

Quantum Money with Classical Verification

Dmitry Gavinsky

NEC Laboratories America, Inc.
Princeton, NJ, U.S.A.

Abstract

We propose and construct a quantum money scheme that allows verification through *classical* communication with a bank. This is the first demonstration that a secure quantum money scheme exists that does not require quantum communication for coin verification.

Our scheme is secure against adaptive adversaries – this property is not directly related to the possibility of classical verification, nevertheless none of the earlier quantum money constructions is known to possess it.

1 Introduction

In 1983 Wiesner [Wie83] proposed a new quantum cryptographic scheme, that later became known as *quantum money*. Informally, a *quantum coin* is a unique object that can be created by a trusted *bank*, then circulated among untrusted *holders*.¹ A holder of a coin should be able to verify it, and the verification must confirm that the coin is authentic if it has been circulated according to the prescribed rules. On the other hand, if a holder wants to counterfeit a coin, that is, to create several objects such that each of them would pass verification, he must fail in doing so with overwhelmingly high probability.

Wiesner has demonstrated that quantum mechanics (as opposed to classical physics) allows money schemes, and the basic principle that made such constructions possible was that of *quantum uncertainty*. The principle states that there are properties of a quantum object that are known to its “manufacturer” but cannot be learnt by an observer who measures the object; nevertheless, those properties can be later “verified” by the manufacturer. Accordingly, a bank can prepare objects with this kind of “secret properties” and let the holders use them as quantum coins – not knowing the secrets, untrusted holders would not be able to forge counterfeits.

1.1 Prior work

In Wiesner’s original construction [Wie83, BBBW83] a coin had to be sent back to the bank in order to get verified. This could be viewed as a possible drawback: a coin might get “stolen”, or intentionally “ruined” by an adversary who had access to the communication channel between a coin holder and the bank.

This problem has been addressed in a number of works. The approach taken by Aaronson [Aar09], Lutomirski et al. [LAF⁺10], Farhi et al. [FGH⁺10] and in the upcoming Aaronson and Christiano [AC12] was to allow the holders to verify quantum coins locally, not having to

¹The notation is still unset in this relatively new area of research. In particular, each coin in our construction will have its own identification number, and some authors would call such items *quantum banknotes*, to emphasize the uniqueness. Also, what we call a *bank* is sometimes addressed as a *mint*.

contact the bank. Clearly, in this situation an adversary can, given unlimited computational resources, produce as many counterfeit coins as he wishes (being able to locally verify implies having a complete description of all objects that would pass the verification, so coin forgery becomes an achievable, albeit possibly computationally-expensive task). What is worse, the present state of mathematical development only allows to *conjecture* that certain tasks are hard for a reasonably powerful model of computation, and a major breakthrough would be required to argue that a scheme of this type is secure, say, against an adversary who can use a Turing machine.

In a different line of research, Tokunaga, Okamoto and Imoto [TOI03] and Mosca and Stebila [MS10] considered the problem of creating quantum money that can be used *anonymously*.² In [TOI03] a coin holder introduces some local randomness into the state of a coin to obtain anonymity. In [MS10] the construction allows multiple identical (but still resistant to counterfeiting) instances of quantum coins. In both of these works quantum communication with a bank is required in order to use the scheme ([MS10] discusses the hypothetical possibility of using computational hardness assumptions to allow local verification).

Relatively recently another limitation of all previously known quantum money schemes has been noticed by Aaronson [Aar09] and by Lutimirski [Lut10]: An adversary can gain even more power from interacting *adaptively* with the bank. No quantum money scheme was known to be resistant to this type of attacks; in fact, [Lut10] has shown a very efficient adaptive attack against one version of Wiesner’s scheme (which was unconditionally secure against non-adaptive adversaries).

1.2 Our results

In this work we propose to use *classical communication with a bank* in order to verify a quantum coin. We construct such a scheme. This is the first demonstration that *a secure quantum money scheme exists that does not require quantum communication for coin verification*.

Some advantages of our construction over the previously known ones are:

- Unlike the original scheme of Wiesner and the constructions of Mosca and Stebila, *our construction does not require quantum communication with a bank* in order to verify a coin.
- We *prove* that our scheme is (unconditionally) secure; security arguments for schemes with local verification require either unproved hardness assumptions or a major mathematical breakthrough (complexity lower bounds). Moreover, to the best of our knowledge, no such scheme has been shown to be secure under so-called “widely believed” unproved assumptions.³
- Unlike the schemes with local verification, our construction remains *secure against computationally unlimited adversary* who obeys the laws of quantum mechanics.

Besides offering possible practical advantages, the concept of quantum money with classical verification gives rise to natural and attractive theoretical questions.

Another advantage of our construction is not directly related to the possibility of quantum verification:

- Our scheme remains secure against an adversary who uses *adaptive multi-round attacks*; no such scheme was known before.

²Note that locally verifiable coins can be viewed as a partial answer to this requirement: when the bank isn’t involved in the verification procedure it cannot “trace” the transactions.

³It has happened that a proposed scheme was broken soon after its publication.

Note that adaptive multi-round attacks are also conceivable in the case of money schemes with quantum verification alone: an adversary can, for example, “split” a coin into two “fragments”, send one of them to the bank and collect the response, and later use the remaining fragment in a way that would depend on the bank’s response to the first fragment. Indeed, Lutomirski [Lut10] has demonstrated a linear-time adaptive attack against one version of Wiesner’s scheme, which was provably secure against non-adaptive adversaries. Before this work it was open whether any quantum money scheme can be resistant to adaptive multi-round attacks.

We call our quantum money scheme \mathcal{Q} . In order to verify a \mathcal{Q} -coin a holder has to contact the bank via a classical communication channel and perform quantum measurements, as directed by the bank, then report the outcomes. In the end the bank either confirms that the coin is valid or rejects it.

Our construction has the following specific properties:

- The coins are *exponentially* hard to counterfeit (cf. Theorem 5.1 and Corollary 5.2).
- The classical communication channel used for verification can be *unencrypted*: e.g., both the bank and the coin holder can broadcast their messages, without compromising security of the scheme.
- Our scheme remains secure against an adversary who uses *adaptive “attempted verifications”* in order to collect information about a coin. Exponentially many such attempts have to be made before one has non-negligible chances to counterfeit a coin.
- The database of the bank is *static*, and therefore many de-centralized “verification branches” can exist that do not have to communicate with one another.
- The number of verifications that a \mathcal{Q} -coin can go through is limited – the number of qubits required to store a coin is polynomial in the number of validity tests via classical communication that the coin can go through during its circulation period (after that it would have to be replaced by the bank). We show that this dependency is *optimal* (cf. Theorem 6.1).

1.3 Related work

Using a different approach, Aaronson and Christiano in the upcoming [AC12] will construct a scheme that uses quantum communication with a bank for verification (like Wiesner’s original scheme) and is resistant against adaptive multi-round attacks.

Very recently some of the ideas proposed in this work have been further developed by Pastawski et al. [PYJ⁺11] and by Molina et al. [MVW12].

2 Who needs quantum money?

The first quantum money scheme was proposed by Wiesner more than 30 years ago (several years before [Wie83] was published). Nevertheless, there seems to remain some confusion about the advantages that quantum money has over possible classical constructions. Below we reproduce a typical “classical money” proposal, then discuss the advantages of Wiesner’s scheme, then further advantages of our construction.

Note that here we are not comparing our scheme to the previously known ones (that was the subject of Section 1.2). Instead, this part (informally) addresses the question posed by its title.

2.1 A classical proposal

Let every coin issued by the bank contain a secret string s , known only to the bank and to the current coin holder. When a coin holder Alice wants to pass her coin to a new coin holder Bob, they run the following protocol:

- Alice sends to the bank the string s and tells the bank that she wants to pass the coin to Bob.
- The bank checks that s is a valid secret string (if not then a forgery attempt has been detected), then erases s from the list of valid strings and adds to the list a newly generated secret string s' .
- The bank sends s' to Bob; henceforth, Bob holds the coin.

2.2 Advantages of Wiesner's scheme

- The bank's database can be static (for the classical scheme to be secure, it is crucial that a new secret string is issued each time a coin is passed along).
- Interaction with the bank does not require 3-party authentication (for the classical scheme to be secure, the bank has to make sure that the *only* recipient of the newly generated secret string is the party named by Alice in the first round).

2.3 Advantages of our scheme

- All the benefits of Wiesner's construction listed above.
- The communication channel can be classical and not encrypted. Moreover, all the messages (both ways) can be openly broadcast.
- In the classical scheme, as well as in Wiesner's scheme, an intruder who pretends to be the bank can steal a valid coin from its fair holder who wants to verify it. Our scheme shields against that.

3 Notation and preliminaries

For $a \in \mathbb{N}$ we denote $[a] \stackrel{\text{def}}{=} \{1, \dots, a\}$. Denote by I_a the identity matrix of rank a . For any finite A we denote by \mathcal{U}_A the uniform distribution over the elements of A .

We will use concentration bounds extensively in our proofs.

Theorem 3.1. (Chernoff bound) *Let X_1, \dots, X_n be mutually independent random variables taking values in $[0, 1]$, such that $\mathbf{E}[X_i] = \mu$ for all $i \in [n]$. Then for any $\lambda > 0$,*

$$\Pr \left[\sum_{i \in [n]} X_i \geq (1 + \lambda)\mu n \right] \leq e^{\frac{-n\lambda^2\mu}{2+\lambda}},$$

and

$$\Pr \left[\sum_{i \in [n]} X_i \leq (1 - \lambda)\mu n \right] \leq e^{\frac{-n\lambda^2\mu}{2}}.$$

We also need a generalization, originally proved by Panconesi and Srinivasan [PS97]. The following version of it is due to Impagliazzo and Kabanets [IK10].

Theorem 3.2. (Generalized Chernoff bound) *Let X_1, \dots, X_n be Boolean random variables, such that for some δ and every $S \subseteq [n]$ it holds that $\Pr [\wedge_{i \in S} X_i = 1] \leq \delta^{|S|}$. Then*

$$\Pr \left[\sum_{i \in [n]} X_i \geq (1 + \lambda)\delta n \right] \leq e^{-2n\lambda^2\delta^2}.$$

Corollary 3.3. *Let X_1, \dots, X_n be Boolean random variables, such that for all $i \in [n]$ and any event C that only depends on $\{X_j | j \neq i\}$ it holds that $\Pr [X_i = 1 | C] \leq \delta$. Then*

$$\Pr \left[\sum_{i \in [n]} X_i \geq (1 + \lambda)\delta n \right] \leq e^{-2n\lambda^2\delta^2}.$$

Proof. For every $S \subseteq [n]$,

$$\Pr [\wedge_{i \in [|S|]} X_{S_i} = 1] = \prod_{i \in [|S|]} \Pr [X_{S_i} = 1 | X_{S_1} = \dots = X_{S_{i-1}} = 1] \leq \delta^{|S|},$$

where S_i is the i 'th least element of S . ■

We will also need the following combinatorial lemma (a rather standard one, e.g., see Lemma 2.2 in Jukna [Juk01]).

Lemma 3.4. *Let A_1, \dots, A_N be subsets of $[n]$ of average size t . Suppose that $|A_i \cap A_j| \leq s$ for every $i \neq j$. Then either $N < 2n/t$ or $s > t^2/2n$ (or both).⁴*

Proof. For $x \in [n]$, let $d(x) \stackrel{\text{def}}{=} |\{i | x \in A_i\}|$. On the one hand,

$$\sum_{x=1}^n d(x) = \sum_{i=1}^N |A_i| = Nt \quad \implies \quad \sum_{x=1}^n d(x)^2 \geq \frac{N^2 t^2}{n}.$$

On the other hand,

$$\sum_{x=1}^n (d(x))^2 = \sum_{i=1}^N \sum_{j=1}^N |A_i \cap A_j| = \sum_{i=1}^N |A_i| + \sum_{j \neq i} |A_i \cap A_j| \leq Nt + N(N-1)s.$$

Therefore,

$$\frac{Nt^2}{n} < t + Ns \quad \implies \quad s > \frac{t^2}{n} - \frac{t}{N},$$

and the result follows. ■

⁴The asymptotic guarantees of our Lemma 3.4 are slightly better than those of Lemma 2.2 in [Juk01] – there the main statement is more general, but the result is weaker in the special case that we are interested in.

4 Our quantum money scheme \mathcal{Q}

One of the main technical ingredients of our construction is a constant-dimensional ($n = 4$) special case of a relational communication problem called *Hidden Matching Problem (HMP)*, first considered by Bar-Yossef, Jayram and Kerenidis [BYJK04] in the context of communication complexity.

Definition 1. Let HMP_4 be as follows. For $x \in \{0, 1\}^4$ and $m, a, b \in \{0, 1\}$, we say that $(x, m, a, b) \in HMP_4$ if $b = \begin{cases} x_1 \oplus x_{2+m} & \text{if } a = 0 \\ x_{3-m} \oplus x_4 & \text{if } a = 1 \end{cases}$.

Intuitively, if we view $x \in \{0, 1\}^4$ as a binary coloring of 4 vertices then a tuple (x, m, a, b) satisfies the relation HMP_4 if and only if b indicates whether x assigns distinct colors to the pair of vertices determined by m and a .

It has been shown in [BYJK04] that if Alice receives x and Bob receives m then Alice can send a short *quantum* message to Bob that would allow him to produce a valid answer (a, b) ; on the other hand, if Alice is only allowed to send classical bits then a much longer message is required. The authors were interested in the asymptotic behavior of quantum and classical communication cost of HMP , and they gave an elegant proof that the gap between the two is exponential.

How can it help us? We want to build a scheme that would be safe against both classical and quantum attacks; moreover, we want to be able to carry out certain communication task (testing validity of a coin) using *only classical* communication. So, why are we interested in something showing that quantum communication is more powerful than classical?

The answer is that the role of quantum communication from [BYJK04] in our case is played by a *quantum coin*: when the bank issues a coin, it sends a quantum message to its future holder. The core of our construction is the observation (apparently, new to this work) that in certain quantum one-way protocol for HMP , a single message from Alice cannot be used by Bob in order to produce valid answers w.r.t. several different values of m . In other words, *the message cannot be “reused”*. This holds in spite of the fact that a message from Alice cannot depend on m , thus using it Bob can produce a valid answer w.r.t. any legitimate value of m .

In our construction we will use a state $|\alpha(x)\rangle$ of 2 qubits (corresponding to the quantum message that Alice would send to Bob in a one-way protocol for HMP_4) that allows its holder, who is given m but doesn't know x , to find an “answer” (a, b) that satisfies HMP_4 with certainty. On the other hand, using the same state in order to find (a_0, b_0) and (a_1, b_1) , such that $(x, m, a_m, b_m) \in HMP_4$ for both $m = 0$ and $m = 1$ would fail with probability at least $1/4$. In other words, our state of 2 qubits will be *useful but not reusable* for producing an answer to HMP_4 .

Let the bank choose $x_1, \dots, x_k \in \{0, 1\}^4$ at random, keep them in secret and produce quantum states $|\alpha(x_1)\rangle, \dots, |\alpha(x_k)\rangle$. A newly issued \mathcal{Q} -coin consists of a piece of paper glued to k quantum registers that hold $|\alpha(x_1)\rangle, \dots, |\alpha(x_k)\rangle$. The piece of paper contains a unique identification tag and k initially unmarked positions, where the i 'th position has to be marked when the corresponding $|\alpha(x_i)\rangle$ is used in the verification protocol.

More formally:

Definition 2. (HMP_4 -states) Let $x \in \{0, 1\}^4$. The corresponding HMP_4 -state is

$$|\alpha(x)\rangle \stackrel{\text{def}}{=} \frac{1}{2} \cdot \sum_{1 \leq i \leq 4} (-1)^{x_i} |i\rangle.$$

Interestingly, the HMP_4 -states (in their multidimensional version) were first considered by Kerenidis and de Wolf [KdW04] in order to prove a lower bound on the length of certain codes, and that was before the Hidden Matching Problem was defined.

Definition 3. (HMP_4 -queries) An HMP_4 -query is an element $m \in \{0, 1\}$. A valid answer to the query w.r.t. $x \in \{0, 1\}^4$ is a pair $(a, b) \in \{0, 1\} \times \{0, 1\}$, such that $(x, m, a, b) \in HMP_4$.

An HMP_4 -state can be used to answer an HMP_4 -query with certainty: If $m = 0$, let

$$v_1 \stackrel{\text{def}}{=} \frac{|1\rangle + |2\rangle}{\sqrt{2}}, v_2 \stackrel{\text{def}}{=} \frac{|1\rangle - |2\rangle}{\sqrt{2}}, v_3 \stackrel{\text{def}}{=} \frac{|3\rangle + |4\rangle}{\sqrt{2}}, v_4 \stackrel{\text{def}}{=} \frac{|3\rangle - |4\rangle}{\sqrt{2}};$$

otherwise ($m = 1$), let

$$v_1 \stackrel{\text{def}}{=} \frac{|1\rangle + |3\rangle}{\sqrt{2}}, v_2 \stackrel{\text{def}}{=} \frac{|1\rangle - |3\rangle}{\sqrt{2}}, v_3 \stackrel{\text{def}}{=} \frac{|2\rangle + |4\rangle}{\sqrt{2}}, v_4 \stackrel{\text{def}}{=} \frac{|2\rangle - |4\rangle}{\sqrt{2}}.$$

Measure $|\alpha(x)\rangle$ in the basis $\{v_1, v_2, v_3, v_4\}$, and let (a, b) be $(0, 0)$ if the outcome is v_1 ; $(0, 1)$ in the case of v_2 ; $(1, 0)$ in the case of v_3 ; $(1, 1)$ in the case of v_4 . Then $(x, m, a, b) \in HMP_4$ always.

Definition 4. (\mathcal{Q} -coins) Let $3|t$. A secret record consists of k entries x_1, \dots, x_k , $x_i \in \{0, 1\}^4$ (i.e., the secret record contains $4k$ bits).

A “fresh” \mathcal{Q} -coin corresponding to the record (x_1, \dots, x_k) consists of

- k quantum registers consisting of 2 qubits each, where the i 'th register contains $|\alpha(x_i)\rangle$;
- a k -bit classical register P , that is initially set to 0^k ;
- a unique identification number.

A bank produces fresh \mathcal{Q} -coins; as a \mathcal{Q} -coin goes through more and more verification protocols, its quantum registers lose their original content (and that shall be reflected in the corresponding bits of P , see below). The identification number of every coin issued by the bank must be unique.

To verify a \mathcal{Q} -coin through classical communication with the bank, its holder runs the following protocol Ver (t is a parameter in the construction of \mathcal{Q} that will be polynomially related to k).

Protocol: When a holder of a valid \mathcal{Q} -coin follows the protocol, verification goes like this:

1. The holder sends the identification number of the \mathcal{Q} -coin to the bank.
2. The bank chooses uniformly at random a set $L_{bn} \subset [k]$ of size t , and sends it to the coin holder.
3. The holder consults with P and chooses uniformly at random a set $L_{hl} \subset L_{bn}$ consisting of $2t/3$ yet unmarked positions. He sends L_{hl} to the bank and marks in P all the elements of L_{hl} as used.
4. The bank chooses at random $2t/3$ values $m_i \in \{0, 1\}$, one for each $i \in L_{hl}$, and sends them to the coin holder.
5. The holder measures the quantum registers corresponding to the elements of L_{hl} in order to produce $2t/3$ pairs (a_i, b_i) , such that $(x_i, m_i, a_i, b_i) \in HMP_4$ for all $i \in L_{hl}$. He sends the list of (a_i, b_i) 's to the bank.
6. The bank checks whether $(x_i, m_i, a_i, b_i) \in HMP_4$ for all $i \in L_{hl}$, in which case it confirms validity of the \mathcal{Q} -coin. Otherwise, the coin is declared to be a counterfeit.

We will say that an instance of \mathcal{V}_{er} has been *passed* or *won* if the bank’s final response was “valid”.

Observe that a fair coin holder fails to pass \mathcal{V}_{er} with exponentially small probability (corresponding to the situation when less than $k/4$ of the coin registers are marked as used, but among the t registers that were uniformly chosen by the bank more than $t/3$ are marked as used). If this happens, a new run of \mathcal{V}_{er} can be started.

It follows from the earlier discussion that both the bank and a fair coin holder can perform their parts of \mathcal{V}_{er} efficiently. Note also that the secret records kept by the bank do not change as a result of executing \mathcal{V}_{er} – that is, the bank’s database is *static*.

Intuitively, adversarial ability to counterfeit a \mathcal{Q} -coin shall imply ability to answer w.r.t. the same quantum register i both to the question $m_i = 0$ and to $m_i = 1$. As we said before, that can be done with probability at most $3/4$; moreover, it turns out that in order to successfully counterfeit a coin the adversary must be able to answer both the HMP_4 -queries w.r.t. a considerable fraction of the coin’s registers, and that will imply exponentially small probability of adversarial success. We will formalize and prove this intuition in Section 5.

We will show (cf. Theorem 5.1 and Corollary 5.2) that only after an adversary has run $e^{\Omega(t^3/k^2)}$ auxiliary instances of \mathcal{V}_{er} , he might be able to counterfeit a \mathcal{Q} -coin with success probability higher than $e^{-\Omega(t^2/k)}$.

Note that every run of \mathcal{V}_{er} “costs” $2t/3$ yet unused quantum registers. As soon as $k/4$ registers have been used, the \mathcal{Q} -coin has to be returned to bank (the bank still would be able to verify its validity and issue a replacement). Accordingly, after $\lceil 3k/st \rceil$ runs of \mathcal{V}_{er} a \mathcal{Q} -coin has to be returned to the bank.

To conclude: Choosing, for example, $t \in \Theta(k^{3/4})$ gives a construction where a coin that consists of $2k$ qubits can go through $\Omega(k^{1/4})$ validity tests via classical communication with the bank, and where it takes $e^{k^{\Omega(1)}}$ time to forge a counterfeit with probability higher than $e^{-k^{\Omega(1)}}$. The bank’s secret database contains $4k$ bits corresponding to every coin, and those records are static (in particular, many de-centralized “verification branches” can exist that do not have to communicate with one another). In Section 6 we will show that these parameters are very close to the best possible.

5 Security of \mathcal{Q}

We are giving “extended security guarantees”, as follows. Instead of only arguing that the first cheating attempt is not likely to succeed, we allow an adversary to use *multiple attacking attempts* – namely, even having been caught cheating in the past, he may continue his attempts. Recall that we allow adaptive attacks, thus something learnt from the earlier attempts might help the adversary in future attacks.

Informally, our security guarantees will be expressed like this: *An exponentially large number of partially completed instances⁵ of \mathcal{V}_{er} are required for an adversary to have non-negligible probability to make a counterfeit coin.*

A high-level view of our security analysis is as follows. First we make preliminary observations regarding possible attacks on the \mathcal{Q} -scheme (Section 5.1), and demonstrate useful properties of HMP_4 -states (Section 5.2). Then we claim that counterfeiting a \mathcal{Q} -coin has its “cost”, in terms of the number of preliminary runs of \mathcal{V}_{er} that are required in order to collect enough auxiliary

⁵By a partially completed protocol we mean an instance of \mathcal{V}_{er} , where the first response from the bank has been received and analyzed by the adversary.

information about the coin (Section 5.3). Finally, we reduce unrestricted attacks to more structured ones and show their limitations (Sections 5.4 and 5.5, respectively). We conclude in Section 5.6 that exponentially many preliminary runs of \mathcal{V}_{er} are required in order to counterfeit a \mathcal{Q} -coin.

5.1 Possible attacks and security guarantees

Our goal will be to show that a \mathcal{Q} -coin is hard to counterfeit. First, we want to argue that in order to establish security of our \mathcal{Q} -scheme it is enough to consider the situation when starting with a fresh authentic \mathcal{Q} -coin, an adversary runs many instances of \mathcal{V}_{er} (probably, in a non-consecutive manner) for this coin⁶, and his goal is to produce two (possibly, entangled) quantum objects that have non-negligible probability to be accepted by the bank as valid coins.

Probably, the most harmful attack on \mathcal{Q} would be the one where an adversary starts with M fresh \mathcal{Q} -coins, and his goal is to produce $M + 1$ quantum objects that are all likely to be accepted as valid \mathcal{Q} -coins.⁷ Let us look at the “two out of one” security guarantee that we give for the \mathcal{Q} -scheme, and see how it implies robustness against “multi-coin” attacks.

Let us call *the first response* a message that the bank sends in step 2 of \mathcal{V}_{er} (that is, a list of t positions). In Section 5.6 we establish the following theorem.

Theorem 5.1. *Let a fresh \mathcal{Q} -coin be given to a computationally unlimited adversary who runs auxiliary instances of \mathcal{V}_{er} for this coin and produces two (possibly, entangled) “counterfeits” ρ_1 and ρ_2 . Then*

$$U \in e^{\Omega(t^3/k^2)}$$

exists, such that if the adversary has received and analyzed the first bank’s responses in at most U instances of \mathcal{V}_{er} , then the probability that both ρ_1 and ρ_2 pass \mathcal{V}_{er} is in

$$e^{-\Omega(t^2/k)}.$$

Corollary 5.2. *Let M fresh \mathcal{Q} -coins be given to a computationally unlimited adversary who analyzes the first bank’s responses in at most U auxiliary instances of \mathcal{V}_{er} , for U as in Theorem 5.1. If the adversary outputs $M + 1$ quantum objects then the probability that all of them pass \mathcal{V}_{er} is in $e^{\ln M - \Omega(t^2/k)}$.*

Proof. If the identification numbers of the $M + 1$ produced quantum objects are not a subset of the identification numbers of M initially given objects then at least one counterfeit has been produced “from scratch”, and it is easy to see that the probability of success in this case is negligible.

Otherwise there is at least one identification number that appears more than once among the $M + 1$ produced quantum objects with probability at least $1/M$. Starting with a single coin, one can emulate the cheating strategy for M coins by locally creating $M - 1$ \mathcal{Q} -coins and running the protocol, locally computing bank’s responses according to \mathcal{V}_{er} w.r.t. any of the $M - 1$ auxiliary coins. If in the end of emulation at least two objects are marked with the same identification number as the given coin then those two objects are returned, otherwise arbitrary output is produced.

If the M -coin counterfeit strategy produces $M + 1$ quantum objects that successfully pass verification with probability ε , then the strategy above succeeds in counterfeiting a single coin with probability at least ε/M , and the corollary follows from Theorem 5.1. ■

⁶Note that every run of \mathcal{V}_{er} can be associated with certain \mathcal{Q} -coin via its identification number, as reported by the coin holder in the first round of a protocol.

⁷Several modifications of this cheating setup can be considered, but it seems that all of them can be reduced to the “ $M + 1$ out of M ” regime.

5.2 Quantum retrieval games

To analyze some useful properties of HMP_λ -states we define the notions of *quantum retrieval games*, *physical projections*, and *selective projections*.

Definition 5. (quantum retrieval games) *Let $k, m, n \in \mathbb{N}$, $\sigma \subseteq [m] \times [n]$, and $\forall a \in [n]$ let $\rho_a \in \mathbb{C}^{k \times k}$ be positive semidefinite such that $\text{tr}(\sum_a \rho_a) = 1$. Then $\mathcal{G} = ((\rho_a)_{a \in [n]}, \sigma)$ is a quantum retrieval game.*

The notion of quantum retrieval games is aimed to model the situation when a mixed quantum state $\sum \rho_a$ is measured in order to extract some information about a . The relation σ describes what knowledge is wanted. We will consider situations when an m -outcome quantum measurement is applied to $\sum \rho_a$, and we say that the game \mathcal{G} has been won if the pair (*outcome of the measurement*, a) is in σ . Formally:

Definition 6. (selective and physical projections) *Let $\mathcal{P} = \{P_i\}_{i=1}^m$ be a set of projections in $\mathbb{C}^{k \times k}$, s.t. $\sum_i P_i \preceq I$. Call \mathcal{P} a selective projection. A selective projection is called physical projections if it satisfies $\sum_i P_i = I$.*

Definition 7. (selective and physical values of a game) *The value of \mathcal{G} w.r.t. \mathcal{P} is defined as*

$$\frac{\sum_{(i,a) \in \sigma} \text{tr}(P_i \rho_a)}{\sum_{i,a} \text{tr}(P_i \rho_a)},$$

and if \mathcal{P} is a selective projection then the value is undefined unless $\sum_{i,a} \text{tr}(P_i \rho_a) > 0$. The selective value of \mathcal{G} is the supremum of the game's value w.r.t. selective projections, and the physical value of \mathcal{G} is the supremum of the game's value w.r.t. physical projections.

Note that for physical projections it holds that $\sum_{i,a} \text{tr}(P_i \rho_a) = 1$ (and the above definition simplifies to $\sum_{(i,a) \in \sigma} \text{tr}(P_i \rho_a)$ in that case). Physical projections are the most general “mechanism” offered by quantum mechanics to extract classical information from a quantum state.

Selective projections are, in general, more powerful than physical projections (they correspond to measurements with “postselection”, and those are not allowed by the laws of quantum mechanics). We will consider selective projections in some of our impossibility statements, that will allow simpler proofs of direct product statements that we will need. Like in the case of physical projections, we will view the elements of \mathcal{P} as outcomes. Clearly, the selective value of a game is always at least as large as its physical value.

A physical projection \mathcal{P} corresponds to some POVM measurement, and the elements of \mathcal{P} are the possible outcomes. When it is applied to some (normalized) $\rho \in \mathbb{C}^{k \times k}$, the i 'th outcome occurs with probability $\text{tr}(P_i \rho)$. If i 'th outcome occurred then the state of the quantum register that originally contained $\rho \in \mathbb{C}^{k \times k}$ becomes $M_i \rho M_i^\dagger$, where $M_i = \sqrt{P_i}$. We view selective projections as a generalization of POVMs where the requirement $\sum_i P_i = I$ is replaced by $\sum_i P_i \preceq I$ and the distribution of outcomes is

$$\mathbf{Pr} [i\text{'th outcome}] \stackrel{\text{def}}{=} \frac{\text{tr}(P_i \rho)}{\sum_j \text{tr}(P_j \rho)}.$$

The class of selective projections is closed w.r.t. compositions and applying admissible quantum transformations.⁸

⁸The class of admissible quantum transformations generalizes the class of unitary transformations to include what can be achieved using auxiliary space.

5.2.1 An HMP_4 -state cannot be used twice

We have seen that each HMP_4 -state can be used to answer at least one HMP_4 -query. To prove that our \mathcal{Q} is secure we will have to argue that an HMP_4 -state cannot be used to answer two complementary HMP_4 -queries with confidence.

Let \mathcal{G}_{HMP_4} be the quantum retrieval game corresponding to answering both the possible HMP_4 -queries using one HMP_4 -state, namely

$$\mathcal{G}_{HMP_4} \stackrel{\text{def}}{=} \left((1/16 \cdot |\alpha(x)\rangle\langle\alpha(x)|)_{x \in \{0,1\}^4}, \sigma_{HM} \right),$$

where

$$\sigma_{HM} \stackrel{\text{def}}{=} \{(x, (a_0, b_0, a_1, b_1)) \mid (x, 0, a_0, b_0), (x, 1, a_1, b_1) \in HMP_4\}.$$

Note that this definition corresponds to the uniform choice of $x \in \{0,1\}^4$.

Lemma 5.3. *The selective value of \mathcal{G}_{HMP_4} is at most $3/4$.⁹*

Proof. Note that $\sum_x 1/16 \cdot |\alpha(x)\rangle\langle\alpha(x)| = 1/4 \cdot I$. Consider a selective projection that produces correct answer to \mathcal{G}_{HMP_4} with probability δ . There must exist an answer (a'_0, b'_0, a'_1, b'_1) that is produced with non-zero probability, and if it is produced then it is correct with probability at least δ when $x \sim \mathcal{U}_{\{0,1\}^4}$. Fix one such answer.

Denote by E the event that $(x, (a'_0, b'_0, a'_1, b'_1)) \in \sigma_{HM}$. By the definition of selective value there exists a projection P such that $\text{tr}(P\rho) > 0$ and if the outcome P occurs then E holds with probability at least δ . We will argue that E cannot be “witnessed” very well by any outcome of measuring ρ .

Observe that E always corresponds to choosing three different coordinates $j_1, j_2, j_3 \in [4]$ and fixing the values of $x_{j_1} \oplus x_{j_2}$ and $x_{j_2} \oplus x_{j_3}$. By symmetry of $|\alpha(x)\rangle$, we can, w.l.g., consider the case of witnessing $x_1 = x_2 = x_3$ via measuring of ρ .

Let $P = \sum_{i \in [k]} |e_i\rangle\langle e_i|$ for some orthonormal $|e_1\rangle, \dots, |e_k\rangle$. We have:

$$\begin{aligned} \delta &\leq \frac{\text{tr}\left(P \sum_{x_1=x_2=x_3} |\alpha(x)\rangle\langle\alpha(x)|\right)}{\text{tr}\left(P \sum_{x \in \{0,1\}^4} |\alpha(x)\rangle\langle\alpha(x)|\right)} \leq \max_{i \in [k]} \left\{ \frac{\sum_{x_1=x_2=x_3} |\langle e_i | \alpha(x) \rangle|^2}{\text{tr}\left(|e_i\rangle\langle e_i| \sum_{x \in \{0,1\}^4} |\alpha(x)\rangle\langle\alpha(x)|\right)} \right\} \\ &= \frac{1}{4} \cdot \sum_{x_1=x_2=x_3} |\langle e_0 | \alpha(x) \rangle|^2, \end{aligned}$$

where $|e_0\rangle \in \{|e_1\rangle, \dots, |e_k\rangle\}$ attains the optimum of the second inequality (under the “0/0 = 0” convention). Let $e_0^{(j)} \stackrel{\text{def}}{=} \langle e_0 | j \rangle$ for $j \in [4]$, then

$$\begin{aligned} \delta &\leq \frac{1}{8} \left(\left(e_0^{(1)} + e_0^{(2)} + e_0^{(3)} + e_0^{(4)} \right)^2 + \left(e_0^{(1)} + e_0^{(2)} + e_0^{(3)} - e_0^{(4)} \right)^2 \right) \\ &= \frac{1}{4} \left(\left(e_0^{(1)} + e_0^{(2)} + e_0^{(3)} \right)^2 + \left(e_0^{(4)} \right)^2 \right) \leq \frac{3}{4}, \end{aligned}$$

because $|e_0\rangle$ is a unit vector. ■

For $k \in \mathbb{N}$, let $\mathcal{G}_{HMP_4}^k$ be the naturally defined “product game” that consists of k independent instances of \mathcal{G}_{HMP_4} .

⁹From the proof it can be seen that the bound is, actually, tight.

Corollary 5.4. *The selective value of $\mathcal{G}_{HMP_4}^k$ is at most $(3/4)^k$.*

Proof. Let $\mathcal{G}_{HMP_4}^{(i)}$ denote the i 'th instances of \mathcal{G}_{HMP_4} “inside” $\mathcal{G}_{HMP_4}^k$. Then

$$\Pr \left[\mathcal{G}_{HMP_4}^k \text{ is won} \right] = \prod_{i \in [k]} \Pr \left[\mathcal{G}_{HMP_4}^{(i)} \text{ is won} \mid \mathcal{G}_{HMP_4}^{(1)}, \dots, \mathcal{G}_{HMP_4}^{(i-1)} \text{ are won} \right], \quad (1)$$

where the probabilities are defined w.r.t. the selective measurement being used to play $\mathcal{G}_{HMP_4}^k$.

Note that each conditional probability that appears on the right-hand side of (1) is at most $3/4$: Otherwise there would exist a selective measurement that used $i - 1$ auxiliary instances $\mathcal{G}_{HMP_4}^{(1)}, \dots, \mathcal{G}_{HMP_4}^{(i-1)}$, and conditioned upon winning these $i - 1$ instances won $\mathcal{G}_{HMP_4}^{(i)}$ with probability higher than $3/4$, contradicting Lemma 5.3. The result follows. \blacksquare

5.3 The cost of counterfeiting a \mathcal{Q} -coin

Unless stated otherwise, let \mathbf{c} be the \mathcal{Q} -coin that an adversary is trying to counterfeit, and let \mathbf{x} be the bank's secret string that describes the structure of \mathbf{c} . We want to argue that in order to achieve his goal, the adversary has to collect certain minimal amount of additional information about the coin, and that task itself is difficult to fulfill.

Let us assume for the rest of our security analysis that the attack under consideration, denoted by \mathcal{C} , runs at most U instances of \mathcal{V}_{er} , all of them initiated by sending the identification number of \mathbf{c} . Informally, \mathcal{C} is successful if in the end it outputs quantum states ρ_1 and ρ_2 (possibly entangled), such that both of them, if given to a trustworthy user, pass \mathcal{V}_{er} with some non-negligible probability. This probability is viewed w.r.t. the randomness present in the construction of \mathbf{c} , in \mathcal{C} itself, and in the final run of \mathcal{V}_{er} .

It is crucial that we consider the probability of both the fakes having been accepted. If instead we were asking what is the smaller of the probabilities that ρ_j passes \mathcal{V}_{er} for $j \in \{0, 1\}$, we would end up with a bound of at least $1/2$: For example, an adversary can toss $j \sim \mathcal{U}_{\{0,1\}}$ and make ρ_j to be \mathbf{c} , and ρ_{1-j} to be anything.

Lemma 5.5. *Consider an attack that completes at most U instances of \mathcal{V}_{er} in order to counterfeit \mathbf{c} . Conditional upon having passed at most $u \leq U$ instances, the success probability of counterfeiting is at most*

$$e^{-\Omega(t)} + e^{u \ln U - \Omega(k)}.$$

Proof. For $j \in \{0, 1\}$, let I_j be a random variable taking the value of the list of HMP_4 -registers that are marked as unused on ρ_j . By the definition of \mathcal{Q} -scheme it should hold that $|I_j| \geq 3k/4$ (otherwise the forgery would be obvious right away).

Consider the run of \mathcal{V}_{er} for the counterfeit contained in ρ_j . For any choice of L_{hl} in step 3 and of “questions” $\{m_i \mid i \in L_{hl}\}$ in step 4, the quantum measurement applied by the coin holder in step 5 can be decomposed into $2^{t/3}$ measurements that access individual registers of ρ_j in order to find answers w.r.t. the corresponding m_i . Let us denote by $P_j^{i,m}$ the measurement applied to ρ_j in order to produce (a_i, b_i) when $m_i = m$.

In step 5 of \mathcal{V}_{er} the holder of ρ_j performs the measurements $\{P_j^{i,m_i} \mid i \in L_{hl}\}$ in order to determine the $2^{t/3}$ pairs (a_i, b_i) that he will report to the bank. Now we make two observations that will be crucial for the proof:

- The only pairs of the measurements that do not commute are

$$\left\{ \left(P_j^{i,m}, P_j^{i,1-m} \right) \mid j \in \{0, 1\}, i \in [k], m \in \{0, 1\} \right\}.$$

- Since the coin holder is now fair to the protocol, the $2^{t/3}$ -set L_{fl} chosen in step 3 of \mathcal{V}_{er} is a uniformly random subset of I_j . The questions $(m_i)_{i \in L_{\text{fl}}}$ are i.i.d. by $\mathcal{U}_{\{0,1\}}$.

Denote by V_j the instance of \mathcal{V}_{er} that tests ρ_j , and accordingly define L_{fl}^j and m_i^j . Let us view choosing $(m_i^j)_{i \in L_{\text{fl}}^j}$ as first taking $\mathbf{m}^j \sim \mathcal{U}_{\{0,1\}^k}$, followed by choosing a random L_{fl}^j and outputting the projection of \mathbf{m}^j to the coordinates in L_{fl}^j . Clearly, the resulting distribution of L_{fl}^j and $(m_i^j)_{i \in L_{\text{fl}}^j}$ are the same. Therefore, we can replace the protocols V_1 and V_2 by a new quantum procedure, somewhat more friendly to analyze.

Let \tilde{V} be the following procedure that either accepts or rejects quantum states ρ_1 and ρ_2 .

1. For $j \in \{0, 1\}$, choose $\mathbf{m}^j \sim \mathcal{U}_{\{0,1\}^k}$.
2. For $j \in \{0, 1\}$ and $i \in I_j$, apply P_j^{i, \mathbf{m}^j} to ρ_j and denote the outcome by (a_i^j, b_i^j) .
3. For $j \in \{0, 1\}$, choose L_{fl}^j as a uniformly random subset of I_j of size $2^{t/3}$.
4. Accept if for all $j \in \{0, 1\}$ and $i \in L_{\text{fl}}^j$ it holds that $(\mathbf{x}_i, \mathbf{m}_i^j, a_i^j, b_i^j) \in \text{HMP}_4$; reject otherwise.

Observe that all P_j^{i, \mathbf{m}^j} 's that can appear in a single run of \tilde{V} commute, and therefore the probability that \tilde{V} accepts exactly equals the probability that both V_1 accepts ρ_1 and V_2 accepts ρ_2 .

Denote $I \stackrel{\text{def}}{=} I_1 \cap I_2$, $I' \stackrel{\text{def}}{=} \{i \in I \mid \mathbf{m}_i^1 \neq \mathbf{m}_i^2\}$, $\tilde{I}_j \stackrel{\text{def}}{=} \{i \in I' \mid (\mathbf{x}_i, \mathbf{m}_i^j, a_i^j, b_i^j) \notin \text{HMP}_4\}$ and $\tilde{I} \stackrel{\text{def}}{=} \tilde{I}_1 \cup \tilde{I}_2$. We will see that \tilde{I} is unlikely to be small, and if it is big then \tilde{V} is unlikely to accept.

Let us first consider the case when the adversary has not run any preliminary protocol and created ρ_1 and ρ_2 from \mathbf{c} alone, without any auxiliary knowledge about \mathbf{x} .

By definition, $|I| \geq k/2$. By uniformity of \mathbf{m}^1 and \mathbf{m}^2 it holds that $\mathbf{E}[|I'|] = |I|/2$, and Chernoff bound (Theorem 3.1) implies

$$\Pr \left[|I'| \leq \frac{k}{5} \right] < e^{-\frac{|I|}{100}} \leq e^{-\frac{k}{200}}. \quad (2)$$

By Lemma 5.3, for every $i_0 \in I'$ it holds that $\Pr [i_0 \notin \tilde{I}] \leq 3/4$; moreover, the same remains true even if we condition upon the content of $\tilde{I} \setminus \{i_0\}$ (otherwise Lemma 5.3 would be contradicted by a selective measurement that uses auxiliary instances of $\mathcal{G}_{\text{HMP}_4}$ in order to win with higher probability, similarly to the proof of Corollary 5.4). Therefore Corollary 3.3 can be used here, resulting in

$$\Pr \left[|\tilde{I}| \leq \frac{|I'|}{5} \right] \leq e^{-\frac{|I'|}{200}}. \quad (3)$$

Clearly,

$$\Pr \left[|\tilde{I}| < \frac{k}{25} \right] \leq \Pr \left[|I'| \leq \frac{k}{5} \right] + \Pr \left[|\tilde{I}| \leq \frac{|I'|}{5} \mid |I'| > \frac{k}{5} \right],$$

which leads, together with (2) and (3), to

$$\Pr \left[|\tilde{I}| < \frac{k}{25} \mid \star \right] \leq e^{-\frac{k}{200}} + e^{-\frac{k}{1000}}, \quad (4)$$

where “ \star ” is the condition that ρ_1 and ρ_2 are created from \mathbf{c} not using any auxiliary input.

Now assume that in order to produce ρ_1 and ρ_2 the adversary has completed at most U instances of \mathcal{V}_{er} , and condition upon at most u of them having been passed successfully. The idea here is to emulate the same attack, letting the adversary guess the bank’s responses locally. In this form the attack uses no auxiliary data from the bank, which makes \star from (4) hold.

According to \mathcal{V}_{er} , the only bank’s message that depends on \mathbf{x} is the final “accept”/“reject” notice. Therefore, if the adversary (who doesn’t know \mathbf{x}) does his best to predict all bank’s responses, such predictions will be statistically indistinguishable from bank’s responses, as long as all “accept”/“reject” verdicts are guessed correctly. The number of different ways to choose at most u “accepts” out of U verdicts is at most $U^u + 1$, and therefore they are guessed correctly with probability at least $\frac{1}{U^u + 1}$. Thus from (4),

$$\frac{\Pr \left[\left| \tilde{I} \right| < k/5 \right]}{U^u + 1} \leq e^{-\frac{k}{200}} + e^{-\frac{k}{1000}}. \quad (5)$$

Now assume that $\left| \tilde{I} \right| \geq k/5$. W.l.g., let $\left| \tilde{I}_1 \right| \geq k/10$. Then the probability that \tilde{V} accepts is upper-bounded by the probability that none of the elements of $L_{\tilde{I}_1}^1$ comes from \tilde{I}_1 , and that is at most $(9/10)^{2t/3} < (14/15)^t$. Together with (5), this implies that

$$\Pr \left[\tilde{V} \text{ accepts} \right] < \left(\frac{14}{15} \right)^t + \left(e^{-\frac{k}{200}} + e^{-\frac{k}{1000}} \right) \cdot (U^u + 1),$$

as required. ■

5.4 Phased attacks

If we could assume that the attack under consideration is *phased*, in a sense that during cheating phase i the i ’th steps of all U auxiliary instances of \mathcal{V}_{er} are executed, that would simplify our analysis considerably. In this part we will show that any attack can be transformed, with a modest loss in the success probability, to the *nearly-phased* form.

Definition 8. (phased and nearly-phased attacks) *Let an attack be using U auxiliary instances of \mathcal{V}_{er} .*

We say that the scenario is phased if it can be viewed as consisting of 6 consecutive phases, such that at phase i the i ’th steps of all U auxiliary instances of \mathcal{V}_{er} are executed.

We call the scenario nearly-phased if it is phased with a relaxation that instead of phases 3 and 4 it has a phase called “3 - 4”, when both the 3’rd and the 4’th steps of the auxiliary instances of \mathcal{V}_{er} are executed.

Intuitively, the difference between the two restrictions is that in a nearly-phased scenario an adversary is allowed, say, to choose the $2t/3$ “playing” registers (out of the t suggested by the bank) in the auxiliary instance 1 of \mathcal{V}_{er} after he has received the $2t/3$ questions m_i relevant to the auxiliary instance 2 of \mathcal{V}_{er} . In the case of phased attacks such behavior is not allowed: the questions m_i relevant to all the auxiliary instances of \mathcal{V}_{er} are available to the adversary only after the choices of “playing” registers have been made w.r.t. all the instances.

The convenience of these definitions comes from the fact that, on the one hand, if an attack is phased then it cannot use in an earlier stage of one auxiliary instances of \mathcal{V}_{er} the output from a later stage of another instances, while on the other hand, only the last bank’s response in \mathcal{V}_{er} provides any information about the string \mathbf{x} . That is, assuming that an attack is (nearly-) phased limits

considerably the possibilities for the adversary to use dependencies between different instances of \mathcal{V}_{er} .

Our claim is the following.

Lemma 5.6. *If an attack exists that initiates at most U and wins at least $u \leq U$ auxiliary instances of \mathcal{V}_{er} with probability at least δ , then there is a nearly-phased scenario that initiates and completes exactly U and wins exactly u instances of \mathcal{V}_{er} with probability larger than*

$$\frac{\delta - U/2^{2t/3}}{U^u}.$$

In the above statement by “initiating” an instance of \mathcal{V}_{er} we mean sending a coin identification number to the bank and getting back a list of t registers (step 2 of \mathcal{V}_{er}).

Proof. The proof idea here is somewhat similar to that of Lemma 5.5 – namely, if the output from a later stage of one auxiliary instance of \mathcal{V}_{er} is used by the adversary in order to decide how to act in an earlier stage of another instance, we would let a “new adversary” guess the future response of the first instance before actually receiving it from the bank, and act in the second instance under the assumption that the guessing has been accurate.

Let \mathcal{C} be the attack, as guaranteed by the lemma condition.

First of all, let us turn it into \mathcal{C}' that always completes U instances of \mathcal{V}_{er} and is likely to win exactly u of them. This first modification is straightforward – \mathcal{C}' would behave as \mathcal{C} , except for the following modifications:

- If, according to \mathcal{C} , no more auxiliary instances are needed but less than U have been run then \mathcal{C}' runs “dummy” instances of \mathcal{V}_{er} (generating uniformly at random all messages that are sent to the bank), in order to make their total number equal U .
- If, according to \mathcal{C} , some instances of \mathcal{V}_{er} are aborted, \mathcal{C}' completes them in a “dummy” way.
- If at some point it occurs that \mathcal{C}' has already won u instances of \mathcal{V}_{er} , then it completes all remaining instances in a “dummy” way.

Note that the way \mathcal{C}' produces ρ_1 and ρ_2 is irrelevant for us now, as here we are only interested in the number of “accepts” among the preliminary runs of \mathcal{V}_{er} .

Clearly, the probability that \mathcal{C}' wins at least u instances is the same as in the case of \mathcal{C} ; on the other hand, \mathcal{C}' wins more than u instances only if at least one “dummy” instance has been won. A single “dummy” instance is won with probability exactly $2^{-2t/3}$, and at least one is won with probability less than $U \cdot 2^{-2t/3}$. Therefore, \mathcal{C}' wins exactly u instances of \mathcal{V}_{er} with probability larger than

$$\delta - U \cdot 2^{-2t/3}. \tag{6}$$

Now let us turn \mathcal{C}' into nearly-phased. The new attack \mathcal{C}'' consists of 5 phases, as follows.

1. Initiate U instances of \mathcal{V}_{er} , sending the identification number of \mathbf{c} to the bank. Index the instances by $1, \dots, U$.
2. Get back U t -tuples, denoting by $T_i^{(1)}$ the response from the i 'th instance of \mathcal{V}_{er} , $i \in [U]$.
- 3 - 4. Let $W \in \{0, 1\}^U$ be a uniformly chosen binary vector of Hamming weight u – this is going to be the adversary's guess regarding the winning instances of \mathcal{V}_{er} . Start emulating \mathcal{C}' skipping the first two steps of each instance of \mathcal{V}_{er} , as those have been processed already (use $T_i^{(1)}$'s as

bank's responses in step 2 of \mathcal{V}_{er}). Skip all interaction with the bank beyond step 4; instead, whenever \mathcal{C}' acts depending on the bank's final response in the i 'th instance of \mathcal{V}_{er} , emulate \mathcal{C}' assuming that the response was W_i (where "0" stands for "reject", and "1" stands for "accept"). Run the emulation until the bank's responses in step 4 have been received in all U instances of \mathcal{V}_{er} . For $i \in [U]$, denote by $T_i^{(2)}$ the $2t/3$ -tuple chosen in step 3 of the i 'th instance of \mathcal{V}_{er} , and by $M_i \in \{0, 1\}^{2t/3}$ the values sent by the bank in step 4 of the i 'th instance.

5. Start a new emulation of \mathcal{C}' , this time skipping steps 1 - 4 of each instance of \mathcal{V}_{er} and respectively using the values $T_i^{(1)}$ and M_i as bank's responses. Do not interact with the bank beyond step 5; instead, whenever \mathcal{C}' acts depending on the bank's final response in the i 'th instance of \mathcal{V}_{er} , emulate \mathcal{C}' assuming that the response was W_i .
6. Receive the final responses from the bank, denote them by $V \in \{0, 1\}^U$.

It is clear from the construction that \mathcal{C}'' is nearly-phased.

Let us analyze the probability that \mathcal{C}'' wins exactly u instances of \mathcal{V}_{er} . It is lower bounded by the probability that $V = W$, and that equals the probability that \mathcal{C}' wins u instances of \mathcal{V}_{er} and the right W has been guessed in the beginning of phase "3 - 4" of \mathcal{C}'' . The string $W \in \{0, 1\}^U$ is uniformly random of Hamming weight u , thus it is correct with probability at least U^{-u} . Therefore, (6) implies that \mathcal{C}'' wins exactly u instances of \mathcal{V}_{er} with probability larger than $(\delta - U \cdot 2^{-2t/3}) \cdot U^{-u}$, as required. \blacksquare

5.5 Phased cheating is slow

In this section we will prove that nearly-phased attacks require many auxiliary instances of \mathcal{V}_{er} in order to win enough of them for Lemma 5.5 to allow non-negligible counterfeiting success probability.

Lemma 5.7. *A nearly-phased attack that initiates and completes U auxiliary instances of \mathcal{V}_{er} wins at least $3k/t$ of them with probability at most*

$$e^{2 \ln U - \Omega(t^2/k)}.$$

As before, by "initiating" an instance of \mathcal{V}_{er} we mean sending a coin identification number to the bank and getting back a list of t registers (step 2 of \mathcal{V}_{er}).

Proof. Let \mathcal{C} be the nearly-phased attack under consideration. For $i \in [U]$, let random variables $T_i^{(1)}$, $T_i^{(2)}$ and M_i describe the transcript of the i 'th instance of \mathcal{V}_{er} , as follows:

- $T_i^{(1)}$ takes the value of the t -tuple chosen by the bank in step 2;
- $T_i^{(2)}$ takes the value of the $2t/3$ -tuple chosen by the adversary in step 3;
- $M_i \in \{0, 1\}^{2t/3}$ contains the $2t/3$ "questions" chosen by the bank in step 4.

For $j \in T_i^{(1)}$, let $T_i^{(1)}[j]$ be the position of j in $T_i^{(1)}$, and similarly define $T_i^{(2)}[j]$. For $j \in T_i^{(2)}$, let $M_i[j]$ be the $T_i^{(2)}[j]$ 'th bit of the value received by M_i – that is, $M_i[j]$ denotes the HMP_4 -query asked in the i 'th instance of \mathcal{V}_{er} w.r.t. the register j .

For $i, j \in [U]$, $i \neq j$, let $S_{i,j}^{(2)} \stackrel{\text{def}}{=} T_i^{(2)} \cap T_j^{(2)}$ (viewed as a set) and

$$\tilde{S}_{i,j}^{(2)} \stackrel{\text{def}}{=} \left\{ M_i[s] \neq M_j[s] \mid s \in S_{i,j}^{(2)} \right\}. \quad (7)$$

That is, $S_{i,j}^{(2)}$ contains the registers of \mathbf{c} that are part of the bank's challenge questions both in the i 'th and in the j 'th auxiliary instances of \mathcal{V}_{er} , and $\tilde{S}_{i,j}^{(2)}$ contains the registers where good answers to the both possible HMP_4 -queries have to be found in order to pass both the i 'th and the j 'th instances of \mathcal{V}_{er} . Note that the attack \mathcal{C} produces answers to all the relevant HMP_4 -queries not having any auxiliary information about \mathbf{x} (\mathcal{C} is phased, and the only bank's responses that contain information about \mathbf{x} are the final ones, which were not available to \mathcal{C} at the earlier phase).

Denote by \mathcal{W}_i the event that the i 'th instance of \mathcal{V}_{er} is won, and by $\mathcal{W}_{i,j}$ the event that both the i 'th and the j 'th instances are won. For every $i \neq j$, Corollary 5.4 implies that

$$\Pr \left[\mathcal{W}_{i,j} \mid T_1^{(2)}, \dots, T_{[U]}^{(2)}, M_1, \dots, M_{[U]} \right] \leq (3/4)^{|\tilde{S}_{i,j}^{(2)}|}. \quad (8)$$

Let $\tilde{r} \in \mathbb{N}$, and denote by $\tilde{\mathcal{E}}$ the event that $\mathcal{W}_{i,j}$ does not hold whenever $|\tilde{S}_{i,j}^{(2)}| \geq \tilde{r}$ (later we will fix \tilde{r} to make $\mathcal{W}_{i,j}$ very likely to occur). Then from (8),

$$\Pr \left[\tilde{\mathcal{E}} \right] \geq 1 - U^2 \cdot (3/4)^{\tilde{r}}. \quad (9)$$

Similarly, let $r \in \mathbb{N}$ (to be fixed later), and denote by \mathcal{E} the event that $\mathcal{W}_{i,j}$ does not hold whenever $|S_{i,j}^{(2)}| \geq r$. Let \mathcal{E}' be the event that $|\tilde{S}_{i,j}^{(2)}| \geq \tilde{r}$ whenever $|S_{i,j}^{(2)}| \geq r$, then from (9),

$$\Pr \left[\mathcal{E} \right] \geq \Pr \left[\mathcal{E}' \right] - U^2 \cdot (3/4)^{\tilde{r}}. \quad (10)$$

When \mathcal{E} holds, any two different elements of the family

$$\mathcal{F} \stackrel{\text{def}}{=} \left\{ T_i^{(2)} \mid \mathcal{W}_i \text{ holds} \right\}$$

share less than r elements. We choose

$$r \stackrel{\text{def}}{=} \left\lfloor \frac{2t^2}{9k} \right\rfloor,$$

then Lemma 3.4 implies that $|\mathcal{F}| < 3k/t$, i.e.,

$$\mathcal{E} \text{ holds} \implies \text{Less than } 3k/t \text{ instances of } \mathcal{V}_{er} \text{ are won.} \quad (11)$$

It remains to show that \mathcal{E} is likely to hold. Fix

$$\tilde{r} \stackrel{\text{def}}{=} \left\lfloor \frac{t^2}{10k} \right\rfloor,$$

and let us see that \mathcal{E}' is very likely to occur.

Before we deal with the nearly-phased case, suppose that \mathcal{C} is phased. In this case there is no dependence between the variables $(T_i^{(2)})_{i=1}^U$ and the variables $(M_i)_{i=1}^U$, and therefore $\tilde{S}_{i,j}^{(2)}$ is a randomly chosen subset of $S_{i,j}^{(2)}$, where each $s \in S_{i,j}^{(2)}$ independently becomes an element of $\tilde{S}_{i,j}^{(2)}$ with probability $1/2$. By Chernoff bound (Theorem 3.1), if $|S_{i,j}^{(2)}| \geq r$ then $\Pr \left[|\tilde{S}_{i,j}^{(2)}| < \tilde{r} \right] \leq e^{-\Omega(r)}$, and therefore

$$\Pr \left[\mathcal{E}' \right] \geq 1 - U^2 \cdot e^{-\Omega(r)}. \quad (12)$$

When \mathcal{C} is nearly-phased, the variables $(T_i^{(2)})_{i=1}^U$ are not necessarily independent from $(M_i)_{i=1}^U$ (the adversary is allowed to choose the $2t/3$ "playing" registers in step 3 of the i 'th instance of \mathcal{V}_{er} ,

depending on M_j received from the bank in step 4 of the j 'th instance of \mathcal{V}_{er} , $j \neq i$). However, we claim that for every $j \neq i$ the values $\{M_i[s] \oplus M_j[s] \mid s \in T_i^{(2)} \cap T_j^{(2)}\}$ are unbiased and mutually independent – and this is all we need for (12) to hold (cf. (7)). Indeed, from the definition of \mathcal{V}_{er} it is clear that at least one of M_i and M_j is chosen by the bank uniformly at random after the values of both $T_i^{(2)}$ and $T_j^{(2)}$ have been set by the choice of the adversary, and therefore $M_i[s] \oplus M_j[s]$ is unbiased.

From (10) and (12),

$$\Pr[\mathcal{E}] \geq 1 - e^{2 \ln U - \Omega(t^2/k)}.$$

Together with (11), this implies the result. ■

5.6 \mathcal{Q} is safe

We are ready to prove the main theorem.

Theorem 5.1: *Let a fresh \mathcal{Q} -coin be given to a computationally unlimited adversary who runs auxiliary instances of \mathcal{V}_{er} for this coin and produces two (possibly, entangled) “counterfeits” ρ_1 and ρ_2 . Then*

$$U \in e^{\Omega(t^3/k^2)}$$

exists, such that if the adversary has received and analyzed the first bank’s responses in at most U instances of \mathcal{V}_{er} , then the probability that both ρ_1 and ρ_2 pass \mathcal{V}_{er} is in

$$e^{-\Omega(t^2/k)}.$$

Proof. From Lemmas 5.6 and 5.7 it follows that an attack \mathcal{C} that receives the first responses in at most U auxiliary instances of \mathcal{V}_{er} can win at least $3k/t$ of them with probability at most

$$e^{O(k/t) \ln U - \Omega(t^2/k)}.$$

Then Lemma 5.5 implies that \mathcal{C} succeeds in counterfeiting the coin \mathbf{c} with probability at most

$$e^{O(k/t) \ln U - \Omega(t^2/k)},$$

and the result follows. ■

6 Optimality of \mathcal{Q}

In this part we consider a generic quantum money scheme with classical verification, where the qubit-size of a coin is K and a secret bank record describing a coin contains R bits.

Let us define the *counterfeiting complexity* of a quantum money scheme as

$$\min_{\varepsilon} \{ \max \{ 1/\varepsilon, \langle \text{time required to counterfeit a coin with success probability at least } \varepsilon \rangle \} \},$$

this definition is a lower bound on what we intuitively mean by “time required to forge a counterfeit”.¹⁰ Note that Theorem 5.1 and Corollary 5.2 imply that the counterfeiting complexity of \mathcal{Q} is exponential both in K and in R .

¹⁰Instead, one might consider the time required to counterfeit a coin with *constant* success probability. The (asymptotic) time complexity of an attack that succeeds with constant probability is an upper bound on the counterfeiting complexity, as defined above. Note that our scheme from the previous section has high counterfeiting complexity, therefore it is secure in the stronger sense. On the other hand, the upcoming (formal) optimality statements will be made w.r.t. attacks that achieve constant success probability, which will make those statements also as strong as possible. Intuition-wise, we find the definition with “flexible ε ” more appealing, that is why we use it in the informal discussion.

First of all, 2^R adversarial verification attempts are enough to exhaustively check all possible bank's records, and therefore $O(2^R)$ is an upper bound on the counterfeiting complexity of any quantum scheme. So, in the case of \mathcal{Q} the length of a bank record as a function of counterfeiting complexity is polynomially-close to optimal.

Can the counterfeiting complexity be super-exponential in K ? We could not find a simple argument against this possibility. The counterfeiting complexity of \mathcal{Q} is exponential in K (which can probably be viewed as “reasonably good”), and we leave the question above as an open problem.

There is one parameter in the construction of \mathcal{Q} that one might like to improve – namely, the number of verification rounds that a new quantum coin can go through before it has to be returned to the bank. In this section we show that no scheme can allow this number to be larger than linear in K , and therefore our construction is polynomially-close to the optimal in this respect also.

Theorem 6.1. *Let T be the number of times that a new coin can be verified via classical communication with the bank before it has to be replaced. Suppose that if a fair user verifies a fresh coin T times in a sequence then all T verifications are passed with probability at least $8/9$. Then either T auxiliary instances of the verification protocol are sufficient for an adversary to counterfeit a coin with probability at least $2/3$, or a coin contains $\Omega(T)$ qubits (or both).*

To prove the theorem we will need the following technical statement (which might be of independent interest).

Lemma 6.2. *Let A and B be discrete random variables, such that there exists a condition that can be satisfied with probability at most α by the value of any random variable independent from A . If the value of B satisfies the condition with probability at least $\beta \geq \alpha$, then*

$$I(A : B) \geq 2(\beta - \alpha)^2.$$

First we prove the theorem, then the lemma.

Proof sketch of Theorem 6.1. Assume that more than T auxiliary instances of the verification procedure are required for an adversary to counterfeit a coin with probability at least $2/3$.

To argue that a coin consists of $\Omega(T)$ qubits, let us show that its “quantum part” has mutual information $\Omega(T)$ with bank's secret record. To make sure that we are not counting information carried by the classical part of a coin, let us assume w.l.g. that the first message of the verification protocols is sent by the coin holder to the bank and contains all the classical information that the coin contained when it was fresh.

Let L_1, \dots, L_T be random variables, respectively taking values of the transcripts of T sequentially executed protocols for coin verification via classical communication (assuming that the coin holder fairly follows the protocol, and that the coin was fresh when the first verification started). For convenience (and w.l.g.), assume that a transcript provides complete information about the action taken by a fair coin holder w.r.t. the coin being verified. Let ρ be the mixed state of a fresh coin whose secret record is unknown, and for every $j \in [T]$ and ℓ_1, \dots, ℓ_j , let $\rho_{\ell_1, \dots, \ell_j}$ be the state of a fresh coin that went through j verification protocols whose transcripts were, respectively, ℓ_1, \dots, ℓ_j .

Let $(\ell_i)_{i=1}^T$ be the values taken by $(L_i)_{i=1}^T$. Denote by \mathcal{R} a random variable describing the bank's secret record corresponding to the coin under consideration. Let $S(\cdot, \cdot)$ denote quantum mutual information, we claim that

$$\forall i \in [T] : \mathbf{E} [S(\rho_{\ell_1, \dots, \ell_{i-1}} : \mathcal{R}) - S(\rho_{\ell_1, \dots, \ell_i} : \mathcal{R})] \geq \frac{1}{2592}, \quad (13)$$

where the expectation is taken w.r.t. the choice of $(L_i)_{i=1}^T$. From (13) it follows that $S(\rho : \mathcal{R}) \in \Omega(T)$, and Holevo's bound implies that ρ consists of $\Omega(T)$ qubits, as required.

To prove (13) we will also use Holevo's bound. For simplicity let us assume that each run of the verification protocol requires exactly one quantum measurement to be performed by the coin holder (the case of many measurements is treated similarly).

Fix $i \in [T]$. Observe that if during the i 'th run of the verification protocol the coin holder would not perform any quantum measurement, instead making the "best guess" responses based on the previous transcripts $\ell_1, \dots, \ell_{i-1}$, then the probability to pass the verification would be less than $7/8$ – otherwise an adversary could, based on the transcripts $\ell_1, \dots, \ell_{i-1}$, prepare two counterfeits that would both pass the verification with probability at least $(7/8)^2 > 2/3$, contradicting the assumptions of the theorem. On the other hand, by making the quantum measurement, as prescribed by the verification protocol, a fair coin holder is able to pass the verification with probability at least $8/9$, also guaranteed by the theorem assumptions.

Conditional on $\ell_1, \dots, \ell_{i-1}$ being the values taken by L_1, \dots, L_{i-1} , the following holds. The acceptance condition of the i 'th verification can be satisfied with probability at most $7/8$ by a random variable that doesn't depend on \mathcal{R} ; at the same time the outcome of the quantum measurement performed by the coin holder satisfies the condition with probability at least $8/9$. Therefore, Lemma 6.2 implies that the expected conditional mutual information between the measurement outcome and \mathcal{R} is at least $1/2592$. Holevo's bound implies (13), and the result follows. \blacksquare *Theorem 6.1*

Proof of Lemma 6.2. W.l.g., assume that the condition under consideration is a function of the value taken by B . Let X and Y be the supports of A and B , respectively. For every $b \in Y$, let $X_b \subseteq X$ be the set of values of A that satisfy the condition when $B = b$. Let μ be the distribution of A , and let μ_b be the distribution of A conditional upon $B = b$. Let $\alpha_b \stackrel{\text{def}}{=} \mu(X_b)$ and $\beta_b \stackrel{\text{def}}{=} \mu_b(X_b)$. The requirements of the lemma assure that

$$\mathbf{E}_{B=b} [\alpha_b] \leq \alpha \quad \text{and} \quad \mathbf{E}_{B=b} [\beta_b] \geq \beta. \quad (14)$$

By definition,

$$I(A : B) = \mathbf{E}_{B=b} [d_{KL}(\mu_b || \mu)] = \mathbf{E}_{B=b} \left[\sum_x \mu_b(x) \cdot \log \left(\frac{\mu_b(x)}{\mu(x)} \right) \right], \quad (15)$$

and this is the value we want to bound from below. We have

$$\begin{aligned} \sum_{x \in X_b} \mu_b(x) \cdot \log \left(\frac{\mu_b(x)}{\mu(x)} \right) &= \alpha_b \sum_{X_b} \frac{\mu_b(x)}{\alpha_b} \cdot \log \left(\frac{\mu_b(x)/\alpha_b}{\mu(x)/\beta_b} \cdot \frac{\alpha_b}{\beta_b} \right) \\ &= \alpha_b \cdot \left(d_{KL} \left(\frac{\mu_b}{\alpha_b} \middle\| \frac{\mu}{\beta_b} \right) + \log \left(\frac{\alpha_b}{\beta_b} \right) \right) \geq \alpha_b \cdot \log \left(\frac{\alpha_b}{\beta_b} \right), \end{aligned}$$

where the inequality follows from non-negativity of $d_{KL}(\cdot || \cdot)$ and the fact that, restricted to X_b , both μ_b/α_b and μ/β_b are probability distributions. Similarly,

$$\sum_{x \notin X_b} \mu_b(x) \cdot \log \left(\frac{\mu_b(x)}{\mu(x)} \right) \geq (1 - \alpha_b) \cdot \log \left(\frac{1 - \alpha_b}{1 - \beta_b} \right),$$

leading to

$$d_{KL}(\mu_b || \mu) \geq \alpha_b \cdot \log \left(\frac{\alpha_b}{\beta_b} \right) + (1 - \alpha_b) \cdot \log \left(\frac{1 - \alpha_b}{1 - \beta_b} \right) = d_{KL}(\mathcal{B}_{\alpha_b} || \mathcal{B}_{\beta_b}) \geq 2(\alpha_b - \beta_b)^2,$$

where \mathcal{B}_p denotes Bernoulli distribution with probability p for outcome “1”, and the last inequality follows from Pinsker’s inequality. Plugging this into (15), we obtain the desired

$$I(A : B) \geq 2 \mathbf{E}_{B=b} [(\alpha_b - \beta_b)^2] \geq 2(\beta - \alpha)^2,$$

where the last inequality follows from (14).

■ *Lemma 6.2*

7 Conclusions

We constructed a quantum money scheme \mathcal{Q} that allows verifying a coin via classical communication with a bank. Thus we are proving existence of secure quantum money schemes that do not require quantum communication for coin verification.

Our scheme has the following properties.

- The coins are exponentially hard to counterfeit, even if an adversary is adaptively using repeated verification attempts in order to collect information about a coin.
- The classical communication channel used for verification can be unsecured.
- The database of the bank is static.
- The dependence between the number of verifications that a \mathcal{Q} -coin can go through and the number of qubits that it contains is optimal, up to a polynomial.

There are (at least) two questions that remain open:

- Is it possible to build *anonymous* quantum money schemes with classical verification, by allowing multiple identical instances of quantum coins, as suggested by Mosca and Stebila [MS10]?
- Is it possible to have the counterfeiting complexity of quantum money super-exponential in the number of qubits that a coin contains (cf. Section 6)?

Acknowledgments

I am grateful to Scott Aaronson and Martin Rötteler for numerous helpful discussions. Dana Moshkovitz has drawn my attention to the result of [IK10]. I have received many helpful comments from Ronald de Wolf, and various comments from anonymous reviewers, mostly helpful. I acknowledge support by ARO/NSA under grant W911NF-09-1-0569.

References

- [Aar09] S. Aaronson. Quantum Copy-Protection and Quantum Money. *Proceedings of the 24th IEEE Conference on Computational Complexity*, pages 229–242, 2009.
- [AC12] S. Aaronson and P. Christiano. Quantum Money from Hidden Subspaces. *Proceedings of the 44th Symposium on Theory of Computing*, 2012.
- [BBBW83] C. H. Bennett, G. Brassard, S. Breidbart, and S. Wiesner. Quantum Cryptography, or Unforgeable Subway Tokens. *Advances in Cryptology – Proceedings of Crypto 82*, pages 267–275, 1983.

- [BYJK04] Z. Bar-Yossef, T. S. Jayram, and I. Kerenidis. Exponential Separation of Quantum and Classical One-Way Communication Complexity. *Proceedings of 36th Symposium on Theory of Computing*, pages 128–137, 2004.
- [FGH⁺10] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor. Quantum Money from Knots. <http://arxiv.org/abs/1004.5127>, 2010.
- [IK10] R. Impagliazzo and V. Kabanets. Constructive Proofs of Concentration Bounds. *Proceedings of APPROX-RANDOM*, pages 617–631, 2010.
- [Juk01] S. Jukna. Extremal Combinatorics With Applications in Computer Science. *Springer-Verlag*, 2001.
- [KdW04] I. Kerenidis and R. de Wolf. Exponential Lower Bound for 2-Query Locally Decodable Codes via a Quantum Argument. *Journal of Computer and System Sciences* 69(3), pages 395–420, 2004.
- [LAF⁺10] 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. *Proceedings of the 1st Symposium on Innovations in Computer Science*, pages 20–31, 2010.
- [Lut10] A. Lutomirski. An Online Attack Against Wiesner’s Quantum Money. <http://arxiv.org/abs/1010.0256>, 2010.
- [MS10] M. Mosca and D. Stebila. Quantum Coins. *Error-Correcting Codes, Finite Geometries and Cryptography – American Mathematical Society*, pages 35–46, 2010.
- [MVW12] A. Molina, T. Vidick, and J. Watrous. Optimal Counterfeiting Attacks and Generalizations for Wiesner’s Quantum Money. <http://arxiv.org/abs/1202.4010>, 2012.
- [PS97] A. Panconesi and A. Srinivasan. Randomized Distributed Edge Coloring via an Extension of the Chernoff-Hoeffding Bounds. *SIAM Journal on Computing* 26(2), pages 350–368, 1997.
- [PYJ⁺11] F. Pastawski, N. Y. Yao, L. Jiang, M. D. Lukin, and J. I. Cirac. Unforgeable Noise-Tolerant Quantum Tokens. <http://arxiv.org/abs/1112.5456>, 2011.
- [TOI03] Y. Tokunaga, T. Okamoto, and N. Imoto. Anonymous Quantum Cash. *ERATO Conference on Quantum Information Science*, 2003.
- [Wie83] S. Wiesner. Conjugate Coding. *SIGACT News*, Vol. 15(1), pages 78–88, 1983.