Quantum money

Yuval Efron, Dean Hirsch

April 25, 2022

1 Introduction

In the 60's, Wiesner [6] came up a with a scheme, leveraging the well known no-cloning theorem, for un-counterfeitable currency. Such a primitive is clearly unattainable in a classical world, as any piece of classical data can be easily copied. Wiesner's scheme consisted of a trusted bank, acting as a mint, that can manufacture and issue currency in the form of quantum state $|\$\rangle$, that can be used by players as currency. The bank is the only entity that can issue valid banknotes, using some secret key sk. In Wiesner's scheme, the bank is also the sole entity that is able to verify the validity of a purported banknote $|\psi\rangle$, using a verification procedure $\mathsf{Ver}(sk, |\psi\rangle)$. Wiesner's scheme has two crucial properties: First, *non-clonability*, i.e. a polynomially bounded quantum adversary $\mathcal{A}$, given a single valid banknote $|\$\rangle$, can not produce two (possibly entangled) states $|\psi_1\rangle, |\psi_2\rangle$ such that $\mathsf{Ver}(sk, |\psi_1\rangle)$, $\mathsf{Ver}(sk, |\psi_2\rangle)$ both accept concurrently with non-negligible probability. The second property, *verfiability*, states that $\mathsf{Ver}(sk, |\$\rangle)$ always accepts a valid banknote $|\$\rangle$, and for any purported banknote $|\psi\rangle$, which is sufficiently far from being a valid banknote, $\mathsf{Ver}(sk, |\psi\rangle)$ rejects w.h.p.

Drawbacks. The pioneering scheme of Wiesner has some evident drawbacks. The prime one being that only the bank is able to verify the validity of banknotes, i.e., each time two parties engage in currency exchange, they must involve the bank in order to verify the soundness of the interaction. Additional drawbacks of the original scheme of Wiesner are the size of the secret verification key, and its interceptability to interactive attacks. The last two drawbacks have relatively easy fixes. The main objective of future work was to design a quantum money scheme that had no dependence on the bank for verification. From that goal, the notion of Public-Key Quantum Money (PKQM) emerged.

1.1 Public-Key Quantum Money

The holy grail of the line of work revolving Quantum Money is to construct PKQM (which will be defined shortly) from standard cryptographic assumptions, such as one-way functions. Roughly, a PKQM scheme is a scheme in which only the bank can manufacture banknotes using a secret key, but all players can verify the validity of purported banknotes using a public key and an efficient verification procedure. Such a scheme should be secure against forgery in the sense that any efficient adversary with a single valid banknote, can not produce two (possibly entangled) purported banknotes such that the verifier accepts both banknotes.

Since Wiesner's seminal work, numerous works have studied the problem from various angles on the way to achieve PKQM from standard assumptions [1–4, 7]. Including considering stronger assumptions [2], and notions slightly weaker than PKQM [4].

In this survey, we give a flavour of the recent progress on the subject by diving deeper into three recent results in this line of work. We elaborate on them in the following paragraphs.

Quantum Money from Hidden Subspaces. The idea of quantum money from hidden subspaces was introduced by Aaronson and Christiano [2], in which they were able to construct a *private key* money scheme that is *unconditionally* secure, and a conditionally secure *public key* money scheme, but the money still requires a classical oracle to be publicly accessible (hence to verifiers and potential counterfeiters).

The authors also provide an explicit suggestion, based on systems of polynomial equations, for implementing the classical oracle. This means the bank distributes an implementation of such an oracle for the public to query, where the implementation is conjectured to not reveal enough information about the money to be useful for any other purpose.

Here, in order to mint a new banknote, the bank randomly selects an $n/2$-dimensional subspace A of $\mathbb{F}_2^n$ for a suitable even n (this is the "hidden subspace"). The note is then $|\$\rangle = \frac{1}{\sqrt{|A|}} \sum_{v \in A} |v\rangle$, and a membership oracle for A is given. That is, there is public access to an oracle U_A for which any $v \in A$ satisfies $U_A|v\rangle = |v\rangle$, and all other v's satisfy $U_A|v\rangle = -|v\rangle$. Likewise, an oracle for $U_{A^\perp}$ is also given in the same way, where $A^\perp \subseteq \mathbb{F}_2^n$ is the subset of vectors v that are orthogonal over $\mathbb{F}_2$ to all elements of A. It is important to remember that over $\mathbb{F}_2$ the dot product is not an inner product, and in particular $A^\perp$ and A can intersect.

Some key observations now allow these quantum states to be used as a quantum money. The authors are also able to prove *unconditionally* that using U_A and $U_{A^\perp}$ as black-box classical oracles, it is impossible for a poly-time counterfeiter to obtain a state close to $|A\rangle \otimes |A\rangle$ given a state $|A\rangle$ even with polynomially-many queries to U_A and $U_{A^\perp}$. In fact, they prove a strictly exponential lower bound for this task.

Franchised Quantum Money (FQM). The notion of FQM, introduced by Roberts and Zhandry [4], is a relaxation of PKQM in the following sense. In FQM, the bank is the sole entity that can mint valid bank notes using a secret key sk. Furthermore, every player P using and exchanging banknotes has their *own secret verification key* ver_P, with which they are able to verify whether any purported bank note is valid. Since verification keys are private, FQM needs to be secure not only against counterfeiting, but also against *sabotage*. Consider an adversary holding a valid banknote $|\$\rangle$, which they then modify slightly, such that one honest player P_1 still accepts the banknote, but when P_1 tries to use $|\$\rangle$ and give it to P_2, P_2 rejects it. Note that security against counterfeiting does not guarantee security against such an attack, as the adversary doesn't need to produce additional banknotes from a given banknote in order to succeed in sabotage. We say that an FQM scheme is secure against sabotage if an adversary cant produce a purported banknote $|\psi\rangle$ that is accepted by one player, but rejected by another. Another concern we address in this survey is that of security against colluding adversaries, i.e. a potential family of adversaries using the combined knowledge of their verification keys in order to counterfeit banknotes.

In this survey, we present and discuss the result of [4], in which FQM is constructed using only the most basic assumption in cryptography, namely, existence of one-way functions.

Quantum Lightning. The notion of Quantum Lightning (QL), introduced by Zhandry [7], is in some sense, a generalization of Quantum money. A QL scheme consists of two procedures $(\mathcal{S}, \mathcal{V})$, the first referred to as a *storm*, and the latter, a verifier. $\mathcal{S}$ generates quantum states $|\lightning\rangle$, which we refer to as *bolts*. For the validity of the scheme, we require the following.

- $\mathcal{V}$ always accepts bolts generated by $\mathcal{S}$, and extracts from a given bolt a classical fingerprint/serial number.
- $\mathcal{V}$'s execution on a given bolt $|\lightning\rangle$ does not perturb $|\lightning\rangle$ by more than a negligible amount.
- Repeated executions of $\mathcal{V}$ on a given bolt $|\lightning\rangle$ always yield the same serial number.

For security of the scheme, we require that any efficient adversarial storm $\mathcal{S}'$ can not produce two (perhaps entangled) bolts $|\lightning_1\rangle$, $|\lightning_2\rangle$ such that $\mathcal{V}$ accepts both bolts and outputs *the same* serial number for both bolts.

Given such a primitive, constructing PKQM is straightforward, banknotes are simply bolts produced by bank using $\mathcal{S}$, with the serial number signed using a classical signature scheme (for which one-way functions are known to suffice [5]), the public verification procedure is simply $\mathcal{V}$. For forgery, any efficient adversary must either forge the bank's signature, or produce two bolts with the same serial number.

For those familiar with classical cryptography, the notion of QL might sound related to Collision Resistant Hash Functions (CRHF), and one would wonder whether a QL scheme can be designed assuming CRHF. It turns out that another key property is required. Consider a uniform superposition over inputs to a CRHF H, on which one then applies H and measures the result y, which is some string in the range of H. Note that the post measurement state $|\psi_y\rangle$ is the uniform superposition over all pre-images of y. The key issue now is whether there is an efficient verification procedure that can *distinguish* between $|\psi_y\rangle$ and some specific $|x\rangle, x \in H^{-1}(y)$. If that is *not* the case, we call H *collapsing*. With this notion in mind, one of the main results of [7] is that if H is a CRHF against quantum adversaries, then it is either collapsing, *or* it can be used to build quantum lightning without any further assumptions. We state the result formally and showcase its proof in Section 2.3

2 Progress Towards PKQM

2.1 Quantum Money from Hidden Subspaces

Here, in order to mint a new banknote, the bank randomly selects an $n/2$-dimensional subspace A of $\mathbb{F}_2^n$ for a suitable even n. The note is then $|\$\rangle = \frac{1}{\sqrt{|A|}} \sum_{v \in A} |v\rangle$, and a membership oracle for A is given. That is, there is public access to the oracle U_A for which any $v \in A$ satisfies $U_A|v\rangle = |v\rangle$, and all other v's satisfy $U_A|v\rangle = -|v\rangle$. Likewise, an oracle for $U_{A^\perp}$ is also given in the same way, where $A^\perp \subseteq \mathbb{F}_2^n$ is the subset of vectors v that are orthogonal over $\mathbb{F}_2$ to all elements of A. It is important to remember that over $\mathbb{F}_2$ the dot product is not an inner product, and in particular $A^\perp$ and A can intersect.

Some key observations now allow these quantum states to be used as a quantum money:

Claim 1. *Using the oracles U_A and $U_{A^\perp}$ it is possible to project states onto A, $A^\perp$.*

First, using the given oracles U_A and $U_{A^\perp}$ one can implement a projection $\mathbb{P}_A$ on A, and likewise $\mathbb{P}_{A^\perp}$ on $A^\perp$. Indeed, the following circuit computes $\mathbb{P}_A$:

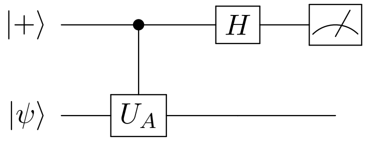

Where we apply the U_A gate conditioned on the control bit being in $|1\rangle$. Now if $|\psi\rangle = \alpha_A|\psi_A\rangle + \alpha_{\bar{A}}|\psi_{\bar{A}}\rangle$ with $|\psi_A\rangle$ being the part of ψ consisting of basis elements from A, and $|\psi_{\bar{A}}\rangle$ consists of the rest, the state of the system after querying the oracle is $\alpha_A|+\rangle \otimes |\psi_A\rangle + \alpha_{\bar{A}}|-\rangle \otimes |\psi_{\bar{A}}\rangle$, so measuring in the $\{|+\rangle, |-\rangle\}$ basis, or equivalently the $\{|0\rangle, |1\rangle\}$ basis after a Hadamard transform, will collapse this to $|+\rangle \otimes |\psi_A\rangle$ with probability $|\alpha_A|^2$. When we will apply $\mathbb{P}_A$ we will do it for states we expect will be in $\mathrm{Span}(\{|v\rangle\}_{v\in A})$ and $A^\perp$, for otherwise the initial banknote would be invalid. Clearly, in these cases we will reject the banknote if we get the wrong measurement, and otherwise $\mathbb{P}_A$ and $\mathbb{P}_{A^\perp}$ will have no effect. Assuming we do not easily reject the notes, $\mathbb{P}_A$ and $\mathbb{P}_{A^\perp}$ can be viewed simply as projections.

Claim 2. $H_2^{\otimes n}|A\rangle = |A^\perp\rangle$.

Proof. Recall that for $u, v \in \mathbb{F}_2^n$ we have $\langle v|H|u\rangle = \frac{1}{\sqrt{2^n}}(-1)^{v\cdot u}$, while the statement $H|A\rangle = |A^\perp\rangle$ is equivalent to $\langle v|H|A\rangle = \begin{cases} 1/\sqrt{|A^\perp|} & v \in A^\perp \\ 0 & \text{otherwise} \end{cases}$.

Indeed:

$$\langle v|H|A\rangle = \frac{1}{\sqrt{|A|}}\sum_{u\in A}\langle v|H|u\rangle = \frac{1}{\sqrt{2^n|A|}}\sum_{u\in A}(-1)^{u\cdot v}$$

Now if $v \in A^\perp$ then all terms are $(-1)^0 = 1$, hence this sums to $\frac{|A|}{\sqrt{|A|2^n}}$. But $\dim A + \dim A^\perp = n \Rightarrow |A| \cdot |A^\perp| = 2^n$, hence this expression is equal to $\frac{1}{\sqrt{|A^\perp|}}$.

On the other hand, it can be seen that if $v \notin A^\perp$ then the sum is zero, by a pairing argument using some $v' \in A$ such that $v \cdot v' = 1$, or indirectly by noticing that there is no norm left to be distributed on other coefficients. □

Claim 3. *The verifier can check that $|A\rangle$ is valid.*

This is done by noting that, since $H_2^{\otimes n}|A\rangle = |A^\perp\rangle$, we have in particular that running $V_A = H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}\mathbb{P}_A$ on the purported banknote $|\psi\rangle$ should output not fail any of the projections, and returns to its original state. Clearly, the state $|A\rangle$ passes this test with certainty. We will now show that states far from $|A\rangle$ (which is the overwhelming majority of states) pass this test with low probability.

Claim 4.

$$Pr(V_A \textit{ accepts } |\psi\rangle) = |\langle\psi|A\rangle|^2$$

Proof. Let $|\psi\rangle$ be any purported banknote, and decompose into $|\psi\rangle = \alpha|A\rangle + \beta|\psi_A\rangle + \gamma|\psi_{\bar{A}}\rangle$, where $|\psi_A\rangle$ is a state in $|A\rangle^\perp \cap \mathrm{Span}\{|v\rangle : v \in A\}$ and $|\psi_{\bar{A}}\rangle$ is a state in $\mathrm{Span}\{|v\rangle : v \notin A\}$, and furthermore $|\alpha|^2 + |\beta|^2 + |\gamma|^2 = 1$. We will show that

$$\Pr(V_A \text{ accepts } |\psi\rangle) = |\alpha|^2 = |\langle\psi|A\rangle|^2$$

. Indeed, passing the $\mathbb{P}_A$ projection happens with probability $|\alpha|^2 + |\beta|^2$, after which the state becomes $\alpha'|A\rangle + \beta'|\psi_A\rangle$ with $\alpha' = \frac{\alpha}{\sqrt{|\alpha|^2+|\beta|^2}}$ and $\beta' = \frac{\beta}{\sqrt{|\alpha|^2+|\beta|^2}}$. It remains to show that the probability of this state passing the remaining test is $|\alpha'|^2$, which in total will make the probability $(|\alpha|^2 + |\beta|^2) \cdot |\alpha'|^2 = |\alpha|^2$. Indeed, write

$$|\psi_A\rangle = \sum_{v \in A} c_v |v\rangle$$

where $\sum_{v \in A} c_v = 0$, since $|\psi_A\rangle$ is orthogonal to $|A\rangle$ by construction. Then

$$\begin{aligned} H|\psi_A\rangle = \sum_{v \in A} c_v H|v\rangle = \frac{1}{\sqrt{2^n}} \sum_{v \in A} c_v \sum_{u \in \mathbb{F}_2^n} (-1)^{u \cdot v} |u\rangle &= \frac{1}{\sqrt{2^n}} \sum_{u \in \mathbb{F}_2^n} \left(\sum_{v \in A} c_v (-1)^{u \cdot v} \right) |u\rangle \\ = \frac{1}{\sqrt{2^n}} \sum_{u \in A^\perp} \left(\sum_{v \in A} c_v (-1)^{u \cdot v} \right) |u\rangle &+ \frac{1}{\sqrt{2^n}} \sum_{u \notin A^\perp} \left(\sum_{v \in A} c_v (-1)^{u \cdot v} \right) |u\rangle \\ = \frac{1}{\sqrt{2^n}} \sum_{u \in A^\perp} \left(\sum_{v \in A} c_v \right) |u\rangle &+ \frac{1}{\sqrt{2^n}} \sum_{u \notin A^\perp} \left(\sum_{v \in A} c_v (-1)^{u \cdot v} \right) |u\rangle \\ &= \frac{1}{\sqrt{2^n}} \sum_{u \notin A^\perp} \left(\sum_{v \in A} c_v (-1)^{u \cdot v} \right) |u\rangle \end{aligned}$$

where in the last transition we used $\sum_{v \in A} c_v = 0$. Therefore, $H|\psi_A\rangle$ is orthogonal to $\mathrm{Span}\{|v\rangle : v \in A^\perp\}$, while $H|A\rangle = |A^\perp\rangle$ *is* in $\mathrm{Span}\{|v\rangle : v \in A^\perp\}$, so the final projection $\mathbb{P}_{A^\perp}$ passes with probability $|\alpha'|^2$, as promised. □

A powerful technique described in [2] also allows proving *unconditionally* that using U_A and $U_{A^\perp}$ as black-box classical oracles, it is impossible for a poly-time counterfeiter to obtain a state close to $|A\rangle \otimes |A\rangle$ given a state $|A\rangle$ even with polynomially-many queries to U_A and $U_{A^\perp}$. An important observation is that, if one can consistently start with a state $|A\rangle$ and maps it to $|A\rangle \otimes |A\rangle$ using queries to U_A and $U_{A^\perp}$, then the same procedure would also work on a different banknote $|A'\rangle$ while querying $U_{A'}$ and $U_{(A')^\perp}$. We next examine how two such runs must diverge.

Let $|\psi_t^A\rangle$ be the state of such an algorithm after t oracle queries, so that $|\psi_0^A\rangle = |A\rangle$ and $|\psi_T^A\rangle = |A\rangle \otimes |A\rangle$, where T is the number of queries the algorithm makes, assumed to be $T \leq \mathrm{poly}(n)$. Then under the assumptions, one has $|\langle \psi_0^A | \psi_0^{A'} \rangle| = |\langle A | A' \rangle|$, but

$$|\langle \psi_T^A | \psi_T^{A'} \rangle| = |(\langle A| \otimes \langle A|)(|A'\rangle \otimes |A'\rangle)| = |\langle A | A' \rangle|^2.$$

Hence, if A and A' are not very close, this implies there is at least a constant reduction in the inner product between states of two different counterfeiting runs.

At the same time, the authors prove an exponential upper bound on the reduction in the inner product between two consecutive queries, with the immediate corollary that then T must be exponential in order for the inner product to be reduced the necessary amount. This method has a stronger corollary, in that one can also not achieve an approximate counterfeiting.

Finally, the authors of [2] also suggest a practical instantiation of the classical oracles, under some hardness assumptions tailored for this. Specifically, the suggest drawing random multi-variate polynomials $\mathbb{F}_2^n \to \mathbb{F}_2$ that vanish on A to form a set p_A of polynomials, and separately polynomials that vanish on $A^\perp$ to form $p_{A^\perp}$, such that with high probability we can expect the zero sets of p_A to be exactly A, and likewise for $p_{A^\perp}$ to be exactly $A^\perp$. Then, checking membership in A or $A^\perp$ can be done by substituting in all polynomials in the respective set. However, it is argued that computing a common root to such a set of polynomials might be hard, and hence p_A and $p_{A^\perp}$ can be made public, and effectively function as membership oracles.

2.2 Franchised Quantum Money (FQM)

The idea of franchised quantum money described in [4] is, in some sense, a simplification of the ideas of Aaronson and Christiano's Quantum Money from Hidden Subspaces [2]. Here, a banknote is again a state of the form $|A\rangle$ for some subspace A of $\mathbb{F}_2^n$ of dimension $\dim A = \frac{n}{2}$. Here, each player who wishes to be a verifier, needs to be explicitly "franchised" by the issuing bank. This can be done at any time, even after the money is minted.

In the suggested scheme of [4], each verifier will have knowledge of subspaces V, W of $\mathbb{F}_2^n$ such that $V \subseteq A$ and $W \subseteq A^\perp$ for any specific banknote $|A\rangle$. That is, every verifier has *some* explicit information about the subspace A, but we will make sure it is not enough information for counterfeit for sabotage, even with some (bounded) number of colluding verifiers.

The verification that $|A\rangle$ is valid is similar to the one in [2]. Namely, the verifier checks that $|A\rangle \in \text{Span}\{|v\rangle : v \in W^\perp\}$ (since $W \subseteq A^\perp \Leftrightarrow A \subseteq W^\perp$) by querying a self-made oracle, and similarly checking that $H_2^{\otimes n}|A\rangle \in \text{Span}\{|v\rangle : v \in V^\perp\}$. Any random state is extremely unlikely to pass the verification, and, crucially, since verifiers have independent partial information, a verifier cannot succeed with high probability to produce two quantum states passing another verifier's verification with non-negligible probability.

Initialization. Here, given n, the bank produces a public-private key pair for some public-key signature scheme, sig_{pub} and sig_{priv}, and also randomly draws n private keys of length n for some symmetric key scheme, $k_1, \ldots, k_n$. Each verifier will receive a subset of these keys, to receive the partial information on A allocated to it. These values ($k_1, \ldots, k_n$ and $sig_{priv,pub}$) will be the same for all minted banknotes.

Minting. When the bank issues a new banknote, it first randomly draws a subspace A of $\mathbb{F}_2^n$ of dimension $\frac{n}{2}$, and computes a random basis B_A of A and a random basis $B_{A^\perp}$ of $A^\perp$ (this can be done, for example, by randomly sampling sets of $\frac{n}{2}$ elements from each of A and $A^\perp$, until each set is linearly independent).

Next, the bases B_A and $B_{A^\perp}$ are encrypted using the keys $k_1, \ldots, k_n$, to produce encryptions $c_1, \ldots, c_{\frac{n}{2}}$ for the vectors in the basis B_A, and $c_{\frac{n}{2}+1}, \ldots, c_n$ that are encryptions of vectors in the basis $B_{A^\perp}$. It is emphasized that the ciphertext c_i was produced using the key k_i.

The bank distributes the banknote $|A\rangle$ together with its metadata $(c_1, c_2, \ldots, c_n)$, along with a signature σ of this metadata by the bank's private key. That is,

$$|\$\rangle = (|A\rangle, c_1, c_2, \ldots, c_n, \sigma)$$

Franchising. For franchising a new verifier, the bank simply chooses random subsets $I \subseteq \left[1, \frac{n}{2}\right]$ of size $|I| = \sqrt{n}$ and $J \subseteq \left[\frac{n}{2}+1, n\right]$ of size $|J| = \sqrt{n}$. The bank then sends to the verifier the keys

$\{k_i : i \in I \cup J\}$, so the verifier will be able to read the corresponding subset of the basis together with I, J (so the verifier knows which ones they can read). The bank finally signs $(I, J, \{k_i : i \in I \cup J\})$ and sends that with its signature to the franchised verifier.

Verification was already described, and its correctness is clear (that is, given a valid banknote $|\$\rangle$, a franchised verifier will always accept it with certainty).

Security. So long as there are less than $\frac{\sqrt{n}}{4}$ colluding parties, it is proved in [4] that the scheme is secure against counterfeiting and against sabotage.

2.3 Quantum Lightning (QL)

The main goal of this section is is to give an overview of the proof of the following theorem from [7].

Theorem 1. *Suppose H is a CRHF, then both of the following hold:*

- *Either H is collapsing, or H can be used to build a QL scheme that is infinitely often secure.*
- *Either H is infinitely often collapsing, or H can be used to build a QL scheme that is secure.*

Before we give an overview of the proof of this theorem, we must explain the notions of *collapsing*, and *infinitely often secure.*

Infinitely often secure. Usually in cryptography, security claims are stated in terms of a game, which an efficient adversary can not win, except for with negligible probability. Note that such a statement can be falsified by any adversary that can win the aforementioned game with *non*-negligible probability, i.e. lower bounded by a polynomial in the security parameter *infinitely often.* Such a strong security requirement does not suit "win-win" type of theorems. Referring back to the theorem at hand, ideally, one would like to prove that either H is collapsing, or H can be used to build a secure QL scheme. However, proving such a result has the following barrier. The proof would need to use the adversary $\mathcal{A}$ given by H being not collapsing to construct a *secure* QL scheme. However, the security of such a scheme would be directly related the winning probability of $\mathcal{A}$ in the collapsing security game (which is explained shortly), and for this a promise of *non-negligible* success probability does not suffice, one in fact needs that probability to be lower bounded by a polynomial for *all* parameters. For this purpose we discuss the relaxed notion of security, referred to as *infinitely often secure*, in which a primitive is considered secure if no efficient adversary can win the security game with probability bounded by some inverse polynomial.

Collapsing CRHF. Intuitively, it is easy for an adversary to obtain a *super-position* of pre-images of some (uncontrollable) input, by running H on a uniform superposition of inputs and then measuring the output. A CRHF is called *collapsing*, if the resulting state is *computationally indistinguishable* from a random input. It turns out that this property of collapsing is problematic for constructing a secure QL scheme, see [7] for further details.

Formally, consider the following game between $\mathcal{A}$ and a challenger.

- Challenger holds input bit b.

- Challenger chooses a random key k, which it gives to $\mathcal{A}$.
- $\mathcal{A}$ creates a uniform superposition of all inputs to H, i.e. $\alpha \sum_x |x\rangle$.
- Challenger computes $H(k, \cdot)$ on the uniform superposition of inputs in order to obtain $\sum_x |x, H(k,x)\rangle$.
- Then, challenger:
  - If $b = 0$, measures the $H(k,x)$ register, in order to get an image string y, the state collapses to be a uniform superposition on the pre-images of y.
  - If $b = 1$, measures the entire state and obtains $x, H(k,x)$, to which the state also collapses.
- Challenger sends remaining state to $\mathcal{A}$.
- $\mathcal{A}$ outputs a guess b' for b

$\mathcal{A}$ wins if $b' = b$. We say that a CRHF H is collapsing if any efficient adversary wins the above game with at most negligible probability. We say that H is infinitely often collapsing, if there is no adversary winning the above game with probability lower bounded by $1/p(\lambda)$ for some polynomial p.

Proof overview. The proofs for the secure vs infinitely often secure cases are almost identical, for ease of notation, we focus on the case where there is an adversary $\mathcal{A}$ winning the collapsing game with probability lower bounded by some inverse polynomial $1/p(n)$, i.e. the second bullet.

The paper [7] shows that in fact, in that case, one may as well assume that there is an adversary $\mathcal{A}$ that wins the collapsing game w.p. at least $1 - 2^{-n}$. I.e., the paper shows how to amplify the success probability of the adversary to be near perfect.

With such an adversary in mind, consider the following QL scheme for a given random key k. We note that since a CRHF primitive consists of a family of functions, so will we have a family of pairs $(\mathcal{S}, \mathcal{V})$ of storms and verifiers.

- The specific pair of storm/verifier is determined by a random hash key k from the family of CRHF. Denote the storm/verifier pair by $(\mathcal{S}_k, \mathcal{V}_k)$.
- $\mathcal{S}_k$ runs $\mathcal{A}$ as denoted above up to the point where $\mathcal{A}$ creates a uniform superposition $|\psi\rangle$ over inputs $H_k(\cdot)$, along with with some private state. Then $\mathcal{S}_k$ applies H_k in superposition on $|\psi\rangle$, and measures the resulting state in order to obtain an image string y and a resulting state of a superposition on pre-images of y as defined above, $|\psi_y\rangle$. It outputs $|\psi_y\rangle$ as the bolt.
- Now, $\mathcal{V}_k$ does the following on a supposed bolt $|\psi\rangle$. It first applies H_k in super-position and measures the *output* register, in order to obtain a classical string y, which it keeps as the serial number. Note that only the output register is measured, and so the state $|\psi\rangle$ is kept intact.

  Now, consider two tests T_0, T_1. In test T_0, $\mathcal{V}_k$ continues the execution of $\mathcal{A}$ (denote this part by $\mathcal{A}_1$) on $|\psi\rangle$ and obtains the guess b'. If $b' = 1$, i.e. $\mathcal{A}$ guesses that the challenger measured both input and output registers ($|\psi\rangle$ is a specific input and not a superposition), $\mathcal{V}_k$ rejects the bolt. Otherwise, if $\mathcal{A}$ guesses 0, i.e. $\mathcal{A}$ thinks that $|\psi\rangle$ is a superposition on pre-images of y, we uncompute $\mathcal{A}_1$ and output the resulting state along with serial number y. One can think of this test as corresponding to the case where the challenger picked $b = 0$.

Test T_1, is similar to T_0, with the following change. $\mathcal{V}_k$ "measures" the input x registers before sending them to $\mathcal{A}_1$. This is not a real measurements, and is done by copying the input registers x into a private register. It then gives $\mathcal{A}_1$ this copied input register along with the original output register from computing H_k on $|\psi\rangle$ and executes $\mathcal{A}_1$ on it to obtain a guess b'. We now flip the rejection condition, if $b' = 0$, we abort and reject. Otherwise, we uncompute $\mathcal{A}_1$ and output the resulting state $|\psi'\rangle$ along with serial number y.

With these tests in mind, $\mathcal{V}_k$ chooses a random but c and performs T_c, and answers accordingly.

Correctness. If $|\psi\rangle$ is a valid bolt, i.e. a uniform superposition of some pre-images of y, then we note that T_0 corresponds exactly to the case where $b = 0$ in the collapsing game described above, in which case $\mathcal{A}_1$ wins with overwhelming probability, and thus the bolt is accepted. Similarly, T_1 is the $b = 1$ case in which also $\mathcal{A}_1$ wins with overwhelming probability. Thus with very high probability, the bolt is accepted. The paper [7] also shows why the resulting state $|\psi'_y\rangle$ is very close to the original state $|\psi_y\rangle$.

Security. The security proof is more involved, and we refer the reader to [7] for further details.

References

[1] Scott Aaronson. Quantum copy-protection and quantum money. In *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE, jul 2009. doi:10.1109/ccc.2009.42.

[2] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '12, page 41–60, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2213977.2213983.

[3] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021*, pages 526–555, Cham, 2021. Springer International Publishing.

[4] Bhaskar Roberts and Mark Zhandry. Franchised quantum money, 2021. doi:10.48550/ARXIV.2110.09733.

[5] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, STOC '90, page 387–394, New York, NY, USA, 1990. Association for Computing Machinery. doi:10.1145/100216.100269.

[6] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, jan 1983. doi:10.1145/1008908.1008920.

[7] Mark Zhandry. Quantum lightning never strikes the same state twice. or: Quantum money from cryptographic assumptions. *J. Cryptol.*, 34(1):6, 2021.