversary was given $\phi_{x'}$ he could easily transform it to ϕ_x by applying C, and then perform the strategy to get two copies with a fidelity higher than the bound proved in the main paper. Therefore the additional power given to the adversary does not allow her to decrease the value of $p_{\mathrm{Ver}_1} + p_{\mathrm{Ver}_2}$.

3. Coherent strategy

We now consider the case of N games created by the bank. The bank creates

$$\begin{aligned}&\frac{1}{2^{Nn}}\sum_{x_1,x_2}|x_1\rangle\langle x_1|_{X_1}\otimes|x_2\rangle\langle x_2|_{X_2} \\ &\quad\otimes|\phi_{x_1}\rangle\langle\phi_{x_1}|_A\otimes|\phi_{x_2}\rangle\langle\phi_{x_2}|_B.\end{aligned} \tag{A3}$$

The X_1 and A registers contain the first $N-1$ secret strings selected by the bank and the corresponding hidden matching states, respectively. The X_2 and B registers contain the Nth secret string selected by the bank and its corresponding hidden matching state. Only the A and B registers are accessible to the adversary. We assume for a contradiction that there exists a strategy available to the adversary such that, conditional on the value in the X_1 register, and conditional on the verifiers obtaining specific outcomes in previous measurements, the value of $p_{\mathrm{Ver}_1} + p_{\mathrm{Ver}_2}$ in the Nth game is decreased below the bound in Eq. (24).

We describe this strategy as follows—upon receiving the states from the bank, the adversary applies the unitary operation S_{ABC} so that the state becomes

$$\begin{aligned}&\frac{1}{2^{Nn}}\sum_{x_1,x_2}|x_1\rangle\langle x_1|_{X_1}\otimes|x_2\rangle\langle x_2|_{X_2} \\ &\quad\otimes S_{ABC}(|\phi_{x_1}\rangle\langle\phi_{x_1}|_A\otimes|\phi_{x_2}\rangle\langle\phi_{x_2}|_B\otimes|0\rangle\langle 0|_C)S^\dagger_{ABC} \\ &= \frac{1}{2^{Nn}}\sum_{x_1,x_2}|x_1\rangle\langle x_1|_{X_1}\otimes|x_2\rangle\langle x_2|_{X_2} \\ &\quad\otimes|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|_{AA'BB'C'}.\end{aligned} \tag{A4}$$

The A, A' registers are the spaces that contain the states that will be sent to Ver_1 and Ver_2 (respectively) for the first $N-1$ games. The B, B' registers are the spaces that contain the states that will be sent to Ver_1 and Ver_2 (respectively) for the Nth game. The C registers are auxiliary registers held by the adversary. We assume that the bank measures the X_1 register, and gets a state, x_1, which satisfies the conditions in the assumption. The state held by the adversary is then

$$\frac{1}{2^n}\sum_{x_2}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|. \tag{A5}$$

The adversary gives the A, A', B, B' parts of the state to the verifiers. The honest verifiers will first make measurements on systems A, A' and a possible postmeasurement state is

$$\frac{1}{2^n}\sum_{x_2}a_{x_1x_2}\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}. \tag{A6}$$

We assume that $\Pi_{AA'}$ is a measurement outcome satisfying the conditions of the assumption, so that the error probabilities on the Nth game are decreased. Here $a_{x_1x_2}$ is the normalization term, $a_{x_1x_2} = 1/\mathrm{Tr}[\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}]$.

The verifiers now each measure Γ_{x_2}, as defined in Eq. (A2), on their B system. The assumption tells us that

$$\frac{1}{2^n}\sum_{x_2}\Big[a_{x_1x_2}\mathrm{Tr}\Big[\Gamma_B^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}\Big] + a_{x_1x_2}\mathrm{Tr}\Big[\Gamma_{B'}^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}\Big]\Big] < p. \quad \text{(A7)}$$

We now aim to prove that this leads to a contradiction.

4. Teleportation strategy

Supposing the above strategy exists, we explore what this enables the adversary to do in the individual case in the hopes of finding a contradiction. We suppose the bank creates

$$\frac{1}{2^n}\sum_{x_2}|x_2\rangle\langle x_2|_{X_2}\otimes|\phi_{x_2}\rangle\langle\phi_{x_2}|_B \quad \text{(A8)}$$

and sends the B part to the adversary. The adversary can simulate the above strategy locally, by creating $|x_1\rangle$, $|\phi_{x_1}\rangle$ and the maximally mixed state on n dimensions $|\Phi\rangle$. After relabelling the registers, the adversary holds the state

$$\frac{1}{2^n}\sum_{x_2}|x_1\rangle\langle x_1|_{X_1}\otimes|x_2\rangle\langle x_2|_{X_2}\otimes|\phi_{x_1}\rangle\langle\phi_{x_1}|_A \otimes|\phi_{x_2}\rangle\langle\phi_{x_2}|_D\otimes|0\rangle\langle 0|_C\otimes|\Phi\rangle\langle\Phi|_{BE}. \quad \text{(A9)}$$

To simulate the strategy in the previous section, the adversary applies S to the A, B, and C registers, followed by a measurement on the resulting A,A' registers. Conditional on measurement outcome $\Pi_{AA'}$, she then applies a generalized Bell measurement on the D and E registers in order to teleport the unknown state $|\phi_{x_2}\rangle$ into the B register which was acted on by S (modulo a teleportation correction). If the appropriate measurement outcome is not found, the adversary does not perform the Bell measurement and instead starts again. The resulting state is

$$\frac{1}{2^n}\sum_{x_2}a_{x_1x_2'}\Pi_{AA'}|\Psi^{x_1x_2'}\rangle\langle\Psi^{x_1x_2'}|\Pi^\dagger_{AA'}. \quad \text{(A10)}$$

Notice the state contains x_2' since the Bell measurement does not faithfully teleport the state, and a correction is required which we have not performed. If the dimension of the hidden matching states is a power of 2, the correction operators are simply tensor products of the Pauli operators [43]. Crucially, all corrections define a bijective mapping between x_2' and x_2, so that as x_2 cycles over all possible values so does x_2', and the probabilities are not affected (all corrections are equally likely, which must be the case so that information is not communicated faster than light).

The state in Eq. (A10) is the same as the state in Eq. (A6), but the measurements applied by the verifiers are correlated with the X_2 register held by the bank. Therefore, the verifiers' failure probabilities are not the same when measuring the two states. Measurements on the state in Eq. (A6) leads to a failure probability of

$$\frac{1}{2^n}\sum_{x_2}\Big[a_{x_1x_2}\mathrm{Tr}\Big[\Gamma_B^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}\Big] + a_{x_1x_2}\mathrm{Tr}\Big[\Gamma_{B'}^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2}\rangle\langle\Psi^{x_1x_2}|\Pi^\dagger_{AA'}\Big]\Big], \quad \text{(A11)}$$

while measurements on the state in Eq. (A10) lead to a failure probability of

$$\frac{1}{2^n}\sum_{x_2}\Big[a_{x_1x_2'}\mathrm{Tr}\Big[\Gamma_B^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2'}\rangle\langle\Psi^{x_1x_2'}|\Pi^\dagger_{AA'}\Big] + a_{x_1x_2'}\mathrm{Tr}\Big[\Gamma_{B'}^{\mathrm{inc},x_2}\,\Pi_{AA'}|\Psi^{x_1x_2'}\rangle\langle\Psi^{x_1x_2'}|\Pi^\dagger_{AA'}\Big]\Big], \quad \text{(A12)}$$

the difference being the appearance of x_2' in the second expression. Nevertheless, the two can be made equal if the verifiers are forced to apply the teleportation correction unitary to their measurement outcomes. In effect, this correction relabels the measurement outcomes so that $\Gamma^{\mathrm{inc},x_2}\to\Gamma^{\mathrm{inc},x_2'}$. Following this correction, the two expressions (A11) and (A12) are equal. This shows that the assumption in Eq. (A7) leads to a contradiction, since it shows an individual attack in the modified scenario can achieve the same error probability as a coherent attack, and the error probabilities achievable in the modified individual scenario are the same as for the unmodified individual scenario.

[1] C. H. Bennett and G. Brassard, Theor. Comput. Sci. **560**, 7 (2014).

[2] A. Broadbent and C. Schaffner, Sesig. Codes Cryptogr. **78**, 351 (2016).

[3] R. Amiri, P. Wallden, A. Kent, and E. Andersson, Phys. Rev. A **93**, 032325 (2016).

[4] J. M. Arrazola, P. Wallden, and E. Andersson, Quantum Inf. Comput. **16**, 0435 (2016).

[5] R. Amiri and E. Andersson, Entropy **17**, 5635 (2015).

[6] T. Lunghi, J. Kaniewski, F. Bussières, R. Houlmann, M. Tomamichel, A. Kent, N. Gisin, S. Wehner, and H. Zbinden, Phys. Rev. Lett. **111**, 180504 (2013).

[7] C. Erven, N. Ng, N. Gigov, R. Laflamme, S. Wehner, and G. Weihs, Nat. Commun. **5**, 3418 (2014).

[8] N. H. Y. Ng, S. K. Joshi, C. C. Ming, C. Kurtsiefer, and S. Wehner, Nat. Commun. **3**, 1326 (2012).

[9] A. Broadbent, J. Fitzsimons, and E. Kashefi, in *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS'09* (IEEE, Piscataway, NJ, 2009), pp. 517–526.

[10] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, Science **335**, 303 (2012).

[11] B. Sanguinetti, G. Traverso, J. Lavoie, A. Martin, and H. Zbinden, Phys. Rev. A **93**, 012336 (2016).

[12] B. A. Bash, A. H. Gheorghe, M. Patel, J. L. Habif, D. Goeckel, D. Towsley, and S. Guha, Nat. Commun. **6**, 8626 (2015).

[13] J. M. Arrazola and V. Scarani, Phys. Rev. Lett. **117**, 250503 (2016).

[14] K. Bradler, T. Kalajdzievski, G. Siopsis, and C. Weedbrook, arXiv:1607.05916.
[15] X. Ma, X. Yuan, Z. Cao, B. Qi, and Z. Zhang, Npj Quantum Inf. **2**, 16021 (2016).
[16] B. Sanguinetti, A. Martin, H. Zbinden, and N. Gisin, Phys. Rev. X **4**, 031056 (2014).
[17] T. Lunghi, J. B. Brask, C. C. W. Lim, Q. Lavigne, J. Bowles, A. Martin, H. Zbinden, and N. Brunner, Phys. Rev. Lett. **114**, 150501 (2015).
[18] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Phys. Rev. Lett. **87**, 167902 (2001).
[19] J. M. Arrazola and N. Lütkenhaus, Phys. Rev. A **89**, 062305 (2014).
[20] F. Xu, J. M. Arrazola, K. Wei, W. Wang, P. Palacios-Avila, C. Feng, S. Sajeed, N. Lütkenhaus, and H.-K. Lo, Nat. Commun. **6**, 8735 (2015).
[21] J.-Y. Guan, F. Xu, H.-L. Yin, Y. Li, W.-J. Zhang, S.-J. Chen, X.-Y. Yang, L. Li, L.-X. You, T.-Y. Chen, Z. Wang, Q. Zhang, and J.-W. Pan, Phys. Rev. Lett. **116**, 240502 (2016).
[22] S. Wiesner, ACM Sigact News **15**, 78 (1983).
[23] D. Gavinsky, in *Proceedings of the 2012 IEEE 27th Annual Conference on Computational Complexity (CCC)* (IEEE, Piscataway, NJ, 2012), pp. 42–52.
[24] M. Georgiou and I. Kerenidis, in *Proceedings of the LIPIcs-Leibniz International Informatics* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2015), Vol. 44.
[25] S. Aaronson, in *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC'09* (IEEE, Piscataway, NJ, 2009), pp. 229–242.
[26] A. Lutomirski, arXiv:1010.0256.
[27] A. Brodutch, D. Nagaj, O. Sattath, and D. Unruh, Quantum Inf. Comput. **16**, 1048 (2016).
[28] Z. Bar-Yossef, T. S. Jayram, and I. Kerenidis, in *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2004), pp. 128–137.
[29] F. Pastawski, N. Y. Yao, L. Jiang, M. D. Lukin, and J. I. Cirac, Proc. Natl. Acad. Sci. USA **109**, 16079 (2012).
[30] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, D. Nagaj, and P. Shor, Phys. Rev. Lett. **105**, 190503 (2010).
[31] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor, in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (ACM, New York, 2012), pp. 276–289.
[32] A. Lutomirski, S. Aaronson, E. Farhi, D. Gosset, A. Hassidim, J. Kelner, and P. Shor, arXiv:0912.3825.
[33] S. Aaronson and P. Christiano, in *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2012), pp. 41–60.
[34] J. M. Arrazola, M. Karasamanis, and N. Lütkenhaus, Phys. Rev. A **93**, 062311 (2016).
[35] More precisely, this is actually the definition of a *perfect* matching.
[36] We are able to apply the arguments in Ref. [23] because, although our scheme uses more than two matchings, when taken pairwise any two matchings within our scheme are independent.
[37] W. Hoeffding, J. Am. Stat. Assoc. **58**, 13 (1963).
[38] S. Croke and A. Kent, Phys. Rev. A **86**, 052309 (2012).
[39] M. Keyl and R. F. Werner, J. Math. Phys. **40**, 3283 (1999).
[40] Note that this assumption holds for all hidden matching quantum money schemes considered, and for any scheme in which the verification protocol involves declaring many single bit values which are later checked. Nevertheless, there may be protocols in which the verification protocol involves checking many m-bit outcomes, in which case the more reasonable assumption would be $\mathrm{Tr}(M_x^{\mathrm{cor}}\Phi) = 1/2^m$. To our knowledge such a scheme does not exist, but if higher error tolerance is desired our proof suggests looking into such schemes.
[41] J. M. Arrazola and N. Lütkenhaus, Phys. Rev. A **90**, 042335 (2014).
[42] G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders, Phys. Rev. Lett. **85**, 1330 (2000).
[43] G. Rigolin, Phys. Rev. A **71**, 032303 (2005).