Quantum Money from Abelian Group Actions

Mark Zhandry
mzhandry@gmail.com
NTT Research

Abstract

We give a candidate construction of public key quantum money, and even a strengthened version called quantum lightning, from abelian group actions, which can in turn be constructed from suitable isogenies over elliptic curves. We prove our scheme is secure under new but plausible assumptions on suitable group actions.

1 Introduction

Quantum money, first envisioned by Wiesner [Wie83], is a system of money where banknotes are quantum states. By the no-cloning theorem, such banknotes cannot be copied, leading to un-counterfeitable currency. A critical feature of quantum money, identified by [Aar09], is *public verification*, allowing anyone to verify while only the mint can create new banknotes. Such public key quantum money is an important central object in the study of quantum protocols, but unfortunately convincing constructions have remained elusive. See Section 1.4 for a more thorough discussion of prior work in the area.

This Work. We construct public key quantum money from abelian group actions, which can be instantiated by suitable isogenies over ordinary elliptic curves. Group actions, and the isogenies they abstract, are one of the leading contenders for post-quantum secure cryptosystems. Our construction could plausibly even be quantum lightning, a strengthening of quantum money with additional applications. Our construction is arguably the first time group actions have been used to solve a classically-impossible cryptographic task that could not already be solved using other standard tools like LWE. Our construction is sketched in Section 1.1 below, and given in detail in Section 3.

While our main construction can be instantiated on a clean abelian group action — often referred to as an "effective group action" (EGA) — many isogenie-based group actions diverge from this convenient abstraction. We therefore provide an alternative candidate scheme which can be instantiated on so-called "restricted effective group actions" (REGAs); see Section 4 for details.

We prove the quantum lightning security of our protocols under the group action analogs of a "knowledge of exponent" assumption and a strengthening of the discrete log assumption. The ordinary discrete logarithm assumption underpins all of cryptography from groups and group actions, and our generalization is a slight strengthening which essentially allows the adversary a constant number of CDH queries. The knowledge of exponent assumption for group actions — which we call "knowledge of group element" — is a new assumption that needs further study.

We conclude in Section 5 with a discussion of possible generalizations and relation to approaches for building quantum money from LWE.

1.1 Our Construction

Abelian Group Actions. We will use additive group notation for abelian groups. An abelian group action consists of an abelian group $\mathbb{G}$ and a set $\mathcal{X}$, such that $\mathbb{G}$ "acts" on $\mathcal{X}$ through the binary relation $* : \mathbb{G} \times \mathcal{X} \to \mathcal{X}$ with the property that $g * (h * x) = (g + h) * x$ for all $g, h \in \mathbb{G}, x \in \mathcal{X}$. We will also assume a *regular* group action, which means that for every $x \in \mathcal{X}$, the map $g \mapsto g * x$ is a bijection.

The main group actions used in cryptography are those arising from isogenies over elliptic curves. For example, see [Cou06, RS06, CLM$^+$18, BKV19, DFK$^+$23]. Group action cryptosystems rely at a minimum on the assumed hardness of discrete logarithms: given $x, y = g * x \in \mathcal{X}$, finding g. For isogeny-based actions, this corresponds to the hard problem of computing isogenies between elliptic curves. Other hard problems are possible, such as analogs of computational/decisional Diffie-Hellman, and more.

The QFT. Our quantum money scheme will utilize the quantum Fourier transform (QFT) over general abelian groups. This is a quantum procedure that maps

$$|g\rangle \mapsto \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{h \in \mathbb{G}} \chi(g, h) |h\rangle \ \ .$$

Here, χ is some potentially complex phase term. In the case of $\mathbb{G}$ being the additive group $\mathbb{Z}_N$, $\chi(g, h)$ is defined as $e^{i2\pi gh/N}$, with a slightly more complicated definition for non-cyclic groups[1]. The main property we need from χ (besides making the QFT unitary) is that it is *bilinear*, in the sense that $\chi(g, h_1 + h_2) = \chi(g, h_1) \cdot \chi(g, h_2)$. It is also symmetric: $\chi(g, h) = \chi(h, g)$.

Our Quantum Money Scheme. Our quantum money scheme is as follows; see Section 3 for additional details.

- Gen: initialize a register in the state $\frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g \in \mathbb{G}} |g\rangle$, which can be computed by applying the QFT to $|0\rangle$. Let $x \in \mathcal{X}$ be arbitrary. Then by computing the group action in superposition, compute $\frac{1}{\sqrt{|\mathbb{G}|}} \sum_{h \in \mathbb{G}} |g\rangle |g * x\rangle$. Next, apply the QFT over $\mathbb{G}$ to the first register. The result is:

  $$\frac{1}{|\mathbb{G}|} \sum_{g,h \in \mathbb{G}} \chi(g, h) |h\rangle |g * x\rangle = \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{h} |h\rangle |\mathbb{G}^h * x\rangle$$

  Here, $|\mathbb{G}^h * x\rangle$ is the state $\frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g \in \mathbb{G}} \chi(g, h) |g * x\rangle$. Note that $|\mathbb{G}^h * x\rangle$ is, up to an overall phase, independent of x.

  Now measure h, in which case the second register collapses to $|\mathbb{G}^h * x\rangle$. Output h as the serial number, and $|\mathbb{G}^h * x\rangle$ as the money state.

- To verify a banknote \$, choose a random $u \in \mathbb{G}$, and initialize a new qubit with $(|0\rangle + |1\rangle)/\sqrt{2}$. Then apply the controlled group action $|b, y\rangle \mapsto \begin{cases} |0, y\rangle & \text{if } b = 0 \\ |1, u * y\rangle & \text{if } b = 1 \end{cases}$. If \$ is the honest

[1] Remember that the group aoperation is +, so gh in the exponent is not the group operation, but instead multiplication in the ring $\mathbb{Z}_N$.

banknote state, then the state of the system becomes:

$$\frac{1}{\sqrt{2}}\left(|0\rangle + \chi(u,h)^{-1}|1\rangle\right)|\mathbb{G}^h * x\rangle$$

We can then measure the first qubit in the basis containing $|0\rangle + \chi(u,h)^{-1}|1\rangle$, which will accept with probability 1 for honest banknote states. We can repeat this process λ times, and accept only if all trials accept. It is possible to show that if all λ trials accept, the result state is $2^{\Theta(\lambda)}$-close to the honest banknote state.

An instantiation using REGAs. For some isogeny-based group actions such as CSIDH [CLM+18], the operation $*$ is only efficiently computable for a very small set $S \subseteq \mathbb{G}$ of group elements. Such group actions are called "restricted effective group actions" (REGAs) [ADMP20]. Above, however, we see that we need to compute the group action on all possible elements in $\mathbb{G}$, both for minting and for verification. We therefore give a variant of the construction above which only uses the ability to compute $*$ for elements in S. We show that we are still able to sample $|\mathbb{G}^h * x\rangle$, but now the serial number has the form $\mathbf{A}^T h + \mathbf{e} \bmod N$ for a known matrix $\mathbf{A}$ and a "small" $e \in \mathbb{Z}^n$ [2]. Under plausible assumptions, the serial number actually hides h [3]. We nevertheless show that we can use such a noisy serial number for verification. For details, see Section 4. The security of our alternate scheme is essentially equivalent to the main scheme.

1.2 The security of our scheme

We do not know how to base the security of our schemes on any standard assumptions on isogenies. However, we prove security under non-standard but plausible assumptions, including a knowledge assumption. Specifically, we show how to adapt the security proof of quantum money over walkable invariants [LMZ23] to our setting, though several important things need to change. Importantly, we note that our scheme *cannot* be seen as an instance of walkable invariants. Here we sketch the proof for our main construction, which is given in full in Section 3.3. The security of our modified construction over REGAs is essentially equivalent, as discussed in Section 4.

We first define a "knowledge of group element" assumption (KGEA), which roughly states that any adversary which can output a set element $y \in \mathcal{X}$ must "know" a group element g such that $y = g * x$. Here, x is some fixed set element that is provided to all parties. The intuition is that, if $\mathcal{X}$ is a sparse set, then perhaps the only way to generate new set elements is to actually use the group action operation. Slightly more formally, we consider a quantum adversary that outputs a set element y. We assume that the adversary performed no measurements besides measuring y to get the final output, meaning that the adversary might have some remaining quantum state $|\psi_y\rangle$. Then the knowledge of group element assumption states that given $|\psi_y\rangle$ and y, it is possible to efficiently compute g such that $y = g * x$. This assumption is analogous to the "knowledge of path" assumption defined by [LMZ23] for walkable invariants.

[2] Here, we are interpreting h a vector in $\mathbb{Z}_N^n$ for some n, N, which is possible since $\mathbb{G}$ is abelian.

[3] This is the Learning with Errors (LWE) problem [Reg05] which is widely believed to be hard for *random* $\mathbf{A}$. In our case, $\mathbf{A}$ is a fixed matrix that depends on the group action, and LWE may or may not be hard for this $\mathbf{A}$. However, a variant of Regev's quantum reduction between LWE and Short Integer Solution (SIS) [Reg05], outlined by [YZ22], shows that if LWE can be solved relative to $\mathbf{A}$, then SIS can be solved for $\mathbf{A}$ as well. It is straightforward to adapt this reduction to solve the Inhomogenous SIS (ISIS) problem, which then allows for computing the group action for all of $\mathbb{G}$. In this case we would have a clean group action and would not need this alternate construction.

We now consider a quantum lightning adversary, which produces two banknotes with the same serial number. We will first purify the adversary, so that it makes no measurements. The state of the adversary can then be written as

$$\sum_h \alpha_h |\phi_h\rangle |\mathbb{G}^h * x\rangle |\mathbb{G}^h * x\rangle = \frac{1}{|\mathbb{G}|} \sum_{h,g_1,g_2} \alpha_h \chi(h, g_1 + g_2) |\phi_h\rangle |g_1 * x\rangle |g_2 * x\rangle \ ,$$

where $\sum_h |\alpha_h|^2 = 1$, and $|\phi_h\rangle$ is any state of the adversary left over after outputting the two banknotes. Now consider the algorithm which produces this state and then measures the final register, to obtain $g_2 * x$, with the remaining state collapsing to

$$|\psi_{g_2 * x}\rangle := \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g_1,h} \alpha_h \chi(h, g_1 + g_2) |\phi_h\rangle |g_1 * x\rangle$$

Applying KGEA to this adversary, there is an algorithm E that can compute g_2 from $g_2 * x$ and $|\psi_{g_2 * x}\rangle$. We would like to use E and $|\psi_{g_2 * x}\rangle$ to solve the discrete log assumption, reaching a contradiction. While E can solve the discrete log of $g_2 * x$, E additionally needs $|\psi_{g_2} * x_\lambda\rangle$, which is correlated with $g_2 * x$. Therefore, E does not immediately represent a discrete log adversary. Inspired by [LMZ23], we want to apply E to a fresh discrete log challenge $g * x$ to recover g. In [LMZ23], simply swapping out the measured value for a fresh challenge works, as in their setting the measured output is independent of the remaining state of the adversary; see Section 1.4 for a brief explanation of their proof. The problem for us is that $g_2 * x$ is correlated with $|\psi_{g_2 * x}\rangle$, and E may not work when given $g * x$ and $|\psi_{g_2 * x}\rangle$ for $g \neq g_2$; perhaps E only works if given the correct $g_2 * x$ for $|\psi_{g_2 * x}\rangle$.

Our solution is to, given $g * x$, transform $|\psi_{g_2 * x}\rangle$ into $|\psi_{g * x}\rangle$; then we can apply E to $g * x$ and $|\psi_{g * x}\rangle$ to recover g. We first note that we can, using our knowledge of g_2 derived from E, apply the map $y \mapsto g_2 * y$ to the last register in $|\psi_{g_2 * x}\rangle$ to obtain

$$\frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g_1,h} \alpha_h \chi(h, g_1 + g_2) |\phi_h\rangle |(g_1 + g_2) * x\rangle = \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g_2',h} \alpha_h \chi(h, g_2') |\phi_h\rangle |g_2' * x\rangle = |\psi_{0 * x}\rangle = |\psi_x\rangle$$

where we used the change of variables $g_1 + g_2 \mapsto g_2'$. Next, we want to move $|\psi_x\rangle$ to $|\psi_{g * x}\rangle$. We could do this if we knew g as well, but this is what we are trying to compute!

Instead we rely on a strengthening of the discrete log assumption, where we give the adversary a call to a computational Diffie-Hellman oracle in order to help it solve the discrete log. That is, we allow the adversary to query on a set element y, and obtain $(-g) * y$. We moreover allow this call to happen in superposition. But we only allow a single such query, or possibly two queries, depending on exactly how we model the query[4]. Then the adversary is given $g * x$ and must compute g. If we apply this oracle to the last register in $|\psi_x\rangle$, we get the state

$$\frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g_2',h} \alpha_h \chi(h, g_2') |\phi_h\rangle |(g_2' - g) * x\rangle = \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{g_2'',h} \alpha_h \chi(h, g + g_2'') |\phi_h\rangle |g_2'' * x\rangle = |\psi_{g * x}\rangle$$

[4] The number of queries depends on whether $(-g) * y$ replaces the value y (the so-called "minimal" oracle [KKVB02]), or is XORed into a supplied response response register (the "standard" oracle). The former requires only a single query, while the latter requires two queries to simulate the former: the first to the function $y \mapsto (-g) * y$, and the second to the function $y \mapsto g * y$.

as desired, where we used the change of variables $g_2' - g \mapsto g_2''$. Now we can apply E to $g * x$ and $|\psi_{g*x}\rangle$ to recover g, thus solving discrete log and reaching a contradiction.

The lingering question is whether our new assumptions hold. For our strengthened discrete log assumption, it appears hard to use one or two queries to such an oracle to actually recover g. One possibility is to view the oracle as a CDH oracle, which computes $(r - s) * x$ from $r * x, s * x$ [5], and then use the quantum equivalence of CDH and discrete log for group actions [GPSV21, MZ22]. In our case, we only give a very limited CDH oracle which only works if $s * x$ was set to $g * x$. But even if we allowed the full generality of a CDH oracle and tried to apply [GPSV21, MZ22], we will fail to compute the discrete log, since [GPSV21, MZ22] require far more than two queries to the oracle. We therefore conjecture our strengthened discrete log problem is hard.

We note that isogenies over elliptic curves will typically allow for sampling certain elements in $\mathcal{X}$ without directly computing them via applying $*$ to other elements. Specifically, it is possible sample elements with small discriminant using the complex multiplication method. This means KGEA as described above is technically false on known isogenies. Nevertheless, such directly sampled elements bare no obvious relation to each other, so it is unclear how to use them to actually break the security of our scheme. Toward rectifying this issue, we give more refined assumptions that avoid this issue while still allowing for proving security. See Section 3.4 for details.

1.3 Further Discussion

In Section 5, we generalize group actions to *quantum* group actions, which replace classical set elements with quantum states, but otherwise behave mostly the same as standard group actions. We give a simple quantum group action based on the Learning with Errors (LWE) problem [Reg05], where we can actually prove that the discrete log problem is hard under LWE. Despite this promising result, we expect that the LWE-based quantum group action will be of limited use. In particular, if we instantiate our quantum money construction over this group, the construction is *insecure*. The reason is that, in this group action, it is impossible to recognize the quantum states of the set. Our security proof crucially relies on such recognition, since it allows us to characterize states accepted by the verifier. Moreover, without recognition, there is an attack: it is possible to fool the verifier with dishonest banknotes that are different from the honest ones and moreover are clonable, thereby breaking security.

Interestingly, we explain that this failed instantiation is actually *equivalent* to a folklore approach toward building quantum money from lattices, which has been more-or-less shown impossible to make secure [LZ19, LMZ23]. The key missing piece in getting the folklore approach to work has been how to efficiently verify honest banknotes — if such verification were possible, the scheme could be readily proven secure. Under our equivalence, this missing piece exactly maps to the problem of recognizing set elements in our quantum group action. For details, see Section 5. We believe this adds to the confidence of our proposal, since in group actions based on isogenies it is possible to recognize set elements, presumably without otherwise compromising hardness.

[5] CDH would usually be defined as computing $(r + s) * x$, but it is basically equivalent to consider the case where it computes $(r - s) * x$. In the case where we allow $\Omega(|\mathbb{G}|)$ CDH queries, we can use addition queries to implement subtraction queries and vice versa.

1.4 Related Work

Public key quantum money. In Wiesner's original scheme, the mint is required to verify banknotes, meaning the mint must be involved in any transaction. The involvement of the mint also leads to potential attacks [Lut10]. Some partial solutions have been proposed, e.g. [BS20a, RZ21]. The dream solution, however, is known as *public key* quantum money [Aar09]. Here, anyone can verify the banknote, while only the mint can create them.

Unlike Wiesner's scheme which is well-understood, secure public key quantum money has remained elusive. While there have been many proposals for public key quantum money [Aar09, AC12, FGH+12, Kan18, Zha19, KSS21, KLS22, LMZ23], they mostly either (1) have been subsequently broken (e.g. [Aar09, AC12, Zha19, KLS22] which were broken by [LAF+10, CPDDF+19, Rob21, LMZ23]), or (2) rely on new cryptographic building blocks that have received little attention from the cryptographic community (e.g. [FGH+12, Kan18, KSS21] from problems on knots or quaternion algebras). The two exceptions are:

- Building on a suggestion of [BDS16], [Zha19] proved that quantum money can be built from post-quantum indistinguishability obfuscation (iO). While iO has received considerable attention and even has a convincing *pre-quantum* instaniation [JLS21], the post-quantum study of iO has been much less thorough. While some post-quantum proposals have been made [GGH15, BGMZ18, BDGM20, WW21], their post-quantum hardness is not well-understood.

- [LMZ23] construct quantum money from isogenies over super-singular elliptic curves. However, there is a crucial missing piece to their proposal, namely generating uniform superpositions over super-singular curves, which is currently unknown how to do. This is closely related to the major open question of obliviously sampling super-singular elliptic curves.

In light of the above, the existence of public key quantum money is largely considered open.

Cryptography from group actions and isogenies. Isogenies were first proposed for use in post-quantum cryptography by Couveignes [Cou06] and Rostovtsev and Stolbunov [RS06]. Isogenies give a Diffie-Hellman-like structure, but importantly are immune to Shor's algorithm for discrete logarithms [Sho94] due to a more restricted structure. This restricted structure, while helping preserve security against quantum attacks, also makes the design of cryptosystems based on them more complex. Thus, significant effort has gone into building secure classical cryptosystems from isogenies and understanding their post-quantum security (e.g. [CJS14, DJP14, CLM+18, BKV19, CK20, DM20, Pei20, BS20b, ADMP20, AMR22, MZ22, MM22, CD23, BGZ23, Rob23]).

Certain isogenies such as the original proposals of [Cou06, RS06] as well as CSIDH and its variants [CLM+18, DFK+23] can be abstracted as abelian group actions. However, many other isogenies (such as SIDH [DJP14] and OSIDH [CK20]) cannot be abstracted as abelian group actions. Even among abelian group actions, we must distinguish between "effective group actions" (EGAs) and *restricted* EGAs (REGAs). The former satisfies the notion of a clean group action, whereas in the latter, the group action can only be efficiently computed for a certain small set of group elements. CSIDH could plausibly be a EGA at certain concrete security parameters, though asymptotically it only achieves quasi-polynomial security[6]. Our alternate construction also works on REGAs, which

[6] With the state-of-the-art, evaluating CSIDH as an EGA would require time approximately $2^{\sqrt[3]{n}}$ on a quantum

can plausibly be instantiated even asymptotically by CSIDH using a quantum computer[7].

While some non-isogeny abelian group actions have been proposed (e.g. [Sti05]), currently all such examples have been broken (e.g. [Shp08]). For this reason, group actions are largely considered synonymous with isogenies, though this may change if more secure group actions are found.

The vast majority of the isogeny and group action literature has focused on post-quantum cryptography — classical protocols that are immune to quantum attacks. To the best of our knowledge, only two prior works have used isogenies/group actions to build quantum protocols for tasks that are *impossible* classically. The first is [AMR22], who build a proof of quantumness [BCM+18]. We note that proofs of quantumness can also be achieved under several "standard" cryptographic tools, such as LWE [BCM+18] or certain assumptions on hash functions [YZ22]. In contrast, no prior quantum money protocol could be based on similar standard building blocks. We also note that [AMR22] currently has no known asymptotic instantiation with better-than-quasi-polynomial security, as it requires a clean group action (EGA). The second quantum protocol based on isogenies is that of [LMZ23], who build quantum money from walkable invariants, and propose an instantiation using isogenies over super-singular elliptic curves. However, such isogenies cannot be described as abelian group actions, and even more importantly their proposal is incomplete, as discussed above. Thus, ours is arguably the first application of group actions or isogenies to obtain classically impossible tasks that could not already be achieved under standard tools.

Relation to [LMZ23]. Aside from using isogenies, our work has strong conceptual similarities to [LMZ23], though also crucial differences that allow us to specify a complete protocol. Here, we give a brief overview of the similarities and differences.

The walkable invariant framework of [LMZ23] is very general, but here we describe a special case of it that would apply to certain group actions, in order to illustrate the differences with our scheme. Consider a group action that is *not* regular, so that the set $\mathcal{X}$ is partitioned into many distinct orbits. For x, y in the same orbit there will exist a unique g such that $y = g * x$, but for x, y in different orbits, there will not exist any group element mapping between them. We will also assume the ability to generate a uniform superposition over $\mathcal{X}$. We finally assume an "invariant", a unique label for each orbit which can be efficiently computed from any element in the orbit.

The minting process generates the uniform superposition over $\mathcal{X}$, and then measure the invariant, which becomes the serial number. The state then collapses to a uniform superposition over a single orbit, which becomes the banknote. This superposition can then be verified as follows. First check that the banknote has support on the right orbit by re-computing the invariant. Then check that the state is in uniform superposition by checking that the state is preserved under action by random group elements; this is accomplished using an analog of the swap test. [LMZ23] prove the security of their scheme under the certain assumptions which, when mapped to the group action setting above, correspond to the discrete log assumption and a knowledge assumption very similar to ours.

Unfortunately, there are no known instantiations of suitable group actions for their scheme. They propose using the set of ordinary elliptic curves as the set, the number of points on the curve as the invariant, and orbits being sets of curves with the same number of points. Isogenies between

computer, while the best quantum attack is time $2^{\sqrt{n}}$. For a thorough discussion, see [Pan23]. By setting $n = \log^3(\lambda)$, one gets polynomial-time evaluation and the best attack taking time $\lambda^{\sqrt{\log(\lambda)}}$.

[7] In order for CSIDH to be a REGA, one needs to compute the structure of the group. While this is hard classically, it is easy with a quantum computer using Shor's algorithm [Sho94]. Since we always assume a quantum computer in this work, we can therefore treat CSIDH as a REGA.

curves are then the action[8], which do not change the number of points on the curve. The problem is that in general curves, it is not possible to efficiently compute the action, since the degree will be too high. The action *can* be computed on smooth-order curves, but these are rare and there is no known way to compute a uniform superposition over such smooth-order curves. For reasons we will not get into here, [LMZ23] propose using instead supersingular curves with non-smooth order, but again these are rare and there is no known way to generate a uniform superposition over such curves.

We resolve the issues with instantiating [LMZ23], without needing the ability to compute uniform superpositions over the set. Our key insight is that, if we can compute the group action efficiently (say because we are in an orbit of smooth-order elliptic curves), then this is enough to sample states that *are* uniform over a given orbit, except for certain phase terms: namely the states $|\mathbb{G}^h * x\rangle$ for uniform h. Then, rather than the serial number indicating which orbit we are in (which is now useless since we are in a single orbit), the serial number is a description of the phase terms, namely h. Despite these changes we are able to nevertheless prove security under similar assumptions as in [LMZ23] when specialized to group actions.

These changes, however, make adapting the security proof of [LMZ23] to our setting somewhat non-trivial. When specializing walkable invariants to suitable group actions, [LMZ23] make two assumptions that essentially correspond to the discrete log and knowledge of group element assumptions. The reason [LMZ23] can reduce to the plain discrete log assumption is that, when they measure a valid banknote, they get a uniform element in the corresponding orbit, independent of any side information the adversary has. In slightly more detail, their banknotes are uniform superpositions, so when they measure both banknotes from a quantum lightning adversary, they obtain two independent elements in the same orbit. Since the points are uniform and independent of the adversary's state, the discrete log assumption applies. But this directly contradicts the knowledge assumption, which states that any adversary that outputs two points in the same orbit must know the discrete log between them. In our setting, because we have the serial number be the phase instead of the orbit, the banknotes can be entangled with each other through the phase term, and this breaks the straightforward adaptation of the proof from [LMZ23]. Fortunately, we are able to give a proof under our strengthened discrete log assumption.

Acknowledgments

We thank Hart Montgomery for many helpful discussions about isogenies.

2 Preliminaries

Here we give our notation and definitions. We assume the reader is familiar with the basics of quantum computation.

2.1 Quantum Fourier Transform over Abelian Groups

Let $\mathbb{G}$ be an abelian group, which we will denote additively. We here define our notation for the quantum Fourier transform over $\mathbb{G}$. Write $\mathbb{G} = \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \mathbb{Z}_{n_k}$ where $\mathbb{Z}_{n_j}$ are the additive cyclic

[8] It is not a proper group action since different orbits will be acted on by different groups.

groups on n_j elements, and associate elements $g \in \mathbb{G}$ with tuples $g = (g_1, \dots, g_k)$ where $g_j \in \mathbb{Z}_{n_j}$. Then define $\chi : \mathbb{G}^2 \to \mathbb{C}$ by

$$\chi_{\mathbb{G}}(g,h) = \prod_{j=1}^{k} e^{i2\pi g_j h_j / n_j}$$

Observe the following:

$$\chi_{\mathbb{G}}(g,h) = \chi_{\mathbb{G}}(h,g) \qquad \chi_{\mathbb{G}}(g_1+g_2,h) = \chi_{\mathbb{G}}(g_1,h) \times \chi_{\mathbb{G}}(g_2,h)$$

$$\chi_{\mathbb{G}}(-g,h) = \chi_{\mathbb{G}}(g,h)^{-1} \qquad \sum_{g\in\mathbb{G}} \chi_{\mathbb{G}}(g,h) = \begin{cases} |\mathbb{G}| & \text{if } h = 1_{\mathbb{G}} \\ 0 & \text{if } h \neq 1_{\mathbb{G}} \end{cases}$$

The quantum Fourier transform (QFT) over $\mathbb{G}$ is the unitary $\mathsf{QFT}_{\mathbb{G}}$ defined as

$$\mathsf{QFT}_{\mathbb{G}}|g\rangle = \frac{1}{\sqrt{|\mathbb{G}|}} \sum_{h\in\mathbb{G}} \chi(g,h)|h\rangle \ .$$

Observe that $\mathsf{QFT}_{\mathbb{G}} = \mathsf{QFT}_{\mathbb{Z}_{n_1}} \otimes \cdots \otimes \mathsf{QFT}_{\mathbb{Z}_{n_k}}$. Therefore, since the standard QFT corresponds to $\mathsf{QFT}_{\mathbb{Z}_{n_j}}$ and can be implemented efficiently, so can $\mathsf{QFT}_{\mathbb{G}}$.

From this point on, we will only work with a single group, so we will drop the sub-script and simply write $\chi(g,h)$, QFT, etc.

2.2 Quantum Money and Quantum Lightning

Here we define quantum money and quantum lightning. In the case of quantum money, we focus on *mini-schemes* [AC12], which are essentially the setting where there is only ever a single valid banknote produced by the mint. As shown in [AC12], such mini-schemes can be upgraded generically to full quantum money schemes using digital signatures.

Syntax. Both quantum money mini-schemes and quantum lightning share the same syntax:

- $\mathsf{Gen}(1^\lambda)$ is a quantum polynomial-time (QPT) algorithm that takes as input the security parameter (written in unary) which samples a classical serial number σ and quantum banknote \$.

- $\mathsf{Ver}(\sigma, \$)$ takes as input the serial number and a supposed banknote, and either accepts or rejects, denoted by 1 and 0 respectively.

Correctness. Both quantum money mini-schemes and quantum lightning have the same correctness requirement, namely that valid banknotes produced by Gen are accepted by Ver. Concretely, there exists a negligible function $\mathsf{negl}(\lambda)$ such that

$$\Pr[\mathsf{Ver}(\sigma,\$) = 1 : (\sigma,\$) \leftarrow \mathsf{Gen}(1^\lambda)] \geq 1 - \mathsf{negl}(\lambda) \ .$$

Security. We now discuss the security requirements, which differ between quantum money and quantum lightning.

Definition 2.1. Consider a QPT adversary $\mathcal{A}$, which takes as input a serial number σ and banknote $\$$, and outputs two potentially entangled states $\$_1, \$_2$, which it tries to pass off as two banknnotes. $(\mathsf{Gen}, \mathsf{Ver})$ is a secure *quantum money mini-scheme* if, for all such $\mathcal{A}$, there exists a negligible $\mathsf{negl}(\lambda)$ such that the following holds:

$$\Pr\left[\mathsf{Ver}(\sigma, \$_1) = \mathsf{Ver}(\sigma, \$_2) = 1 : \substack{(\sigma,\$)\leftarrow\mathsf{Gen}(1^\lambda)\\ (\$_1,\$_2)\leftarrow\mathcal{A}(\sigma,\$)}\right] \leq \mathsf{negl}(\lambda) \ .$$

Definition 2.2. Consider a QPT adversary $\mathcal{B}$, which takes as input the security parameter λ, and outputs a serial number σ and two potentially entangled states $\$_1, \$_2$, which it tries to pass off as two banknnotes. $(\mathsf{Gen}, \mathsf{Ver})$ is a secure *quantum lightning* scheme if, for all such $\mathcal{B}$, there exists a negligible $\mathsf{negl}(\lambda)$ such that the following holds:

$$\Pr\left[\mathsf{Ver}(\sigma, \$_1) = \mathsf{Ver}(\sigma, \$_2) = 1 : (\sigma, \$_1, \$_2) \leftarrow \mathcal{B}(1^\lambda)\right] \leq \mathsf{negl}(\lambda) \ .$$

Quantum lightning trivially implies quantum money: any quantum money adversary $\mathcal{A}$ can be converted into a quantum lightning adversary $\mathcal{B}$ by having $\mathcal{B}$ run both Gen and $\mathcal{A}$. But quantum lightning is potentially stronger, as it means that even if the serial number is chosen adversarially, it remains hard to devise two valid banknotes. This in particular means there is some security against the mint, which yields a number of additional applications, as discussed by [Zha19].

Remark 2.3. One limitation of quantum lightning as defined above is that it cannot hold against non-uniform attackers with quantum advice, as such attackers could have $\sigma, \$_1, \$_2$ hard-coded in their advice. The situation is analogous to the case of collision resistance, where unkeyed hash functions cannot be secure against non-uniform attackers. This limitation be remedied by either insisting on only uniform attackers or attackers with classical advice. Alternatively, one can work in a trusted setup model, where a trusted third party generates a common reference string that is then inputted into $\mathsf{Gen}, \mathsf{Ver}$. A third option is to use the "human ignorance" approach [Rog06], in which we would formalize security proofs as explicitly transforming a quantum lightning adversary into an adversary for some other task, the latter adversary existing but is presumably unknown to human knowledge. We will largely ignore these issues throughout this work, but occasionally make brief remarks about what the various approaches would look like.

2.3 Group Actions

An (abelian) group action consists of a family of (abelian) groups $\mathbb{G} = (\mathbb{G}_\lambda)_\lambda$ (written additively), a family of sets $\mathcal{X} = (\mathcal{X}_\lambda)_\lambda$, and a binary operation $* : \mathbb{G}_\lambda \times \mathcal{X}_\lambda \to \mathcal{X}_\lambda$ satisfying the following properties:

- **Identity:** If $0 \in \mathbb{G}_\lambda$ is the identity element, then $0 * x = x$ for any $x \in \mathcal{X}_\lambda$.
- **Compatibility:** For all $g, h \in \mathbb{G}_\lambda$ and $x \in \mathcal{X}_\lambda$, $(g + h) * x = g * (h * x)$.

We will additionally require the following properties:

- **Efficiently computable:** There is a QPT procedure $\mathsf{Construct}$ which, on input 1^λ, outputs a description of $\mathbb{G}_\lambda$ and an element $x_\lambda \in \mathcal{X}_\lambda$. The operation $*$ is also computable by a QPT algorithm.

- **Efficiently Recognizable:** There is a QPT procedure Recog which recognizes elements in $\mathcal{X}_\lambda$. That is, for any λ and any string y (not necessarily in $\mathcal{X}_\lambda$), $\mathsf{Recog}(1^\lambda, y)$ accepts y with overwhelming probability if $y \in \mathcal{X}_\lambda$, and rejects with overwhelming probability if $y \notin \mathcal{X}_\lambda$.

- **Regular:** For every $y \in \mathcal{X}_\lambda$, there is exactly one $g \in \mathbb{G}_\lambda$ such that $y = g * x_\lambda$.

Cryptographic group actions. At a minimum, a cryptographically useful group action will satisfy the following discrete log assumption:

Assumption 2.4. The *discrete log assumption* (DLog) holds on a group action $(\mathbb{G}, \mathcal{X}, *)$ if, for all QPT adversaries $\mathcal{A}$, there exists a negligible λ such that

$$\Pr[\mathcal{A}(g * x_\lambda) = g : g \leftarrow \mathbb{G}_\lambda] \leq \mathsf{negl}(\lambda) \ .$$

We will always assume a group action that satisfies the DLog assumption, and this assumption provides intuition for what may and may not be hard on a group action. However, the discrete log assumption will not be sufficient for justifying the security of our construction.

Remark 2.5. For simplicity, we model the group actions as being deterministically computed from the security parameter. We could alternatively imagine the group actions being probabilistic, in which case they would be set up by some probabilistic procedure. The parameters would then be part of a common reference string that is supplied to all parties, including the adversary.

3 Our Quantum Lightning Scheme

Here, we give our basic quantum lightning construction, which assumes a cryptographic group action.

Construction 3.1. Let $\mathsf{Gen}, \mathsf{Ver}$ be the following QPT procedures:

- $\mathsf{Gen}(1^\lambda)$: Initialize quantum registers $\mathcal{S}$ (for serial number) and $\mathcal{M}$ (for money) to states $|0\rangle_\mathcal{S}$ and $|0\rangle_\mathcal{M}$, respectively. Then do the following:
  - Apply $\mathsf{QFT}_{\mathbb{G}_\lambda}$ to $\mathcal{S}$, yielding the joint state $\frac{1}{\sqrt{|\mathbb{G}_\lambda|}} \sum_{g \in \mathbb{G}_\lambda} |g\rangle_\mathcal{S} |0\rangle_\mathcal{M}$.
  - Apply in superposition the map $|g\rangle_\mathcal{S}|y\rangle_\mathcal{M} \mapsto |g\rangle_\mathcal{S}|y \oplus (g * x_\lambda)\rangle_\mathcal{M}$. The joint state of the system $\mathcal{S} \otimes \mathcal{M}$ is then $\frac{1}{\sqrt{|\mathbb{G}_\lambda|}} \sum_{g \in \mathbb{G}_\lambda} |g\rangle_\mathcal{S}|g * x_\lambda\rangle_\mathcal{M}$.
  - Apply $\mathsf{QFT}_{\mathbb{G}_\lambda}$ to $\mathcal{S}$ again, yielding $\frac{1}{|\mathbb{G}_\lambda|} \sum_{g,h \in \mathbb{G}_\lambda} \chi(g,h)|h\rangle_\mathcal{S}|g * x_\lambda\rangle_\mathcal{M}$
  - Measure $\mathcal{S}$, giving the serial number $\sigma := h$. The $\mathcal{M}$ register then collapses to the banknote $\$ = |\mathbb{G}_\lambda^h * x_\lambda\rangle := \frac{1}{\sqrt{|\mathcal{G}_\lambda|}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)|g * x_\lambda\rangle_\mathcal{M}$. Output $(\sigma, \$)$.

- $\mathsf{Ver}(\sigma, \$)$: First verify that the support of $\$$ is contained in $\mathcal{X}_\lambda$, by applying the assumed algorithm for recognizing $\mathcal{X}_\lambda$ in superposition. Then repeat the following λ times:
  - Initialize a new register $\mathcal{H}$ to $(|0\rangle_\mathcal{H} + |1\rangle_\mathcal{H})/\sqrt{2}$.
  - Choose a random group element $u \in \mathbb{G}_\lambda$.

– Apply to $\mathcal{H} \otimes \mathcal{M}$ in superposition the map

$$\mathsf{Apply}|b\rangle_{\mathcal{H}}|y\rangle_{\mathcal{M}} \mapsto \begin{cases} |0\rangle_{\mathcal{H}}|y\rangle_{\mathcal{M}} & \text{if } b = 0 \\ |1\rangle_{\mathcal{H}}|(-u) * y\rangle_{\mathcal{M}} \text{ [9]} & \text{if } b = 1 \end{cases}$$

In the case that \$ is the correct banknote state $|\mathbb{G}_\lambda^h * x_\lambda\rangle$, the result of applying Apply is:

$$\frac{1}{\sqrt{2|\mathbb{G}_\lambda|}} \left(|0\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)|g * x_\lambda\rangle_{\mathcal{M}} + |1\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)|(g-u) * x_\lambda\rangle_{\mathcal{M}} \right)$$
$$= \frac{1}{\sqrt{2|\mathbb{G}_\lambda|}} \left(|0\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)|g * x_\lambda\rangle_{\mathcal{M}} + |1\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g+u,h)|g * x_\lambda\rangle_{\mathcal{M}} \right)$$
$$= \frac{1}{\sqrt{2|\mathbb{G}_\lambda|}} \left(|0\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)|g * x_\lambda\rangle_{\mathcal{M}} + |1\rangle_{\mathcal{H}} \sum_{g \in \mathbb{G}_\lambda} \chi(g,h)\chi(u,h)|g * x_\lambda\rangle_{\mathcal{M}} \right)$$
$$= \frac{1}{\sqrt{2}} \left(|0\rangle_{\mathcal{H}} + \chi(u,h)|1\rangle_{\mathcal{H}}\right) |\mathbb{G}_\lambda^h * x_\lambda\rangle$$

– Measure $\mathcal{H}$ in the basis $B_{h,u} := \{(|0\rangle_{\mathcal{H}} + \chi(u,h)|1\rangle_{\mathcal{H}})/\sqrt{2}, (|0\rangle_{\mathcal{H}} - \chi(u,h)|1\rangle_{\mathcal{H}})/\sqrt{2}\}$, giving a bit $b_u \in \{0,1\}$. Discard the $\mathcal{H}$ register. In the case that \$ is the correct banknote state $|\mathbb{G}_\lambda^h * x_\lambda\rangle$, b_u will be 0 with probability 1, and $\mathcal{M}$ will be left in the original banknote state.

If all the b_u are 0 and the support of \$ is contained in $\mathcal{X}_\lambda$, then accept. If any of the b_u are 1, or if the support is not contained in $\mathcal{X}_\lambda$, reject. We see that for the correct banknote, Ver accepts with probability 1.

Remark 3.2. If using a probabilistic setup of the group action, there are two options. The first is to have Gen set up the group action, and have the parameters be included in the serial number. The second is to have a trusted third party set up the group action, and publish the parameters in a common reference string (CRS). If the goal is only quantum money security, then the former option is always possible, since the security experiment uses an honestly generated serial number. If the goal is quantum lightning security, the former option may not be possible, as the adversary computes the serial number; it may be that there are bad choices of parameters for the group action (and hence the CRS inside the serial number) which make it easy to forge banknotes. Therefore, for quantum lightning security, we would expect using a trusted setup to generate a CRS containing the group action parameters.

3.1 Accepting States of the Verifier

Above we showed that honest banknote states are accepted by the verifier. We now prove that, roughly, honest banknote states are the *only* states accepted by the verifier, with overwhelming probability.

[9] Note that we used the "minimal" oracle here for the group action computation, having $(-u) * y$ replace y, instead of being written to a response register as in the standard quantum oracle. However, since the computation $y \mapsto (-u) * y$ is efficiently reversible (by $y \mapsto u * y$), we can easily implement the minimal oracle efficiently by first computing $|(-u) * y\rangle_{\mathcal{M}'}$ in a new register $\mathcal{M}'$, then uncomputing $|y\rangle_{\mathcal{M}}$ using the efficient inverse (so it now contains $|0\rangle_{\mathcal{M}}$), and finally swapping $\mathcal{M}'$ with $\mathcal{M}$.

Theorem 3.3. *Let $|\psi\rangle$ be a state over $\mathcal{M}$. Then* $\Pr[\mathsf{Ver}(h, |\psi\rangle) = 1] = \|\langle\psi|\mathbb{G}_\lambda^h * x_\lambda\rangle\|^2(1-2^{-\lambda})+2^{-\lambda}$.

In other words, we can treat $\mathsf{Ver}(h, |\psi\rangle)$ as projecting onto $|\mathbb{G}_\lambda^h * x_\lambda\rangle$, incurring only a negligible error. The remainder of this subsection is devoted to proving Theorem 3.3.

Lemma 3.4. *For $h' \neq h$, $\langle\mathbb{G}_\lambda^{h'} * x_\lambda|\mathbb{G}_\lambda^h * x_\lambda\rangle = 0$*

Proof.

$$\begin{aligned}\langle\mathbb{G}_\lambda^{h'} * x_\lambda|\mathbb{G}_\lambda^h * x_\lambda\rangle &= \frac{1}{|\mathbb{G}_\lambda|}\sum_{g,g'\in\mathbb{G}_\lambda}\chi(g',h')^{-1}\chi(g,h)\langle g' * x_\lambda|g * x_\lambda\rangle\\ &= \frac{1}{|\mathbb{G}_\lambda|}\sum_{g\in\mathbb{G}_\lambda}\chi(g,h')^{-1}\chi(g,h) = \frac{1}{|\mathbb{G}_\lambda|}\sum_{g\in\mathbb{G}_\lambda}\chi(g,h-h') = 0\end{aligned}$$

□

Let $|\psi\rangle$ be a a state with support on $\mathcal{X}$. Since the $|\mathbb{G}^{h'} * x_\lambda\rangle$ are orthogonal and the number of h' equals the size of $\mathcal{X}$, the set $\{|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle\}_{h'}$ forms a basis for the set of states with support on $\mathcal{X}$. We can then write $|\psi\rangle = \sum_{h'}\alpha_{h'}|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle$ where $\sum_{h'}\|\alpha_{h'}\|^2 = 1$. We then have $\|\alpha_h\|^2 = \|\langle\psi|\mathbb{G}_\lambda^h * x_\lambda\rangle\|^2$.

Consider a single iteration of Ver on serial number h, which samples a random u, initializes $\mathcal{H}$ to $(|0\rangle + |1\rangle)/\sqrt{2}$, applies the map Apply, and then measures $\mathcal{H}$ is basis $B_{h,u}$ to get outcome b. Let $|\psi'\rangle$ be the post-measurement state of $\mathcal{M}$ conditioned on $b = 0$.

Lemma 3.5. *Conditioned on u, $p := \Pr[b_u = 0] = \frac{1}{4}\sum_{h'}\|\alpha_{h'}\|^2\|1+\chi(u,h-h')\|^2$, and $|\psi'\rangle = \frac{1}{\sqrt{p}}\sum_{h'}\alpha_{h'}\frac{1+\chi(u,h-h')}{2}|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle_{\mathcal{M}}$.*

Proof. By adapting the correctness proof above, we see that the state after applying Apply (but before measurement) is:

$$|\phi\rangle = \sum_{h'\in\mathbb{G}_\lambda}\alpha_{h'}\frac{1}{\sqrt{2}}\left(|0\rangle_{\mathcal{H}} + \chi(u,h')|1\rangle_{\mathcal{H}}\right)|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle_{\mathcal{M}}$$

Then p is length squared of the projection of $|\phi\rangle$ onto $(|0\rangle_{\mathcal{H}} + \chi(u,h)|1\rangle_{\mathcal{H}})/\sqrt{2}$. Therefore, $p = \frac{1}{4}\sum_{h'}\|\alpha_{h'}\|^2\|1+\chi(u,h')^{-1}\chi(u,h)\|^2 = \frac{1}{4}\sum_{h'}\|\alpha_{h'}\|^2\|1+\chi(u,h-h')\|^2$. Before re-normalization, the state of $\mathcal{M}$ conditioned on $b = 0$ is then $\sum_{h'}\alpha_h\frac{1+\chi(u,h-h')}{2}|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle_{\mathcal{M}}$. Re-normalization gives $|\psi'\rangle$. □

We now iterate, replacing $\alpha_{h'}$ with $\alpha_{h'}\frac{1+\chi(u,h-h')}{2}/\sqrt{p}$. This means that after λ trials, conditioned on trial i using u_i and giving measurement outcome b_i, we have that

$$p_{\mathsf{final}} := \Pr[b_1 = b_2 = \cdots = b_\lambda = 0] = \frac{1}{4^\lambda}\sum_{h'}\|\alpha_{h'}\|^2\prod_{1=1}^{\lambda}\|1+\chi(u_i,h-h')\|^2$$

We now average over u to get $\mathbb{E}[p_{\mathsf{final}}]$, the overall probability that Ver accepts $|\psi\rangle$.

$$\begin{aligned}
\mathbb{E}[p_{\mathsf{final}}] &= \frac{1}{(4|\mathbb{G}_\lambda|)^\lambda} \sum_{h', u_1, \cdots, u_\lambda} \|\alpha_{h'}\|^2 \prod_{i=1}^{\lambda} \|1 + \chi(u_i, h - h')\|^2 \\
&= \sum_{h'} \|\alpha_{h'}\|^2 \prod_{i=1}^{\lambda} \left(\frac{1}{4|\mathbb{G}_\lambda|} \sum_u \|1 + \chi(u, h - h')\|^2 \right) \\
&= \sum_{h'} \|\alpha_{h'}\|^2 \prod_{i=1}^{\lambda} \left(\frac{1}{4|\mathbb{G}_\lambda|} \sum_u 2 + \chi(u, h - h') + \chi(u, h - h')^{-1} \right) \\
&= \|\alpha_h\|^2 + 2^\lambda \sum_{h' \neq h} \|\alpha_{h'}\|^2 = \|\alpha_h\|^2 + 2^{-\lambda}(1 - \|\alpha_h\|^2) \\
&= \|\alpha_h\|^2 (1 - 2^{-\lambda}) + 2^{-\lambda}
\end{aligned}$$

This completes the proof of Theorem 3.3. □

3.2 Computing the Serial Number

Here, we show that, given a valid banknote $\$ = |\mathbb{G}_\lambda^h * x_\lambda\rangle$ with unknown serial number h, it is possible to efficiently compute h. This result is not needed anywhere else in the paper, but is included in case it may be useful for future work building on our construction.

Theorem 3.6. *There exists a QPT algorithm* Findh *and a negligible function* $\mathsf{negl}(\lambda)$ *such that, on input* $|\mathbb{G}_\lambda^h * x_\lambda\rangle$*, outputs* h *with probability at least* $1 - \mathsf{negl}(\lambda)$*.*

Proof. Recall from the description of Ver that, for a given u and given $|\mathbb{G}_\lambda^h * x_\lambda\rangle$, we can compute the state $|\tau_{u,h}\rangle|\mathbb{G}_\lambda^h * x_\lambda\rangle$ where $|\tau_{u,h}\rangle := \frac{1}{\sqrt{2}}(|0\rangle_{\mathcal{H}} + \chi(u,h)|1\rangle_{\mathcal{H}})$. Since this process still gives us $|\mathbb{G}_\lambda^h * x_\lambda\rangle$, we can repeat the process, computing $|\tau_{u_i,h}\rangle$ for many different u_i.

A naive solution is to compute many copies of $|\tau_{u,h}\rangle$ for some $u \in \mathbb{G}_\lambda$, and then do state tomography to recover $\chi(u,h)$. If $\mathbb{G}_\lambda$ were cyclic, then $\chi(u,h)$ will uniquely determine h. The problem is that, since $\mathbb{G}_\lambda$ is exponentially large, the distance between $\chi(u,h)$ as h varies will be exponentially small. This means doing state tomography to a sufficiently small error to recover h would require exponentially-many samples and therefore be inefficient. However, by choosing the u_i carefully and being a bit more thoughtful, we can recover h in polynomial time.

Our strategy will still be to compute many copies of $|\tau_{u,h}\rangle$ for some u and do state tomography to recover an estimate $\hat{\chi}(u,h)$ for $\chi(u,h)$. In time $\mathsf{poly}(\lambda, 1/\epsilon, \log(1/\delta))$, we can guarantee that $\Pr[\|\hat{\chi}(u,h) - \chi(u,h)\| < \epsilon] \geq 1 - \delta$, for any desired inverse-polynomial ϵ and exponentially-small δ. We then do this for many different carefully chosen u, which allows us to correct the errors arising from tomography, as we now explain.

The cyclic case. Suppose $\mathbb{G}_\lambda$ is cyclic, and is therefore isomorphic to the additive group $\mathbb{Z}_N$. In this case, $\chi(u,h) = e^{i2\pi uh/N} = \omega_N^{uh}$, where $\omega_N = e^{i2\pi/N}$.

Now when we do state tomorgraphy and recover $\hat{\chi}(u,h)$, we learn an estimate of $uh \bmod N$. In more detail, given real number a and real number R, we let $a \bmod R$ denote the unique value of $a - Rk$ for integer k that lies in $(-R/2, R/2]$. Since we know $\|\chi(u,h)\| = 1$, we can assume, by normalizing if necessary, that $\|\hat{\chi}(u,h)\|$ is also 1. Therefore, $\hat{\chi}(u,h) = e^{i\theta}$ for some $\theta \in (-\pi, \pi]$.

Then by the tomography guarantee, we have $|\ [\theta - (2\pi uh/N)] \bmod 2\pi\ | \leq \epsilon$, or equivalently $|\ [N\theta/2\pi - uh] \bmod N\ | \leq \epsilon N/2\pi$, except with negligible probability

This means we reduce the computation of h to the following classical task: we get to choose arbitrary $u_i \in \mathbb{Z}_N$ for $i = 1, \dots, n$. In response, we learn $u_i h + e_i \bmod N$, where e_i is some random variable in $[-\epsilon N/2\pi, \epsilon N/2\pi]$. In vector notation, we can write this as choosing a vector $\mathbf{u} \in \mathbb{Z}_N^n$, and receiving $h\mathbf{u} + \mathbf{e} \bmod N$, where $\mathbf{e}$ is a vector whose components are independent random variables that are guaranteed to be in $[-\epsilon N/2\pi, \epsilon N/2\pi]$. The goal is to compute h.

This looks very similar to a 1-dimensional version of the LWE problem [Reg05] (or more accurately, bounded distance decoding) except that in our case we get to choose the vector $\mathbf{u}$ in whatever way so as to make the task *easy*. We can then use known techniques to find h. In particular, we can choose $\mathbf{u} = (1, 2, 4, , 8, \cdots, 2^{n-1})$ where $n = \lceil \log_2 N \rceil$. This is known as the gadget "matrix"[10]. Importantly, $\mathbf{u}$ has an efficiently computable "trapdoor". That is, write $N = \sum_{i=0}^{n-1} 2^i \times N_i$ for bits N_i, and let

$$\mathbf{A} = \begin{pmatrix} 2 & -1 & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 2 & -1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 2 & -1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 2 & -1 \\ N_0 & N_1 & N_2 & N_3 & N_4 & \cdots & N_{n-2} & N_{n-1} \end{pmatrix}$$

Then $\mathbf{A}$ is full rank over the integers, but satisfies $\mathbf{A} \cdot \mathbf{u} \bmod N = 0^n$. Set $\epsilon = \pi/3n$. Thus, given $\mathbf{v} := \mathbf{u}h + \mathbf{e} \bmod N$, we can compute

$$\mathbf{A}^{-1} \cdot (\mathbf{A} \cdot \mathbf{v} \bmod N) = \mathbf{A}^{-1} \cdot (\mathbf{A} \cdot \mathbf{e} \bmod N) = \mathbf{A}^{-1} \cdot (\mathbf{A} \cdot \mathbf{e}) = \mathbf{e} \ .$$

Above, we used the fact that the entries of $\mathbf{A} \cdot \mathbf{e}$ have absolute value at most $n \times \max_{i,j} |\mathbf{A}_{i,j}| \times \max_j |\mathbf{e}_j| = n \times 2 \times \epsilon N/2\pi \leq N/3 < N/2$, meaning that reduction mod N has no effect.

Once we compute $\mathbf{e}$, we can then compute $h\mathbf{u} = \mathbf{v} - \mathbf{e}$, and then h is just the first component.

The general case. We cow consider the case of general groups. Let $\mathbb{G}_\lambda = \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \cdots \times \mathbb{Z}_{n_k}$. Write $h = (h_1, \cdots, h_k)$. By choosing $u = (u_1, 0, \cdots, 0)$, the task of computing h_1 reduces to the case where $\mathbb{G}_\lambda = \mathbb{Z}_{n_1}$, which can be solved via the algorithm above. Likewise, we can compute $h_2, \cdots, h_k$, and hence h. □

3.3 Security

Here, we prove the security of our scheme under two new but plausible assumptions on group actions. For now, we assume that it is impossible to obliviously sample set elements. This is false on group actions based on elliptic curves, but our proof here is simpler and provides the main intuition. In Section 3.4, we address the case where oblivious sampling is possible.

The Knowledge of Group Element Assumption (KGEA). This assumption states, informally, that any algorithm that produces a set element y must "know" g such that $y = g * x_\lambda$. We first discuss the case where it is infeasible to sample set elements in the group action. Later, we will

[10] In our case the matrix has width 1, whereas in general applications the matrix will have many columns.

discuss how to model the assumption when there is a such a sampling algorithm. In the classical setting, the KGEA assumption would be formalized as follows:

Assumption 3.7. The *classical knowledge of group element assumption* (C-KGEA) holds on a group action $(\mathbb{G}, \mathcal{X}, *)$ if the following is true. For any probabilistic polynomial time (PPT) adversary $\mathcal{A}$, there exists a PPT "extractor" $\mathcal{E}$ and a negligible ϵ such that:

$$\Pr\left[y \in \mathcal{X} \wedge y \neq g * x_\lambda : {}^{y \leftarrow \mathcal{A}(1^\lambda; r)}_{g \leftarrow \mathcal{E}(1^\lambda, r)}\right] \leq \epsilon(\lambda) \ .$$

Above, r are the random coins given to $\mathcal{A}$, which are also given to $\mathcal{E}$, and the probability is taken over uniform r and any additional randomness of $\mathcal{E}$.

In other words, if $\mathcal{A}$ outputs any set element, it must "know" how to derive that set element from x_λ, since it can compute g such that $y = g * x_\lambda$ using $\mathcal{E}$ and its random coins. Note that once the random coins are fixed, $\mathcal{A}$ is deterministic.

As observed by [LMZ23], when moving to the quantum setting, the problem with Assumption 3.7 is that quantum algorithms do not have to flip random coins to generate randomness, and instead their output may be a measurement applied to a quantum state, the result being inherently randomized even if the quantum state is fixed. Thus there is no meaningful way to give the same random coins to $\mathcal{E}$.

The solution used in [LMZ23] is to, instead of giving $\mathcal{E}$ the same inputs as $\mathcal{A}$, give $\mathcal{E}$ the remaining state of $\mathcal{A}$ at the *end* of the computation. This requires some care, since an algorithm can of course forget any bit of information by simply throwing it away. A more sophisticated way to lose information is to perform other measurements on the state, say measuring in the Fourier basis. The solution in [LMZ23] is to require that $\mathcal{A}$ makes no measurements at all, *except* for measuring the final output. Note that the Principle of Delayed Measurement implies that it is always possible without loss of generality to move all measurements to the final output. Then $\mathcal{E}$ is given both the output and the remaining quantum state of $\mathcal{A}$, and tries to compute g. Note that in the classical setting, if we restrict to *reversible* $\mathcal{A}$, this formulation of giving $\mathcal{E}$ the final state of $\mathcal{A}$ is equivalent to given $\mathcal{E}$ the randomness, since the randomness can be computed by reversing $\mathcal{A}$. Similar to how we can assume a quantum $\mathcal{A}$ makes all its measurements at the end, in we can always assume without loss of generality that a classical $\mathcal{A}$ is reversible. Thus, in the classical setting these two definitions coincide. Adapting to our setting, this approach yields the following assumption:

Assumption 3.8. The *quantum knowledge of group element assumption* (Q-KGEA) holds on a group action $(\mathbb{G}, \mathcal{X}, *)$ if the following is true. For any quantum polynomial time (QPT) adversary $\mathcal{A}$ which performs no measurements except for its final output, there exists a QPT extractor $\mathcal{E}$ and negligible ϵ such that

$$\Pr\left[y \in \mathcal{X} \wedge y \neq g * x_\lambda : {}^{(y, |\psi\rangle) \leftarrow \mathcal{A}(1^\lambda)}_{g \leftarrow \mathcal{E}(y, |\psi\rangle)}\right] \leq \epsilon(\lambda) \ .$$

Above, y is considered as the output of $\mathcal{A}$, and the only measurements applied to $\mathcal{A}$ is the measurement of y to obtain the output.

In group actions based on elliptic curves, it is possible to directly sample set elements. While set elements generated in this way have no obvious relation to other set elements, the ability to generate set elements without applying the group action would technically contradict the KGEA assumptions as defined in Assumptions 3.7 and 3.8. This same issue was present in the knowledge of path assumption in [LMZ23]. In Section 3.4, we discuss a different approach to remedy this issue that seems more robust. For now we proceed with the basic setting where we assume Q-KGEA as in Assumption 3.8 is true.

The Discrete Log Assumption, with Help. We now define a strengthening of the Discrete Log assumption (Assumption 2.4), which allows the adversary limited query access to a computational Diffie Hellman (CDH) oracle.

Assumption 3.9. We say that the *Discrete Log with a single minimal CDH query* assumption (DLog/1-minCDH) assumption holds if the following is true. For any QPT adversary $\mathcal{A}$ playing the following game, parameterized by λ, there is a negligible ϵ such that $\mathcal{A}$ wins with probability at most $\epsilon(\lambda)$:

- The challenger, on input λ, chooses a random $g \in \mathbb{G}_\lambda$. It sends λ to $\mathcal{A}$
- $\mathcal{A}$ submits a superposition query $\sum_{y\in\mathcal{X},z\in\{0,1\}^*} \alpha_{y,z}|y,z\rangle$. Here, y is a set element that forms the query, and z is the internal state of the adversary when making the query. The challenger responds with $\sum_{y\in\mathcal{X},z\in\{0,1\}^*} \alpha_{y,z}|(-g)*y,z\rangle$ [11].
- The challenger sends $g * x$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs a guess g' for g. It wins if $g' = g$.

Note that Assumption 3.9 uses a "minimal" oracle for the CDH oracle, meaning is replaces y with $(-g) * y$. This is only a possibility because $y \mapsto (-g) * y$ is reversible; otherwise the query would not be unitary. The minimal oracle, however, is somewhat non-standard. So we here define a slightly different assumption which uses "standard" oracles:

Assumption 3.10. We say that the *Discrete Log with a double standard CDH query* assumption (DLog/2-stdCDH) assumption holds if the following is true. For any QPT adversary $\mathcal{A}$ playing the following game, parameterized by λ, there is a negligible ϵ such that $\mathcal{A}$ wins with probability at most $\epsilon(\lambda)$:

- The challenger, on input λ, chooses a random $g \in \mathbb{G}_\lambda$. It sends λ to $\mathcal{A}$.
- $\mathcal{A}$ submits a superposition query $\sum_{y\in\mathcal{X},w,z\in\{0,1\}^*} \alpha_{y,w,z}|y,w,z\rangle$. Here, y is a set element that forms the query, w is a string that forms the response register, and z is the internal state of the adversary when making the query. The challenger responds with $\sum_{y\in\mathcal{X},w,z\in\{0,1\}^*} \alpha_{y,w,z}|y,w \oplus [(-g)*y],z\rangle$.
- $\mathcal{A}$ submits a second superposition query $\sum_{y\in\mathcal{X},w,z\in\{0,1\}^*} \alpha_{y,w,z}|y,w,z\rangle$. The challenger responds with $\sum_{y\in\mathcal{X},w,z\in\{0,1\}^*} \alpha_{y,w,z}|y,w\oplus[g*y],z\rangle$.
- The challenger sends $g * x$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs a guess g' for g. It wins if $g' = g$.

Lemma 3.11. *If DLog/2-stdCDH (Assumption 3.10) holds in a group action, then so does DLog/1-minCDH (Assumption 3.9).*

[11] Note that this operation is unitary and efficiently computable since $y \mapsto (-g) * y$ is efficiently computable and efficiently reversible given g.

Proof. Consider a supposed adversary $\mathcal{A}$ for DLog/1-minCDH with non-negligible winning probability ϵ. We construct a new adversary $\mathcal{B}$ for DLog/2-stdCDH with the same non-negligible winning probability as follows. $\mathcal{B}$ runs $\mathcal{A}$ until $\mathcal{A}$ makes its superposition query $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|y\rangle_{\mathcal{Y}}|z\rangle_{\mathcal{Z}}$. $\mathcal{B}$ then initializes a new register $\mathcal{R}$ with the state $|0\rangle$. $\mathcal{B}$ then submits $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|y\rangle_{\mathcal{Y}}|0\rangle_{\mathcal{R}}|z\rangle_{\mathcal{Z}}$ as its first query. In response, it receives $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|y\rangle_{\mathcal{Y}}|(-g)*y\rangle_{\mathcal{R}}|z\rangle_{\mathcal{Z}}$. Now it swaps the roles of $\mathcal{R}$ and $\mathcal{Y}$, and makes its second query on $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|(-g)*y\rangle_{\mathcal{Y}}|y\rangle_{\mathcal{R}}|z\rangle_{\mathcal{Z}}$. In response it receives $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|(-g)*y\rangle_{\mathcal{Y}}|0\rangle_{\mathcal{R}}|z\rangle_{\mathcal{Z}}$. Then it discards the $\mathcal{R}$ register, and sends the resulting state $\sum_{y\in\mathcal{X},z\in\{0,1\}^*}\alpha_{y,z}|(-g)*y\rangle_{\mathcal{Y}}|z\rangle_{\mathcal{Z}}$ to $\mathcal{A}$.

Afterward, when $\mathcal{B}$ receives $g*x$, it forwards it to $\mathcal{A}$, and then outputs whatever g' that $\mathcal{A}$ outputs. Thus, we see that $\mathcal{B}$ correctly simulates the view of $\mathcal{A}$, and thus the probability $\mathcal{B}$ wins is the same as $\mathcal{A}$, namely ϵ. □

From this point forward, we will use DLog/1-minCDH as our assumption; Lemma 3.11 then shows that we could have instead used DLog/2-stdCDH.

The security proof. We are now ready to formally state and prove security.

Theorem 3.12. *Assuming Q-KGEA (Assumption 3.8) and DLog/1-minCDH (Assumption 3.9) both hold on a group action $(\mathbb{G},\mathcal{X},*)$, then Construction 3.1 is a quantum lightning scheme.*

Remark 3.13. Before proving Theorem 3.12, we briefly discuss how to handle the case of non-uniform attackers, since in this setting quantum lightning is insecure without some modifications. Note that even against non-uniform attackers, DLog/1-minCDH still plausibly holds. However, Q-KGEA certainly does not, as a non-uniform attacker may have a y hard-coded for which it does not know the discrete log with x_λ. As discussed in Section 2, there are several possibilities.

- The first is to restrict to non-uniform attackers that only have classical advice. While classical advice does not appear to be useful in breaking Construction 3.1, it still allows for breaking Q-KGEA; thus while our scheme may be secure in this setting, the security proof would be vacuous.

- The second is to use a probabilistically generated group action, and define Q-KGEA and DLog/1-minCDH accordingly. For quantum money security, it would suffice to have Gen create the parameters of the group action and then include them in the serial number, since the serial number is generated honestly. For quantum lightning security, we would instead need the parameters to be generated by a trusted third party and then placed in a common random string (CRS).

- The final option is to use the human ignorance approach [Rog06], where we explicitly state our security theorem as transforming a quantum lightning adversary into a Q-KGEA adversary; while such Q-KGEA adversaries exist in the non-uniform setting without a CRS, they are presumably unknown to human knowledge. As a consequence, a quantum lightning attacker, while existing, would likewise be unknown to human knowledge.

For simplicity, state and prove Theorem 3.12 in the uniform setting; either probabilistically generating the group action or using human ignorance would require straightforward modifications.

We now are ready to prove Theorem 3.12.

Proof. Consider a QPT quantum lightning adversary $\mathcal{A}$ which breaks security with non-negligible success probability ϵ. Since an adversary can always tell if it succeeded by running Ver, we can run $\mathcal{A}$ multiple times to boost the probability of a successful break. In particular, we can run $\mathcal{A}$ for $\lambda\epsilon$, and at except with probability $1-2^{-\Theta(\lambda)}$, at least one of the runs will succeed. This allows us to conclude without loss of generality that $\mathcal{A}$ has success probability $1-2^{-\Theta(\lambda)}$. By Theorem 3.3, we also know that if $\mathcal{A}$ outputs a serial number h, the states outputted are exponentially close to two copies of $|\mathbb{G}_\lambda^h * x_\lambda\rangle$.

For simplicity in the following proof, we will assume the probability of passing verification is actually 1; it is straightforward to adapt the proof to the case of negligible error.

Next, we purify $\mathcal{A}$, and assume that before measurement, $\mathcal{A}$ outputs a pure state $|\psi\rangle$. By our assumption that the success probability is 1, $|\psi\rangle$ will have the form

$$|\psi\rangle = \sum_h \alpha_h|\phi_h\rangle|\mathbb{G}_\lambda^h * x_\lambda\rangle|\mathbb{G}_\lambda^h * x_\lambda\rangle = \frac{1}{|\mathbb{G}_\lambda|}\sum_h \alpha_h|\phi_h\rangle\chi(h, g_1+g_2)|g_1 * x\rangle_{\mathcal{M}_1}|g_2 * x\rangle_{\mathcal{M}_2} \ .$$

Above, $|\phi_h\rangle$ are arbitrary normalized states representing whatever state the adversary contains after outputting its banknotes, and $\sum_h \|\alpha_h\|^2 = 1$.

Now consider the adversary $\mathcal{B}$ which first constructs $|\psi\rangle$, and then measures the register $\mathcal{M}_2$ to obtain $y_2 = g_2 * x$.

Claim 3.14. *g_2 is uniform in $\mathbb{G}$.*

Proof. Consider additionally measuring $\mathcal{M}_1$ in the basis $\{|\mathbb{G}_\lambda^h * x_\lambda\rangle\}$. This this measurement is on a different register than the measurement on $\mathcal{M}_2$, measuring $\mathcal{M}_1$ does not affect the output distribution of $\mathcal{M}_2$ (though the results may be correlated). But the measurement on $\mathcal{M}_1$ determines h, and conditioned on h, $\mathcal{M}_2$ collapses to $|\mathbb{G}_\lambda^h * x_\lambda\rangle$. Regardless of what h is, measuring $|\mathbb{G}_\lambda^h * x_\lambda\rangle$ gives a uniformly random element in $\mathcal{X}$. Thus, even without measuring $\mathcal{M}_1$, the measurement of $\mathcal{M}_2$ gives a uniform element in $\mathcal{X}$. □

Therefore, after measuring $\mathcal{M}_2$, the state $|\psi\rangle$ then collapses to

$$|\psi_{g_2 * x_\lambda}\rangle := \frac{1}{\sqrt{|\mathbb{G}_\lambda|}}\sum_h \alpha_h|\phi_h\rangle\chi(h, g_1+g_2)|g_1 * x\rangle_{\mathcal{M}_1} \ .$$

Claim 3.15. *There is a QPT procedure* Map *such that* $\mathsf{Map}(g, |\psi_y\rangle) = |\psi_{g*y}\rangle$.

Proof. Map simply applies the map $y \mapsto (-g) * y$ to $\mathcal{M}_1$ in superposition. Then we have that:

$$\begin{aligned}\mathsf{Map}(g, |\psi_{g_2 * x_\lambda}\rangle) &= \frac{1}{\sqrt{|\mathbb{G}_\lambda|}}\sum_h \alpha_h|\phi_h\rangle\chi(h, g_1+g_2)|(g_1 - g) * x\rangle_{\mathcal{M}_1} \\ &= \frac{1}{\sqrt{|\mathbb{G}_\lambda|}}\sum_h \alpha_h|\phi_h\rangle\chi(h, g_1'+g+g_2)|g_1' * x\rangle_{\mathcal{M}_1} = |\psi_{(g+g_2)*y}\rangle = |\psi_{g*(g_2*y)}\rangle\end{aligned}$$

Above we used the change of variables $g_1' = g_1 - g$. □

Now we invoke Q-KGEA (Assumption 3.8) on the adversary $\mathcal{B}$. Since $\mathcal{B}$ always outputs a valid set element, this means there is another QPT algorithm $\mathcal{E}$ such that

$$\Pr[\mathcal{E}(g_2 * x_\lambda, |\psi_{g_2 * x_\lambda}\rangle) = g_2] \geq 1 - \mathsf{negl}(\lambda)$$

Above, the probability is over $g_2 * x_\lambda$, as well as any randomness incurred when executing $\mathcal{E}$. We note by a simple random self-reduction that we can insist the above probability holds for *all* $g_2 * x_\lambda$, where the randomness is only over $\mathcal{E}$. Indeed, given $|\psi_{g_2 * x_\lambda}\rangle, g_2 * x_\lambda$, we can choose a random g and compute $g_2' * x_\lambda$ as $g * (g_2 * x_\lambda)$ where $g_2' = g + g_2$. Likewise, we can compute $|\psi_{g_2' * x_\lambda}\rangle$ as $\mathsf{Map}(g, |\psi_{g_2 * x_\lambda}\rangle)$. This gives a random instance on which to apply $\mathcal{E}$, giving g_2' with probability $1 - \mathsf{negl}(\lambda)$, regardless of g_2. Then we can compute $g_2 = g_2' - g$. We thus compute g_2 with overwhelming probability, even in the worst case. We will therefore assume without loss of generality that this is the case for $\mathcal{E}$.

For simplicity, we will actually assume that the probability is 1; it is straightforward to handle the case the probability is negligibly close to 1. By the Gentle Measurement Lemma [Win99], $\mathcal{E}$ can compute g_2 without altering the state $|\psi_{g_2 * x}\rangle$. Thus, by combining $\mathcal{B}$ and $\mathcal{E}$, we can compute both $|\psi_{g_2 * x}\rangle$ and g_2 with probability 1. We can then compute $\mathsf{Map}(-g_2, |\psi_{g_2 * x_\lambda}\rangle) = |\psi_{x_\lambda}\rangle$.

We now describe a new algorithm $\mathcal{C}$ which breaks DLog/1-minCDH (Assumption 3.9). $\mathcal{C}$ works as follows:

- It constructs $|\psi_{x_\lambda}\rangle$ as above.
- It makes its query to the DLog/1-minCDH challenger, setting $\mathcal{M}_1$ as the query register. This query simulates the operation $\mathsf{Map}(g, \cdot)$, where g is the group element chosen by the challenger. Thus, at the end of the query, $\mathcal{C}$ has $|\psi_{g * x_\lambda}\rangle$.
- Now upon receiving $g * x_\lambda$ from the challenger, run $\mathcal{E}(g * x_\lambda, |\psi_{g * x_\lambda}\rangle)$. By the guarantees of $\mathcal{E}$, the output will be g.

Thus we see that $\mathcal{C}$ breaks the DLog/1-minCDH assumption. This completes the security proof. □

3.4 Security under existence of obliviously sampled elements

As previously mentioned, the ability to obliviously sample set elements in group actions based on elliptic curves means the KGEA assumption as stated is false. One possible remedy, used in [LMZ23], explicitly assumes a probabilistic classical procedure $S()$ for obliviously sampling set elements, and modifies the KGEA assumption so that the extractor either outputs (1) an explanation relative to x_λ *or* (2) an explanation relative to some input y together with the random coins r that are fed into S so that $y = S(r)$. The problem with this approach is that the assumption depends on explicitly modeling the oblivious sampling procedure, and if another oblivious sampling procedure is found, it would contradict the assumption.

In order to give a more robust proof, we here devise an alternate solution. Our key idea is to observe that, while obliviously sampling elements strictly speaking violates the KGEA assumption, it does not seem to yield any assistance in actually breaking a quantum money scheme based on group actions, since the elements obliviously sampled will be unrelated to anything else. More generally, we can consider a general cryptographic game that an adversary may play with a challenger. For "nice" games (which we will define shortly), in particular games that only use the group action interface and do not themselves obliviously sample elements, it seems that giving the adversary the ability to obliviously sample elements is no help. We therefore postulate that, for any adversary $\mathcal{A}$ that wins such a nice game, there is a different adversary $\mathcal{A}'$ for which the KGEA assumption can be appled, yielding an extractor *for that* $\mathcal{A}'$. Thus, even if the original $\mathcal{A}$ can obliciously sample elements, we essentially assume that $\mathcal{A}'$ cannot, and therefore $\mathcal{E}$ is possible. We now make this intuition precise.

Generic Group Action Games. We first introduce the notion of generic group action games. Note that we will only be interested in *games* that are given by generic algorithms; we will always treat the adversary as non-generic.

Briefly, a generic group action game is given by an interactive algorithm ("challenger") Ch. Ch is limited to only performing group action computations that are "generic" and only interacts with the group action through oracles implementing the group action interface. Specifically, a generic algorithm is an oracle-aided algorithm $\mathcal{B}$ that has access to oracles $\mathsf{GA} = (\mathsf{Start}, \mathsf{Act}, \mathsf{Mem})$. Here, Start is the oracle that takes as input the empty query, and outputs a string $\tilde{x}$ representing x_λ. Act is the oracle that takes as input a group element $g \in \mathbb{G}_\lambda$ and a string $\tilde{y}$ representing a set element y, and outputs a string $\tilde{z}$ representing $z = g * x$. Finally, Mem is a membership testing oracle, that tests is a given string $\tilde{x}$ represents an actual set element. From a generic game, we obtain a standard model game by implementing the oracles $\mathsf{Start}, \mathsf{Act}, \mathsf{Mem}$ with the algorithms for an actual group action: Start outputs the actual set element x_λ, Act is the group action $*$, and Mem is the membership tester for the set $\mathcal{X}_\lambda$. For a concrete group action $(\mathbb{G}, \mathcal{X}, *)$, we denote this standard-model game by $\mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}$.

Notice that a generic group action game cannot obliviously sample elements, since it is not given any interface to the group action other than the group action itself.

For any algorithm $\mathcal{A}$, we say the algorithm $\delta(\lambda)$-breaks $\mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}$ if $\mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}(1^\lambda)$ outputs 1 with probability at least $\delta(\lambda)$ when interacting with $\mathcal{A}$.

We say that Ch is one-round if it sends a single classical string to $\mathcal{A}$, and then receives a single quantum message from $\mathcal{A}$, before deciding if $\mathcal{A}$ wins.

Our modified KGEA assumption. We now give our modified KGEA assumption.

Assumption 3.16. The *quantum modified knowledge of group element assumption* (Q-mKGEA) holds on a group action $(\mathbb{G}, \mathcal{X}, *)$ if the following is true. Consider a one-round generic group action game Ch and any quantum polynomial time (QPT) adversary $\mathcal{A}$ that $1-\delta$-breaks $\mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}$ for a negligible δ. Write the message from $\mathcal{A}$ to $\mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}$ as $\rho_{1,2}$, as a joint system over two registers $1, 2$. Consider measuring the first register, to obtain a set element y. Denote this as $(y, |\psi\rangle) \leftarrow \mathcal{A}'(1^\lambda) \Leftrightarrow \mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}(1^\lambda)$. Then for all such $\delta, \mathcal{A}, \mathsf{Ch}$, there exists another negligible δ', a QPT $\mathcal{A}'$ that also $1-\delta'$-breaks $\mathsf{Ch}^(\mathbb{G}, \mathcal{X}, *)$, and moreover there exists a QPT extractor $\mathcal{E}$ and negligible ϵ such that

$$\Pr\left[y \in \mathcal{X} \wedge y \neq g * x_\lambda : \substack{(y,|\psi\rangle) \leftarrow \mathcal{A}'(1^\lambda) \Leftrightarrow \mathsf{Ch}^{(\mathbb{G},\mathcal{X},*)}(1^\lambda) \\ g \leftarrow \mathcal{E}(y,|\psi\rangle)}\right] \leq \epsilon(\lambda) \ .$$

Intuitively, this assumption says that if $\mathcal{A}$ wins some game, we might not be able to apply the KGEA extractor to it. However, there is some other $\mathcal{A}'$ that also wins the game, and that we *can* apply the KGEA extractor to.

Theorem 3.17. *Assuming Q-mKGEA (Assumption 3.16) and DLog/1-minCDH (Assumption 3.9) both hold on a group action $(\mathbb{G}, \mathcal{X}, *)$, then Construction 3.1 is a quantum lightning scheme.*

Proof. The proof is the same as the proof of Theorem 3.12, except that we observe that the quantum lightning experiment is a generic group action game. As such, after obtaining an adversary $\mathcal{A}$ that wins the quantum lightning game with probability $1 - \mathsf{negl}$, we immediately switch to the adversary $\mathcal{A}'$ that is assumed to exist by Assumption 3.16, and apply the extractor $\mathcal{E}$ to the output of $\mathcal{A}'$. □

4 A Construction for REGAs

In this section, we give a construction for the case where the group action can only be computed efficiently for a small “base” set of group elements. Such group actions are known as “restricted effective group actions” (REGAs).

4.1 Some additional background

Before giving the construction, we here provide some additional background that will be necessary for understanding the construction.

Groups. Let $\mathbb{G}$ be a group (written additively), and N an integer such that $N \times g = 0$ for all $g \in \mathbb{G}$. $N = |\mathbb{G}|$ will do. Then $\mathbb{G}$ is a subgroup of $\mathbb{Z}_N^n$ for some positive integer n. Let W be the set of vectors in $\mathbb{Z}_N^n$ such that $\mathbf{w} \cdot g = 0 \bmod N$ for all $g \in \mathbb{G}$. W is then a group, and we can therefore consider the group $(\mathbb{Z}_N^n)/W$ defined using the equivalence relation $\sim$, where $\mathbf{u}_1 \sim \mathbf{u}_2$ if $\mathbf{u}_1 - \mathbf{u}_2 \in W$. $(\mathbb{Z}_N^n)/W$ is isomorphic to $\mathbb{G}$; let $\phi : \mathbb{G} \to (\mathbb{Z}_N^n)/W$ be an isomorphism. Note that for $g \in \mathbb{G} \subseteq \mathbb{Z}_N^n$ and $h \in \mathbb{G}$, $g \cdot \phi(h) \bmod N$ is well-defined by taking any representative $h' \in \phi(h)$ and computing $g \cdot h' \bmod N$.

Under this notation, we can re-define $\chi(g, h)$ as $e^{i2\pi g\cdot\phi(h)/N}$, which is equivalent to the definition in Section 2.

We associate $\mathbb{Z}_N$ with the interval $[-\lfloor (N-1)/2 \rfloor, \lceil (N-1)/2 \rceil]$ in the obvious way, and likewise associate $\mathbb{Z}_N^n$ with the hypercube $[-\lfloor (N-1)/2 \rfloor, \lceil (N-1)/2 \rceil]^n$. This gives rise to a notion of norm on $\mathbb{Z}_N^n$ by taking the norm in $\mathbb{Z}^n$.

Lemma 4.1. *Let $\mathbb{G}$ be a subgroup of $\mathbb{Z}_N$. Then the number of elements $g \in \mathbb{G}$ such that $|g| \geq N/4$ is exactly $|\mathbb{G}| + 1 - 2\lceil |\mathbb{G}|/4 \rceil$. In particular, if $\mathbb{G} \neq \{0\}$, then there is at least one element $g \in \mathbb{G}$ has $|g| \geq N/4$.*

Proof. First, it suffices to consider $|\mathbb{G}| = N$, in other words $\mathbb{G} = \mathbb{Z}_N$: we can then lift to $N = t|\mathbb{G}|$, where $\mathbb{G}$ is embedded into $\mathbb{Z}_N$ by multiplying each element in $\mathbb{G}$ by t (where multiplication is over the integers). Since N is also multiplied by t, this preserves the number of elements with $|g| \geq N/4$.

When $\mathbb{G} = \mathbb{Z}_N$, we are then simply asking for the number of elements in $[-\lfloor (|\mathbb{G}|-1)/2 \rfloor, \lceil (|\mathbb{G}|-1)/2 \rceil]$ with absolute value at least $|\mathbb{G}|/4$. In other words, it is the combined size of the intervals $[\lceil |\mathbb{G}|/4 \rceil, \lceil (|\mathbb{G}|-1)/2 \rceil]$ and $[-\lfloor (|\mathbb{G}|-1)/2 \rfloor, -\lceil |\mathbb{G}|/4 \rceil]$, giving a total of $(\lceil (|\mathbb{G}|-1)/2 \rceil - \lceil |\mathbb{G}|/4 \rceil + 1) + (\lfloor (|\mathbb{G}|-1)/2 \lfloor - \lceil |\mathbb{G}|/4 \rceil + 1) = |\mathbb{G}| + 1 - 2\lceil |\mathbb{G}|/4 \rceil$. □

Lemma 4.2. *Let $\mathbf{A} \in \mathbb{Z}_N^{n\times m}$ be a matrix. Let $\mathbb{G}$ be the subgroup of $\mathbb{Z}_N^n$ generated by the columns of $\mathbf{A}$. Let B, C be positive integers such that $8BCm < N$. Suppose there is a distribution $\mathcal{D}$ on $[-B, B]^m$ such that $\mathbf{A} \cdot \mathbf{x}$ for $x \leftarrow \mathcal{D}$ is negligibly close to uniform in $\mathbb{G}$. Then the function $f : \mathbb{G} \times [-C, C] \to \mathbb{Z}_N^m$ given by $f(g, \mathbf{e}) = \mathbf{A}^T \cdot \phi(g) + \mathbf{e}$ is injective.*

Proof. Note that $\mathbf{A}^T \cdot \phi(g)$ is well defined since it is independent of the representative of $\phi(g)$. Consider a potential collision in f: $\mathbf{A}^T \cdot \phi(g_1) + \mathbf{e}_1 = \mathbf{A}^T \cdot \phi(g_2) + \mathbf{e}_2$. By subtracting, this gives a non-zero pair $(g = g_1 - g_2, \mathbf{e} = \mathbf{e}_1 - \mathbf{e}_2)$ where $\mathbf{e} \in [-2C, 2C]$ such that $\mathbf{A}^T \cdot \phi(g) + \mathbf{e} = 0$ or equivalently $\mathbf{A}^T \cdot \phi(g) = -\mathbf{e}$. Now consider sampling $\mathbf{x} \leftarrow \mathcal{D}$, meaning $\mathbf{u} = \mathbf{A} \cdot \mathbf{x}$ is negligibly close to uniform in $\mathbb{G}$. Then $\mathbf{u}^T \cdot \phi(g) = \mathbf{x}^T \cdot \mathbf{A}^T \cdot \phi(g) = -\mathbf{x}^T \cdot \mathbf{e}$. On one hand, $\mathbf{u}^T \cdot \phi(g)$ is statistically close to uniform in a subgroup $\mathbb{G}'$ of $\mathbb{Z}_N$, and $\mathbb{G}'$ is different from $\{0\}$ since $g \neq 0$. By Lemma 4.1, the probability $|\mathbf{u}^T \cdot \phi(g)| \geq N/4$ is $|\mathbb{G}'| + 1 - 2\lceil |\mathbb{G}'|/4 \rceil > 0$ since $|\mathbb{G}'| \geq 2$. On the other hand, $| - \mathbf{x}^T \cdot \mathbf{e}| < 2mBC \leq N/4$ always. This means the distributions of $\mathbf{u}^T \cdot \phi(g)$ and $-\mathbf{x}^T \cdot \mathbf{e}$ must be non-negligibly far, a contradiction. □

Discrete Gaussians. The *discrete Gaussian distribution* is the distribution over $\mathbb{Z}$ defined as:

$$\Pr[x] = \mathcal{D}_\sigma(x) := C_\sigma e^{2\pi x^2/\sigma^2},$$

where C_σ is the normalization constant $C_\sigma = \sum_{x \in \mathbb{Z}} e^{2\pi x^2/\sigma^2}$, so that $\mathcal{D}_\sigma$ defined a probability distribution. We will also define a truncated variant, denoted

$$\mathcal{D}_{\sigma,B}(x) := \begin{cases} C_{\sigma,B} e^{2\pi x^2/\sigma^2} & \text{if } |x| \leq B \\ 0 & \text{otherwise} \end{cases},$$

where again $C_{\sigma,B}$ is an appropriately defined normalization constant. For large B, we can treat the truncated and un-truncated Gaussians as essentially the same distribution:

Fact 4.3. *For $\sigma \geq \omega(\sqrt{\log \lambda})$ and $B \geq \sigma \times \omega(\sqrt{\log \lambda})$, the distributions $\mathcal{D}_\sigma$ and $\mathcal{D}_{\sigma,B}$ are negligibly close*

For a vector $\mathbf{r} \in \mathbb{Z}^m$, we write $\mathcal{D}_{\sigma,B}(\mathbf{r}) = \prod_{i=1}^m \mathcal{D}_{\sigma,B}(r_i)$.
The *discrete Gaussian superposition* is the quantum state

$$|\mathcal{D}_\sigma\rangle := \sum_{x \in \mathbb{Z}} \sqrt{\mathcal{D}_\sigma(x)} |x\rangle$$

As we will generally need to restrict to finite-precision, we also consider the truncated variant

$$|\mathcal{D}_{\sigma,B}\rangle := \sum_{x \in [-B,B]} \sqrt{\mathcal{D}_{\sigma,B}(x)} |x\rangle$$

Again, for large enough B, we can treat the truncated and un-truncated Gaussian superpositions as essentially the same state:

Fact 4.4. *For $\sigma \geq \omega(\sqrt{\log \lambda})$ and $B \geq \sigma \times \omega(\sqrt{\log \lambda})$, the $\| |\mathcal{D}_\sigma\rangle - |\mathcal{D}_{\sigma,B}\rangle \|$ is negligible.*

By adapting classicsal lattice sampling algorithms, the states $|\mathcal{D}_{\sigma,B}\rangle$ can be efficiently constructed.

Fourier transform pairs. Fix an integer N. We will associate the set $\mathbb{Z}_N$ with the integers $[-\lfloor (N-1)/2 \rfloor, \lceil (N-1)/2 \rceil]$. Denote by QFT_N the Quantum Fourier Transform $\mathsf{QFT}_{\mathbb{Z}_N}$. We now recall some basic facts about quantum Fourier transforms.

$$\mathsf{QFT}_N^m \sum_{\mathbf{r} \in \mathbb{Z}_N^m : \mathbf{A} \cdot \mathbf{r} = \mathbf{s}} |\mathbf{r}\rangle = N^{m/2-n} \sum_{\mathbf{t} \in \mathbb{Z}_N^n} e^{i2\pi \mathbf{t} \cdot \mathbf{s}/N} |\mathbf{A}^T \cdot \mathbf{t}\rangle \text{ for } \mathbf{A} \in \mathbb{Z}_N^{n \times m}$$

$$\mathsf{QFT}_N^m \sum_{\mathbf{r}} \alpha_{\mathbf{r}} \beta_{\mathbf{r}} |\mathbf{r}\rangle = \frac{1}{N^{m/2}} \sum_{\mathbf{t},\mathbf{u}} \hat{\alpha}_{\mathbf{t}} \hat{\beta}_{\mathbf{u}} |\mathbf{u} + \mathbf{t}\rangle \text{ for } {\scriptstyle \sum_{\mathbf{t}} \hat{\alpha}_{\mathbf{t}} |\mathbf{t}\rangle = \mathsf{QFT}_N^m \sum_{\mathbf{r}} \alpha_{\mathbf{r}} |\mathbf{r}\rangle \atop \sum_{\mathbf{u}} \hat{\beta}_{\mathbf{u}} |\mathbf{u}\rangle = \mathsf{QFT}_N^m \sum_{\mathbf{r}} \beta_{\mathbf{r}} |\mathbf{r}\rangle}$$

$$\mathsf{QFT}_N |\mathcal{D}_{\sigma, \lfloor (N-1)/2 \rfloor}\rangle \approx |\mathcal{D}_{N/\sigma, \lfloor (N-1)/2 \rfloor}\rangle \text{ for } {\scriptstyle N \geq \sigma \times \omega(\sqrt{\log \lambda}) \atop \sigma \geq \omega(\sqrt{\log \lambda})}$$

Above, $\approx$ means the two states are negligibly close.

4.2 The Construction

Let $\mathbb{G}_\lambda, \mathcal{X}_\lambda, *$ be a REGA, and $\mathcal{T} = (g_1, \ldots, g_m)$ a set such that $*$ can be efficiently computed for g_i and g_i^{-1}. We can associate $\mathbb{G}_\lambda$ with a subgroup of $\mathbb{Z}_N^n$ for some integers N, n. We can likewise associate the list $\mathcal{T}$ with the matrix $\mathbf{A} = (g_1, \cdots, g_m) \in \mathbb{Z}_N^{n \times m}$.

We will make the following assumption about the structure of $\mathcal{T}$, which is typical in the isogeny literature.

Assumption 4.5. There is a polynomial B and a distribution $\mathcal{D}^*$ on $[-B, B]^m$ such that for $\mathbf{x} \leftarrow \mathcal{D}$, $\sum_{i=1}^m x_i g_i = \mathbf{A} \cdot \mathbf{x}$ is statistically close to a uniform element in $\mathbb{G}$

Numerous examples of such $\mathcal{D}^*$ have been proposed, such as discrete Gaussians [DG19], or uniform vectors in small balls relative to different norms [CLM+18, NOTT20].

Let $C = N/8Bm$, which then satisfies the conditions of Lemma 4.2. Thus, for $\mathbf{e}$ with entries in $[-C, C]^m$, the map $(g, \mathbf{e}) \mapsto \mathbf{A}^T \cdot \phi(g) + \mathbf{e}$ is injective.

Let $\sigma \geq 16Bm/\epsilon \times \omega(\sqrt{\log \lambda})$ and $B' \geq \sigma \times \omega(\sqrt{\log \lambda})$ be polynomials. We will assume $N \geq 2B'$, which is always possible since we can take N to be arbitrarily large. We will also for simplicity assume N is even. This assumption is not necessary but will simplify some of the analysis, and is moreover without loss of generality since we can always make N larger by multiplying it by arbitrary factors.

Construction 4.6. $\mathsf{Gen}(1^\lambda)$: Initialize quantum registers $\mathcal{S}$ (for serial number) and $\mathcal{M}$ (for money) to states $|\mathcal{D}_{\sigma,B'}\rangle_{\mathcal{S}}^{\otimes m}$ and $|0\rangle_{\mathcal{M}}$, respectively. Then do the following:

- Apply in superposition the map $|\mathbf{r}\rangle_{\mathcal{S}}|y\rangle_{\mathcal{M}} \mapsto |\mathbf{r}\rangle_{\mathcal{S}}|y \oplus [(\sum_{i=1}^m r_i g_i) * x_\lambda])\rangle_{\mathcal{M}}$. The joint state of the system $\mathcal{S} \otimes \mathcal{M}$ is then

$$\sum_{\mathbf{r} \in \mathbb{Z}_N^m} \sqrt{\mathcal{D}_{\sigma',B}(\mathbf{r})}|\mathbf{r}\rangle_{\mathcal{S}}|(\sum_{i=1}^m r_i g_i) * x_\lambda\rangle_{\mathcal{M}} = \sum_{g \in \mathbb{G}_\lambda} \left(\sum_{\mathbf{r} \in \mathbb{Z}_N^m : \mathbf{A} \cdot \mathbf{r} = g} \sqrt{\mathcal{D}_{\sigma,B'}(\mathbf{r})}|\mathbf{r}\rangle_{\mathcal{S}} \right) |g * x_\lambda\rangle_{\mathcal{M}}$$

- Apply $\mathsf{QFT}_{\mathbb{Z}_N^m}$ to $\mathcal{S}$. Using the QFT rules given above, this yields the state negligibly close to:

$$\begin{aligned} &\frac{1}{N^n} \sum_{g \in \mathbb{G}_\lambda} \left(\sum_{\mathbf{s}, \mathbf{e} \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{N/\sigma, N/2-1}(\mathbf{e})} e^{i2\pi(g \cdot \mathbf{s})} |\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}\rangle_{\mathcal{S}} \right) |g * x_\lambda\rangle_{\mathcal{M}} \\ &= \frac{1}{|\mathbb{G}_\lambda|} \sum_{g \in \mathbb{G}_\lambda} \left(\sum_{h \in \mathbb{G}_\lambda, \mathbf{e} \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{N/\sigma, N/2-1}(\mathbf{e})} e^{i2\pi(g \cdot \phi(h))} |\mathbf{A}^T \cdot \phi(h) + \mathbf{e}\rangle_{\mathcal{S}} \right) |g * x_\lambda\rangle_{\mathcal{M}} \\ &= \frac{1}{\sqrt{|\mathbb{G}_\lambda|}} \sum_{g \in \mathbb{G}_\lambda} \left(\frac{1}{\sqrt{|\mathbb{G}_\lambda|}} \sum_{h \in \mathbb{G}_\lambda, \mathbf{e} \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{N/\sigma, N/2-1}(\mathbf{e})} \chi(g, h) |\mathbf{A}^T \cdot \phi(h) + \mathbf{e}\rangle_{\mathcal{S}} \right) |g * x_\lambda\rangle_{\mathcal{M}} \end{aligned}$$

- Measure $\mathcal{S}$, giving the serial number $\mathbf{t} := \mathbf{A}^T \cdot \phi(h) + \mathbf{e}$. $\mathbf{e}$ is distributed negligibly close to $\mathcal{D}_{N/\sigma}$, meaning with overwhelming probability each entry is in $[-N/16Bm, N/16Bm] = [-C/2, C/2] \subseteq [-C, C]$. This means, to within negligible error, $\mathbf{t}$ uniquely determines $\phi(h)$ and hence h. Therefore, the $\mathcal{M}$ register then collapses to a state negligibly close to

$$\frac{1}{\sqrt{|\mathbb{G}_\lambda|}} \sum_{g \in \mathbb{G}_\lambda} \chi(g, h)|g * x_\lambda\rangle_{\mathcal{M}} =: |\mathbb{G}_\lambda^h * x_\lambda\rangle$$

Note that h is unknown. Output $(\mathbf{t}, |\mathbb{G}_\lambda^h * x_\lambda\rangle)$

$\mathsf{Ver}(\mathbf{t}, \$)$: First verify that the support of \$ is contained in $\mathcal{X}_\lambda$, by applying the assumed algorithm for recognizing $\mathcal{X}_\lambda$ in superposition. Then repeat the following λ times:

- Initialize a new register $\mathcal{H}$ to $(|0\rangle_{\mathcal{H}} + |1\rangle_{\mathcal{H}})/\sqrt{2}$.
- Choose a random element $\mathbf{x} \leftarrow \mathcal{D}^*$.
- Apply to $\mathcal{H} \otimes \mathcal{M}$ in superposition the map

$$\mathsf{Apply}|b\rangle_{\mathcal{H}}|y\rangle_{\mathcal{M}} \mapsto \begin{cases} |0\rangle_{\mathcal{H}}|y\rangle_{\mathcal{M}} & \text{if } b = 0 \\ |1\rangle_{\mathcal{H}}|(-\sum_i x_i g_i) * y\rangle_{\mathcal{M}} & \text{if } b = 1 \end{cases}$$

  Since the entries of $\mathbf{x}$ are bounded by B which is polynomial, this step is efficient.

- Measure $\mathcal{H}$ in the basis $B_{\mathbf{t},\mathbf{x}} := \{(|0\rangle_{\mathcal{H}} + e^{i2\pi \mathbf{x}^T \cdot \mathbf{t}/N}|1\rangle_{\mathcal{H}})/\sqrt{2}, (|0\rangle_{\mathcal{H}} - e^{i2\pi \mathbf{x}^T \cdot \mathbf{t}/N}|1\rangle_{\mathcal{H}})/\sqrt{2}\}$, giving a bit $b_u \in \{0,1\}$. Discard the $\mathcal{H}$ register.
- Accept if at least a fraction $7/8$ of the $b_u = 0$ and the support of \$ is contained in $\mathcal{X}_\lambda$; otherwise reject.

4.3 Accepting States of the Verifier

We now analyze the correctness of the construction.

Theorem 4.7. *Let $|\psi\rangle$ be a state over $\mathcal{M}$. Then* $\Pr[\mathsf{Ver}(h, |\psi\rangle) = 1] = \|\langle\psi|\mathbb{G}_\lambda^h * x_\lambda\rangle\|^2(1 - 2^{-\Omega(\sqrt{\lambda})} + 2^{-\Omega(\sqrt{\lambda})}$.

Proof. For simplicity, we analyze the case of $|\psi\rangle = |\mathbb{G}_\lambda^{h'} * x_\lambda$, which form a basis for superpositions over $\mathcal{X}_\lambda$. In this case, Theorem 4.7 states that $|\mathbb{G}_\lambda^h * x_\lambda\rangle$ is accepted with probability $1 - 2^{\Omega(\sqrt{\lambda})}$, while $|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle$ for $h' \neq h$ is accepted with probability $2^{\Omega(\sqrt{\lambda})}$. By a similar approach as in Theorem 3.3, we can extend the analysis to all states.

If we let $u = \mathbf{A} \cdot \mathbf{x} = \sum_i x_i g_i$, then by the same analysis as in Construction 3.1, we have that applying Apply to the state $|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle$ results in the state

$$\begin{aligned} &\frac{1}{\sqrt{2}}\left(|0\rangle_{\mathcal{H}} + \chi(u, h')|1\rangle_{\mathcal{H}}\right)|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle \\ &= \frac{1}{\sqrt{2}}\left(|0\rangle_{\mathcal{H}} + e^{i2\pi u \cdot \phi(h')/N}|1\rangle_{\mathcal{H}}\right)|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle \\ &= \frac{1}{\sqrt{2}}\left(|0\rangle_{\mathcal{H}} + e^{i2\pi \mathbf{x}^T \cdot \mathbf{A}^T \cdot \phi(h')/N}|1\rangle_{\mathcal{H}}\right)|\mathbb{G}_\lambda^{h'} * x_\lambda\rangle \end{aligned}$$

Conditioned on sampling u, $\Pr[b_u = 0]$ is the inner product squared of $\left(|0\rangle_{\mathcal{H}} + e^{i2\pi\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h')/N}|1\rangle_{\mathcal{H}}\right)/\sqrt{2}$ with the basis state $\left(|0\rangle_{\mathcal{H}} + e^{i2\pi\mathbf{x}\cdot\mathbf{t}/N}|1\rangle_{\mathcal{H}}\right)/\sqrt{2}$. This is:

$$\begin{aligned}\Pr[b_u = 0] &= \frac{1}{4}\left\|1 + e^{i2\pi(\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h') - \mathbf{x}^T\cdot\mathbf{t})/N}\right\|^2\\ &= \frac{1}{2}\left(1 + \cos\left[2\pi(\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h') - \mathbf{x}^T\cdot(\mathbf{A}^T\cdot\phi(h) + \mathbf{e}))/N\right]\right)\\ &= \frac{1}{2}\left(1 + \cos\left[2\pi(\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h' - h) + \mathbf{x}^T\cdot\mathbf{e})/N\right]\right)\end{aligned}$$

In the case $h = h'$, $\Pr[b_u = 0] = \frac{1}{2}\left(1 + \cos\left[2\pi\mathbf{x}^T\cdot\mathbf{e}/N\right]\right)$. We have that $|2\pi\mathbf{x}^T\cdot\mathbf{e}/N| \leq \pi/8$. Using the fact that $\cos(x) \geq 1 - x^2/2$, we therefore have that $\Pr[b_u = 0] \geq 1 - \pi^2/256 = 0.9614\ldots = 7/8 + \Omega(1)$. Then via standard concentration inequalities, after λ trials, except with probability $2^{-\Omega(\sqrt{\lambda})}$, at least $7/8$ of the b_u will be 0. Therefore, Ver accepts with probability $1 - 2^{-\Omega(\sqrt{\lambda})}$.

On the other hand, if $g \neq g'$, then $\mathbf{x}^T\mathbf{A}^T$ is statistically close to uniform in $\mathbb{G}_\lambda$, and so $\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h' - h)$ is statistically close to uniform in a non-trivial subgroup $\mathbb{G}'$ of $\mathbb{Z}_N$. By Lemma 4.1 and our assumption that N is even, at least half of the elements of $\mathbb{Z}_N$ are at least $N/4$ in absolute value. In particular, this means $\Pr[|\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h' - h)| \geq N/4] \geq 1/2 - \mathsf{negl}$. On the other hand, $|\mathbf{x}^T\cdot\mathbf{e}| \leq N/16$ always. This means $\|\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h' - h) + \mathbf{x}^T\cdot\mathbf{e}\| \geq N/4 - N/16$ with probability at least $1/2 - \mathsf{negl}$. In this case, we can use that $\cos(\pi/2 + x) \leq |x|$ to bound $\cos\left[2\pi(\mathbf{x}^T\cdot\mathbf{A}^T\cdot\phi(h' - h) + \mathbf{x}^T\cdot\mathbf{e})/N\right] \leq 2\pi/16 = \pi/8$, meaning $\Pr[b_u = 0] \leq 1/2 + \pi/16$. Averaging over all u, we therefore have that: $\Pr[b_u = 0] \leq \frac{3}{4} + \pi/32 + \mathsf{negl} = 0.8481\ldots = 7/8 - \Omega(1)$. Then via standard concentration inequalities, after λ trials, except with probability $2^{-\Omega(\sqrt{\lambda})}$, fewer than $7/8$ of the b_u will be 0. therefore, Ver accepts with probability $2^{-\Omega(\sqrt{\lambda})}$. □

4.4 Security

Here, we state the security of Construction 4.6.

Assumptions. We first need to define slight variants of our assumptions, in order to be consistent with the more limited structure of a REGA. For example, in the ordinary Discrete Log assumption (Assumption 2.4), the challenger computes $y = g * x$ for a random g, and adversary produces g. But the adversary cannot even tell if it succeeded since it cannot compute the action of g in general. Instead, the adversary is required not to compute g, but instead to compute any short $\mathbf{x}$ such that $g = \sum_i x_i g_i$. The adversary can then check that it has a solution by computing the action of g using its knowledge of $\mathbf{x}$. We analogously update each of our assumptions to work with the limited ability to compute the group action on REGAs.

As above, let $\mathbb{G}_\lambda, \mathcal{X}_\lambda, *$ be a REGA, and $\mathcal{T} = (g_1, \ldots, g_m)$ a set such that $*$ can be efficiently computed for g_i and g_i^{-1}. Let $\mathcal{D}^*, B$ be as in Assumption 4.5.

Assumption 4.8. The *REGA quantum knowledge of group element assumption* (REGA-Q-KGEA) holds on a group action $(\mathbb{G}, \mathcal{X}, *)$ if the following is true. For any quantum polynomial time (QPT) adversary $\mathcal{A}$ which performs no measurements except for its final output, there exists a polynomial C, a QPT extractor $\mathcal{E}$ with outputs in $[-C, C]^m$, and negligible ϵ such that

$$\Pr\left[y \in \mathcal{X} \wedge y \neq g * x_\lambda : \substack{(y,|\psi\rangle)\leftarrow\mathcal{A}(1^\lambda)\\ \mathbf{x}\leftarrow\mathcal{E}(y,|\psi\rangle)} g \leftarrow \sum_i x_i g_i\right] \leq \epsilon(\lambda) \ .$$

We can likewise define a *modified* REGA KGEA assumption (REGA-Q-mKGEA), in the same spirit as Assumption 3.16.

Assumption 4.9. We say that the *REGA Discrete Log with a single minimal CDH query* assumption (REGA-DLog/1-minCDH) assumption holds if the following is true. For any QPT adversary $\mathcal{A}$ playing the following game, parameterized by λ, there is a negligible ϵ such that $\mathcal{A}$ wins with probability at most $\epsilon(\lambda)$:

- The challenger, on input λ, chooses a random $g \in \mathbb{G}_\lambda$. It sends λ to $\mathcal{A}$
- $\mathcal{A}$ submits a superposition query $\sum_{y \in \mathcal{X}, z \in \{0,1\}^*} \alpha_{y,z} |y, z\rangle$. Here, y is a set element that forms the query, and z is the internal state of the adversary when making the query. The challenger responds with $\sum_{y \in \mathcal{X}, z \in \{0,1\}^*} \alpha_{y,z} |(-g) * y, z\rangle$.
- The challenger sends $g * x$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs a $\mathbf{x} \in \mathbb{Z}^m$, encoded in unary. It wins if $g = \sum_i x_i g_i$.

Note that the challenger in Assumption 4.9 is inefficient on a REGA. However, under Assumption 4.5, the challenger can be made efficient by first sampling $\mathbf{y} \leftarrow \mathcal{D}^*$ and then computing $g = \sum_i y_i g_i$.

Theorem 4.10. *Assuming REGA-DLog/1-minCDH (Assumption 4.9) and REGA-Q-KGEA (Assumption 4.8) (or more generally, REGA-Q-mKGEA) both hold on a group action* $(\mathbb{G}, \mathcal{X}, *)$*, then Construction 4.6 is a quantum lightning scheme.*

We only sketch the proof. Like in the proof of Theorems 3.12 and 3.17, we can assume the adversary wins the quantum lightning experiment with probability $1 - \mathsf{negl}(\lambda)$. In order for a supposed note \$ to be accepted relative to serial number $\mathbf{t}$ with overwhelming probability, $\mathbf{t}$ must have the form $\mathbf{t} = \mathbf{A}^T \cdot \phi(h) + \mathbf{e}$ for "short" $\mathbf{e}$, and \$ must be negligibly close to $|\mathbb{G}_\lambda^h * x_\lambda\rangle$. Therefore, a quantum lightning adversary outputs two copies of $|\mathbb{G}_\lambda^h * x_\lambda\rangle$ for some h. The security reduction of Theorem 3.12 did not rely on knowing h, just that the adversary outputted two copies of $|\mathbb{G}_\lambda^h * x_\lambda\rangle$. Hence, a near-identical proof holds for Construction 4.6. The only difference is that when the extractor $\mathcal{E}$ outputs a group element, it instead outputs a small linear combination of the g_i giving that group element, and then the DLog/1-minCDH adversary uses this small representation to compute the action by that group element.

5 Further Discussion

5.1 Quantum Group Actions

Here, we consider a generalization of group actions where set elements are replaced with quantum states.

An quantum (abelian) group action consists of a family of (abelian) groups $\mathbb{G} = (\mathbb{G}_\lambda)_\lambda$ (written additively), a family $\mathcal{X} = (\mathcal{X}_\lambda)_\lambda$ of sets $\mathcal{X}_\lambda$ of quantum states over a system $\mathcal{M}_\lambda$, and an operation $*$. We will require that the states in $\mathcal{X}_\lambda$ are orthogonal. $*$ is a quantum algorithm that takes as input a group element $g \in \mathbb{G}_\lambda$ and a quantum state $|\psi\rangle$ over $\mathcal{M}_\lambda$, and outputs another state over $\mathcal{M}_\lambda$. $*$ satisfies the following properties:

- **Identity:** If $0 \in \mathbb{G}_\lambda$ is the identity element, then $|0\rangle * |\psi\rangle = |\psi\rangle$ for any $|\psi\rangle \in \mathcal{X}_\lambda$.
- **Compatibility:** For all $g, h \in \mathbb{G}_\lambda$ and $|\psi\rangle \in \mathcal{X}_\lambda$, $(g + h) * |\psi\rangle = g * (h * |\psi\rangle)$.

We can also relax the above properties to only hold to within negligible error, and/or relax the orthogonality requirement to being near-orthogonal. We will additionally require the following properties:

- **Efficiently computable:** There is a pseudodeterministic QPT procedure $\mathsf{Construct}$ which, on input 1^λ, outputs a description of $\mathbb{G}_\lambda$ and an specific element $|\psi_\lambda\rangle \in \mathcal{X}_\lambda$. The operation $*$ is also computable by a QPT algorithm.
- **Efficiently Recognizable:** There is a QPT procedure Recog which recognizes elements in $\mathcal{X}_\lambda$. That is, $\mathsf{Recog}(1^\lambda, \cdot)$ projects onto the span of the states in $\mathcal{X}_\rangle$.
- **Regular:** For every $|\phi\rangle \in \mathcal{X}_\lambda$, there is exactly one $g \in \mathbb{G}_\lambda$ such that $|\phi\rangle = g * |\psi_\lambda\rangle$.

Again, we can also relax the above properties to only hold to within negligible error.

Cryptographic group actions. At a minimum, a cryptographically useful quantum group action will satisfy the following discrete log assumption:

Assumption 5.1. The *discrete log assumption* (DLog) holds on a quantum group action $(\mathbb{G}, \mathcal{X}, *)$ if, for all QPT adversaries $\mathcal{A}$, there exists a negligible λ such that

$$\Pr[\mathcal{A}(g * |\psi_\lambda\rangle) = g : g \leftarrow \mathbb{G}_\lambda] \leq \mathsf{negl}(\lambda) \ .$$

Note that if we do not insist on orthogonality of the states in $\mathcal{X}_\lambda$, then it is trivial to construct a quantum group action in which DLog holds: simply have all $|\psi\rangle \in \mathcal{X}_\lambda$ be identical, or negligibly close. Then it will be information-theoretically impossible to determine g. Orthogonality essentially says that the group action is classical, except that the basis for the set elements is potentially different than the computational basis.

5.2 Quantum Group Actions From Lattices

Here, we describe a simple quantum group action from lattices.

The group $\mathbb{G}_{\mathsf{LWE,N,n,m},\sigma}$ will be set to $\mathbb{Z}_N^n$ for some integers N, n. We will fix a short wide matrix $\mathbf{A} \in \mathbb{Z}_N^{n\times m}$; we can think of $\mathbf{A}$ as being sampled randomly and included in a common reference string. Note that $\mathbb{G}$ is independent of σ, but we include it for notational consistency.

The set $\mathcal{X}_{\mathsf{LWE,N,n,m},\sigma}$ will be the set of states $|\psi_{\mathbf{s}}\rangle = \sum_{\mathbf{e}\in\mathbb{Z}_N^n} \sqrt{\mathcal{D}_{\sigma,N/2}(\mathbf{e})}|\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}\rangle$. In other words, we take the discrete Gaussian vector superposition of some width, and add the vector $\mathbf{A}^T \cdot \mathbf{s}$.

$\mathbb{G}_{\mathsf{LWE,N,n,m},\sigma}$ acts on $\mathcal{X}_{\mathsf{LWE,N,n,m},\sigma}$ in the following obvious way: $\mathbf{r} * |\psi_{\mathbf{s}}\rangle = |\psi_{\mathbf{r}+\mathbf{s}}\rangle$, which can be computed by simply adding $\mathbf{A}^T \cdot \mathbf{r}$ in superposition.

We have the following theorem:

Theorem 5.2. *Let σ, σ_0 be non-negative real numbers such that σ/σ_0 is super-logarithmic. Assuming the Learning with Errors problem is hard for noise distribution $\mathcal{D}_{\sigma_0}$, discrete logarithms are hard in the group action $(\mathbb{G}_{\mathsf{LWE,N,n,m},\sigma}, \mathcal{X}_{\mathsf{LWE,N,n,m},\sigma}, *)$.*

Proof. The learning with errors assumption states that it is hard to compute $\mathbf{s}$ given $\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}$ with $\mathbf{e}$ sampled from $\mathcal{D}_{\sigma_0}$. We need to show that it is hard to compute $\mathbf{s}$ given the analogous superposition over $\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}$, where here $\mathbf{e}$ comes from the Gaussian superposition $|\mathcal{D}_\sigma\rangle$. The idea is a simple application of noise flooding: given $\mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}$, compute the state $|\psi'_\mathbf{s}\rangle := \sum_{\mathbf{e}' \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{\sigma,N/2}(\mathbf{e}')} |\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} + \mathbf{e}'\rangle$. Since σ/σ_0 is super-polynomial, $\mathbf{e} + \mathbf{e}'$ where $\mathbf{e}' \leftarrow \mathcal{D}_{\sigma,N/2}$ is negligibly close to a Gaussian centered at 0. Therefore, $|\psi'_\mathbf{s}\rangle$ is negligibly close to $|\psi_\mathbf{s}\rangle$. Plugging into a supposed DLog adversary then gives $\mathbf{s}$, breaking LWE. □

Unfortunately, this LWE-based group action is missing a crucial feature: it is not possible to recognize states in $\mathcal{X}$. In particular, the states in $\mathcal{X}$ are indistinguishable from states of the form $\sum_{\mathbf{e} \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{\sigma,N/2}(\mathbf{e})} |\mathbf{v} + \mathbf{e}\rangle$, where $\mathbf{v}$ is an arbitrary vector in $\mathbb{Z}_N^m$. As we will see in the next sub-section, the inability to recognize $\mathcal{X}$ will prevent us from using this group action to instantiate our quantum money scheme.

5.3 Relation to Quantum Money Approaches based on Lattices

Here, we see that our quantum money scheme is conceptually related to a folklore approach to building quantum money from lattices. This approach has not been able to work; in our language, the reason is exactly due to the inability to recognize $\mathcal{X}_{\mathsf{LWE},\mathsf{N},\mathsf{n},\mathsf{m},\sigma}$.

The approach is the following. Let $\mathbf{A} \in \mathbb{Z}_N^{n \times m}$ be a random short wide matrix over $\mathbb{Z}_n$. To mint a banknote, construct the discrete Gaussian superposition $|\mathcal{D}_\sigma\rangle^{\otimes m}$ in register $\mathcal{M}$. Then compute and measure $\mathbf{A} \cdot \mathbf{x}$ applied to $\mathcal{M}$. The result is a vector $\mathbf{h} \in \mathbb{Z}_N^n$, which will be the serial number, and $\mathcal{M}$ collapses to a superposition $|\$_\mathbf{h}\rangle \propto \sum_{\mathbf{x}:\mathbf{A}\cdot\mathbf{x}=\mathbf{h}} \sqrt{\mathcal{D}_\sigma(\mathbf{x})} |\mathbf{x}\rangle$ of short vectors $\mathbf{x}$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{h}$. This is the banknote. A simple argument shows that it is impossible to construct two copies of $|\$_\mathbf{h}\rangle$ for the same $\mathbf{h}$: given such a pair, measure each to get $\mathbf{x}, \mathbf{x}'$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{A} \cdot \mathbf{x}' = \mathbf{h}$. Then subtract to get a short vector $\mathbf{x} - \mathbf{x}'$ such that $\mathbf{A} \cdot (\mathbf{x} - \mathbf{x}') = 0^n$. We can conclude $\mathbf{x} - \mathbf{x}'$ is non-zero with overwhelming probability, since the measurement of $|\$_\mathbf{h}\rangle$ has high entropy. Such a non-zero short kernel vector would solve the Short Integer Solution (SIS) problem, which is widely believed to be hard and is the foundation of lattice-based cryptography.

Unfortunately, the above approach is broken. The problem is that there is no way to actually verify banknotes. One can verify that a banknote has support on short vectors with $\mathbf{A} \cdot \mathbf{x} = \mathbf{h}$, but it is impossible to verify that the banknote is in superposition. If one could solve the Learning with Errors (LWE) problem, it would be possible to verify banknotes as follows: first perform the QFT to the banknote state. If an honest banknote, the QFT will give a state negligibly close to

$$|\$'_\mathbf{h}\rangle := \frac{1}{N^{n/2}} \sum_{\mathbf{s},\mathbf{e} \in \mathbb{Z}_N^n} \sqrt{\mathcal{D}_{N/\sigma}(\mathbf{e})} e^{i2\pi \mathbf{h}\cdot\mathbf{s}/N} |\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}\rangle \ . \tag{5.1}$$

The second step is to simply apply the supposed LWE solver to this state in superposition, ensuring that the state has support on vectors of the form $\mathbf{A}^T \cdot \mathbf{s} + \mathbf{e}$ for small $\mathbf{e}$.

Unfortunately, LWE is likely hard. In fact, it is quantumly equivalent to SIS [Reg05], meaning if one could verify banknotes using an LWE solver, then SIS is easy. Not only does this mean we are reducing from an easy problem, but it would be possible to turn such a SIS algorithm into an attack.

Without the ability to verify that banknotes are in supeprosition, the attacker can simply measure a banknote to get $\mathbf{x}$, and then pass off $|\mathbf{x}\rangle$ as a fake banknote that will pass verification. Since $\mathbf{x}$ is trivially copied, this would break security. Interestingly, [LZ19] prove that, no matter what efficient verification procedure is used, even if the verification diverged from the LWE-based approach above, this attack works. [LMZ23] extend this to a variety of potential schemes based on similar ideas, including a recent proposed instantiation of this approach by [KLS22].

We now see how the above approach is essentially equivalent to our construction of quantum money from group actions, instantiated over our LWE-based quantum group action. The inability to recognize $\mathcal{X}$ is the reason this instantiation is insecure, despite natural hardness assumptions presumably holding on the group action.

We consider the quantum group action $(\mathbb{G}_{\mathsf{LWE,N,n,m,N}/\sigma}, \mathcal{X}_{\mathsf{LWE,N,n,m,N}/\sigma}, *)$, where σ is from the folklore construction above. When applied to $(\mathbb{G}_{\mathsf{LWE,N,n,m,N}/\sigma}, \mathcal{X}_{\mathsf{LWE,N,n,m,N}/\sigma}, *)$, a banknote in our scheme, up to negligibly error from truncating discrete Gaussians, is the state $|\$'_{\mathbf{h}}\rangle$ from Equation 5.1 above, where the serial number is $\mathbf{h}$. Thus, we see that our quantum money scheme is simply the folklore construction but moved to the Fourier domain. The attack on the folklore construction can therefore easily be mapped to an attack on our scheme: if the adversary is given $|\$'_{\mathbf{h}}\rangle$, it measures in the Fourier domain (which is the primal domain for the folklore construction) to get a short vector $\mathbf{x}$ such that $\mathbf{A} \cdot \mathbf{s} = \mathbf{h}$. Then it switched back to the primal domain, giving the state

$$\frac{1}{N^{m/2}} \sum_{\mathbf{u}} e^{i2\pi \mathbf{e} \cdot \mathbf{x}} |\mathbf{x}\rangle$$

This is a state that lies outside the span of $\mathcal{X}$. However, no efficient verification procedure can distinguish it from an honest banknote state.

Two features that distinguish isogeny-based group actions from the LWE-based action above. The first is the ability to recognize elements in $\mathcal{X}$. Suppose it were possible to recognize elements of $\mathcal{X}$ in the LWE-based action, and we had the verifier check to see if the banknote belonged to the span of the elements in $\mathcal{X}$. In the language of quantum group actions, this check would prevent the attacker from sending $\frac{1}{N^{m/2}} \sum_{\mathbf{u}} e^{i2\pi \mathbf{e} \cdot \mathbf{x}} |\mathbf{x}\rangle$, which lies outside the span of $\mathcal{X}$. In the language of the folklore construction, this check would correctly distinguish between an honest banknote and the easily clonable state $|\mathbf{x}\rangle$ in the attack. If such a check were possible, the proof sketched above would work to base the security of the scheme on SIS. Unfortunately, such a check is computationally intractable under the decision LWE problem, which is equivalent to SIS and most likely hard.

The issue of recognizing set elements is also crucial in our security arguments. Indeed, the first step in our proof was to characterize the states accepted by the verifier, showing that only honest banknote states are accepted. This step in the proof fails in the LWE-based scheme, which would prevent the proof from going through. Thus, even though the scheme based on LWE is broken, it does not contradict our DLog/1-minCDH and Q-KGEA assumptions holding on the LWE-based group action.

The second difference, is that, with the LWE-based group action, taking the QFT of money states gives elements with meaningful structure: short vectors $\mathbf{x}$ such that $\mathbf{A} \cdot \mathbf{x} = \mathbf{h}$. This structure and it's relation to the original money state are what enables the attack. In contrast, taking the QFT of money states over $\mathcal{X}$ coming from isognies will give terms with no discernible structure.

We believe the above perspective adds to the confidence in our proposal. Indeed, in the LWE-based scheme, the key missing piece is recognizing set elements; if not for this missing piece the scheme *could* be proven secure. By switching to group actions based on isogenies, we add the missing piece. The hope is that even though the source of hardness is now from hard problems on isogenies over elliptic curves instead of lattices, by adding the missing piece we can finally obtain a secure scheme.

References

[Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity*, CCC '09, pages 229–242, Washington, DC, USA, 2009. IEEE Computer Society.

[AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 41–60. ACM Press, May 2012.

[ADMP20] Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 411–439. Springer, Heidelberg, December 2020.

[AMR22] Navid Alamati, Giulio Malavolta, and Ahmadreza Rahimi. Candidate trapdoor claw-free functions from group actions with applications to quantum protocols. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 266–293. Springer, Heidelberg, November 2022.

[BCM$^+$18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th FOCS*, pages 320–331. IEEE Computer Society Press, October 2018.

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for iO: Circular-secure LWE suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. https://eprint.iacr.org/2020/1024.

[BDS16] Shalev Ben-David and Or Sattath. Quantum tokens for digital signatures, 2016. https://arxiv.org/abs/1609.09047.

[BGMZ18] James Bartusek, Jiaxin Guan, Fermi Ma, and Mark Zhandry. Return of GGH15: Provable security against zeroizing attacks. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 544–574. Springer, Heidelberg, November 2018.

[BGZ23] Dan Boneh, Jiaxin Guan, and Mark Zhandry. A lower bound on the length of signatures based on group actions and generic isogenies. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 507–531. Springer, Heidelberg, April 2023.

[BKV19] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 227–247. Springer, Heidelberg, December 2019.

[BS20a] Amit Behera and Or Sattath. Almost public quantum coins. Cryptology ePrint Archive, Report 2020/452, 2020. https://eprint.iacr.org/2020/452.

[BS20b] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 493–522. Springer, Heidelberg, May 2020.

[CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 423–447. Springer, Heidelberg, April 2023.

[CJS14] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, 2014.

[CK20] Leonardo Colò and David Kohel. Orienting supersingular isogeny graphs. *Journal of Mathematical Cryptology*, 14:414–437, 10 2020.

[CLM$^+$18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 395–427. Springer, Heidelberg, December 2018.

[Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291, 2006. https://eprint.iacr.org/2006/291.

[CPDDF$^+$19] Marta Conde Pena, Raul Durán Díaz, Jean-Charles Faugère, Luis Hernández Encinas, and Ludovic Perret. Non-quantum cryptanalysis of the noisy version of aaronson–christiano's quantum money scheme. *IET Information Security*, 13(4):362–366, 2019.

[DFK$^+$23] Luca De Feo, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Simon-Philipp Merz, Lorenz Panny, and Benjamin Wesolowski. SCALLOP: Scaling the CSI-FiSh. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 345–375. Springer, Heidelberg, May 2023.

[DG19] Luca De Feo and Steven D. Galbraith. SeaSign: Compact isogeny signatures from class group actions. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 759–789. Springer, Heidelberg, May 2019.

[DJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014.

[DM20] Luca De Feo and Michael Meyer. Threshold schemes from isogeny assumptions. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 187–212. Springer, Heidelberg, May 2020.

[FGH$^+$12] Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, and Peter W. Shor. Quantum money from knots. In Shafi Goldwasser, editor, *ITCS 2012*, pages 276–289. ACM, January 2012.

[GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Heidelberg, March 2015.

[GPSV21] Steven Galbraith, Lorenz Panny, Benjamin Smith, and Frederik Vercauteren. Quantum equivalence of the dlp and cdhp for group actions. *Mathematical Cryptology*, 1(1):40–44, Jun. 2021.

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73. ACM Press, June 2021.

[Kan18] Daniel M. Kane. Quantum money from modular forms, 2018. https://arxiv.org/abs/1809.05925.

[KKVB02] Elham Kashefi, Adrian Kent, Vlatko Vedral, and Konrad Banaszek. Comparison of quantum oracles. *Phys. Rev. A*, 65:050304, May 2002.

[KLS22] Andrey Boris Khesin, Jonathan Z Lu, and Peter W Shor. Publicly verifiable quantum money from random lattices, 2022. https://arxiv.org/abs/2207.13135v2.

[KSS21] Daniel M. Kane, Shahed Sharif, and Alice Silverberg. Quantum money from quaternion algebras. Cryptology ePrint Archive, Report 2021/1294, 2021. https://eprint.iacr.org/2021/1294.

[LAF$^+$10] Andrew Lutomirski, Scott Aaronson, Edward Farhi, David Gosset, Jonathan A. Kelner, Avinatan Hassidim, and Peter W. Shor. Breaking and making quantum money: Toward a new quantum cryptographic protocol. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 20–31. Tsinghua University Press, January 2010.

[LMZ23] Jiahui Liu, Hart Montgomery, and Mark Zhandry. Another round of breaking and making quantum money: How to not build it from lattices, and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 611–638. Springer, Heidelberg, April 2023.

[Lut10] Andrew Lutomirski. An online attack against wiesner's quantum money, 2010. https://arxiv.org/abs/1010.0256.

[LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 326–355. Springer, Heidelberg, August 2019.

[MM22] Luciano Maino and Chloe Martindale. An attack on SIDH with arbitrary starting curve. Cryptology ePrint Archive, Report 2022/1026, 2022. https://eprint.iacr.org/2022/1026.

[MZ22] Hart Montgomery and Mark Zhandry. Full quantum equivalence of group action DLog and CDH, and more. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 3–32. Springer, Heidelberg, December 2022.

[NOTT20] Kohei Nakagawa, Hiroshi Onuki, Atsushi Takayasu, and Tsuyoshi Takagi. L_1-norm ball for CSIDH: Optimal strategy for choosing the secret key space. Cryptology ePrint Archive, Report 2020/181, 2020. https://eprint.iacr.org/2020/181.

[Pan23] Lorenz Panny. Csi-fish really isn't polynomial-time, 2023. https://yx7.cc/blah/2023-04-14.html.

[Pei20] Chris Peikert. He gives C-sieves on the CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 463–492. Springer, Heidelberg, May 2020.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[Rob21] Bhaskar Roberts. Security analysis of quantum lightning. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 562–567. Springer, Heidelberg, October 2021.

[Rob23] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 472–503. Springer, Heidelberg, April 2023.

[Rog06] Phillip Rogaway. Formalizing human ignorance. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 06*, volume 4341 of *LNCS*, pages 211–228. Springer, Heidelberg, September 2006.

[RS06] Alexander Rostovtsev and Anton Stolbunov. Public-Key Cryptosystem Based On Isogenies. Cryptology ePrint Archive, Report 2006/145, 2006. https://eprint.iacr.org/2006/145.

[RZ21] Bhaskar Roberts and Mark Zhandry. Franchised quantum money. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 549–574. Springer, Heidelberg, December 2021.

[Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th FOCS*, pages 124–134. IEEE Computer Society Press, November 1994.

[Shp08] Vladimir Shpilrain. Cryptanalysis of stickel's key exchange scheme. In Edward A. Hirsch, Alexander A. Razborov, Alexei Semenov, and Anatol Slissenko, editors, *Computer Science – Theory and Applications*, pages 283–288, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[Sti05] E. Stickel. A new method for exchanging secret keys. In *Third International Conference on Information Technology and Applications (ICITA'05)*, volume 2, pages 426–430, 2005.

[Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983.

[Win99] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theor.*, 45(7):2481–2485, November 1999.

[WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Heidelberg, October 2021.

[YZ22] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *63rd FOCS*, pages 69–74. IEEE Computer Society Press, October / November 2022.

[Zha19] Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 408–438. Springer, Heidelberg, May 2019.