Another Round of Breaking and Making Quantum Money:
How to Not Build It from Lattices, and More

Jiahui Liu[1], Hart Montgomery[2], and Mark Zhandry[3]

[1] Univerty of Texas at Austin, jiahui@cs.utexas.edu[0000−0003−4380−8168]
[2] Linux Foundation, hmontgomery@linuxfoundation.org[0000−0002−8907−5791]
[3] NTT & Princeton University, mzhandry@gmail.com[0000−0001−7071−6272]

Abstract. This work provides both negative and positive results for publicly verifiable quantum money.

- In the first part, we give a general theorem, showing that a certain natural class of quantum money schemes from lattices cannot be secure. We use this theorem to break the recent quantum money proposal of Khesin, Lu, and Shor([KLS22]).
- In the second part, we propose a framework for building quantum money and quantum lightning we call *invariant money* which abstracts and formalizes some ideas of quantum money from knots [FGH+12] and its precedent work [LAF+10]. In addition to formalizing this framework, we provide concrete hard computational problems loosely inspired by classical knowledge-of-exponent assumptions, whose hardness would imply the security of *quantum lightning*, a strengthening of quantum money where not even the bank can duplicate banknotes.
- We discuss potential instantiations of our framework, including an oracle construction using cryptographic group actions and instantiations from rerandomizable functional encryption, isogenies over elliptic curves, and knots.

1 Introduction

1.1 Motivation

Quantum information promises to revolutionize cryptography. In particular, the no cloning theorem of quantum mechanics opens the door to *quantum cryptography*: cryptographic applications that are simply impossible classically. The progenitor of this field, due to Wiesner [Wie83], is quantum money: quantum digital currency that cannot be counterfeited due to the laws of physics. Since Wiesner's proposal, many applications of quantum information to cryptography have been proposed, including quantum key distribution (QKD) [BB87], randomness expansion [Col09, CY14, BCM+18], quantum copy protection [Aar09, AL21, ALL+21, CLLZ21], quantum one-time programs [BGS13], and much more.

Throughout the development of quantum cryptography, quantum money has remained a central object, at least implicitly. Indeed, the techniques used

for quantum money are closely related to those used in other applications. For example, the first message in the BB84 quantum QKD protocol [BB87] is exactly a banknote in Wiesner's scheme. The techniques used by [BCM+18] to prove quantumness using classical communication have been used to construct quantum money with classical communication [RS19]. The subspace states used by [AC12] to construct quantum money were recently used to build quantum copy protection [ALL+21].

The Public Verification Barrier. Wiesner's scheme is only privately verifiable, meaning that the mint is needed to verify. This results in numerous weaknesses. Improper verification opens the scheme to active attacks [Lut10]. Moreover, private verification is not scalable, as the mint would be required to participate in every single transaction. Wiesner's scheme also requires essentially perfect quantum storage, since otherwise banknotes in Wiesner's scheme will quickly decohere and be lost.

All these problems are readily solved with *publicly verifiable* quantum money[4], where anyone can verify, despite the mint being the sole entity that can mint notes. Public verification immediately eliminates active attacks, and solves the scaling problem since the transacting users can verify the money for themselves. Aaronson and Christiano [AC12] also explain that public verifiability allows for also correcting any decoherance, so users can keep their banknotes alive indefinitely.

Unfortunately, constructing convincing publicly verifiable quantum money has become a notoriously hard open question. Firstly, some natural modifications to Wiesner's quantum money scheme will not give security under public verification [FGH+10]. Aaronson [Aar09], and later Aaronson and Christiano [AC12] gave publicly verifiable quantum money relative to quantum and classical oracles, respectively. Such oracle constructions have the advantage of provable security, but it is often unclear how to instantiate them in the real world[5]: in both [Aar09] and [AC12], "candidate" instantiations were proposed, but were later broken [LAF+10, CPDDF+19]. Another candidate by Zhandry [Zha19] was broken by Roberts [Rob21]. Other candidates have been proposed [FGH+12, Kan18, KSS21], but they all rely on new, untested assumptions that have received little cryptanalysis effort. The one exception, suggested by [BDS16] and proved by [Zha19], uses indistinguishability obfuscation (iO) to instantiate Aaronson and Christiano's scheme [AC12]. Unfortunately, the post-quantum security of iO remains poorly understood, with all known constructions of post-quantum iO [GGH15, BGMZ18, BDGM20, WW21] being best labeled as candidates, lacking justification under widely studied assumptions.

Thus, it remains a major open question to construct publicly verifiable quantum money from standard cryptographic tools. Two such post-quantum

[4] Sometimes it is also referred to as public-key quantum money. We may use the two terms interchangeably.

[5] Quantum oracles are quantum circuits accessible only as a black-box unitary. They are generally considered as strong relativizing tools when used in proofs. Classical oracles are black-box classical circuits, a much weaker tool.

tools we will investigate in this work are the two most influential and well-studied: lattices and isogenies over elliptic curves.

This public verification barrier is inherited by many proposed applications of quantum cryptography. For example, quantum copy protection for any function whose outputs can be verified immediately implies a publicly verifiable quantum money scheme. As such, all such constructions in the standard model [ALL+21, CLLZ21] require at a minimum a computational assumption that implies quantum money.[6]

Quantum Money Decentralized: Quantum Lightning An even more ambitious goal is a publicly verifiable quantum money where the bank/mint itself should *not* be capable of duplicating money states. To guarantee unclonability, the scheme should have a "collision-resistant" flavor: no one can (efficiently) generate two valid money states with the same serial number. This notion of quantum money appeared as early in [LAF+10]; the name "quantum lightning" was given in [Zha19].

Quantum lightning has broader and more exciting applications: as discussed in [Zha19, Col19, CS20, AGKZ20], it can be leveraged as verifiable min-entropy, useful building blocks to enhance blockchain/smart contract protocols and moreover, it could lead to decentralized cryptocurrency without a blockchain.

Quantum money has a provably secure construction from iO, a strong cryptographic hammer but still a widely used assumption. On the other hand, quantum lightning from even *relatively standard-looking* assumptions remains open. Some existing constructions [Kan18, KSS21] use strong oracles such as quantum oracles, with conjectured instantiations that did not go through too much cryptanalysis. [FGH+12] is another candidate built from conjectures in knot theory. But a correctness proof and security reduction are not provided in their paper.

Collapsing vs. Non-Collapsing With a close relationship to quantum money, collapsing functions [Unr16] are a central concept in quantum cryptography. A collapsing function f says that one should not be able to distinguish a superposition of pre-images $\frac{|x_1\rangle+|x_2\rangle\cdots|x_k\rangle}{\sqrt{k}}$, from a measured pre-image $|x_i\rangle, i \in [k]$ for some image $y = f(x_i)$, for all $i \in [k]$.

While collapsing functions give rise to secure post-quantum cryptography like commitment schemes, its precise opposite is necessary for quantum money: if no verification can distinguish a money state in a superposition of many supports from its measured state, a simple forgery comes ahead. Hence, investigating the collapsing/non-collapsing properties of hash functions from lattices and isogenies will provide a win-win insight into quantum money and post-quantum security of existing cryptographic primitives.

[6] This holds true even for certain weaker versions such as copy *detection*, also known as infinite term secure software leasing.

2 Our Results

In this work, we give both negative and positive results for publicly verifiable quantum money.

Breaking Quantum Money. Very recent work by Khesin, Lu, and Shor [KLS22] claims to construct publicly verifiable quantum money from the hardness of worst-case lattice problems, a standard assumption. Our first contribution is to identify a fatal flaw in their security proof, and moreover show how to exploit this flaw to forge unlimited money. After communicating this flaw and attack, the authors of [KLS22] have retracted their paper.[7]

More importantly, we show that a general class of *natural* money schemes based on lattices *cannot* be both secure and publicly verifiable. We consider protocols where the public key is a short wide matrix $\mathbf{A}^T$, and a banknote with serial number $\mathbf{u}$ is a superposition of "short" vectors $\mathbf{y}$ such that $\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u} \bmod q$. Our attack works whenever $\mathbf{A}^T$ is uniformly random. We also generalize this to handle the case where $\mathbf{A}^T$ is uniform conditioned on having a few public short vectors in its kernel. This generalization includes the Khesin-Lu-Shor scheme as a special case. Our result provides a significant barrier to constructing quantum money from lattices.

Along the way, we prove that the SIS hash function is *collapsing* [Unr16] for all moduli, resolving an important open question in the security of post-quantum hash functions.[8]

Invariant Money/Lightning. To complement our negative result, we propose a new framework for building quantum money, based on invariants. Our framework abstracts some of the ideas behind the candidate quantum money from knots in [FGH+12] and behind [LAF+10]. Our main contributions here are two-fold:

- We propose a (classical) oracle construction that implements our framework assuming the existence of a quantum-secure cryptographic group action and a relatively modest assumption about *generic* cryptographic group actions. We then give proposals for instantiating our invariant framework on more concrete assumptions. The first is based on isogenies over elliptic curves[9]; the second is based on rerandomizable functional encryption with certain properties; finally, we also discuss the quantum money from knots construction in [FGH+12] with some modifications.
- In order to gain confidence in our proposals, we for the first time formalize abstract properties of the invariant money under which security can be proved. Concretely, we prove that a certain mixing condition is sufficient to characterize the states accepted by the verifier, and in particular prove

[7] We thank the authors of [KLS22] for patiently answering our numerous questions about their work, which was instrumental in helping us identify the flaw.

[8] Previously, [LZ19] showed that SIS was collapsing for a super-polynomial modulus.

[9] The recent attacks [CD22, MM22, Rob22] on SIDH do not apply to the isogeny building blocks we need. We will elaborate in the full version

correctness[10]. We also propose "knowledge of path" security properties for abstract invariant structures which would be sufficient to justify security. These knowledge of path assumptions are analogs of the "knowledge of exponent" assumption on groups proposed by Damgård [Dam92]. Under these assumptions, we are even able to show that the invariants give quantum *lightning* [LAF+10, Zha19], the aforementioned strengthening of quantum money that is known to have additional applications.
Note that the knowledge of exponent assumption in groups is quantumly broken on groups due to the discrete logarithm being easy. However, for many of our assumptions, which are at least conjectured to be quantum-secure, the analogous knowledge of path assumption appears plausible, though certainly more cryptanalysis is needed to gain confidence. The main advantage of our proposed knowledge of path assumption is that it provides a concrete cryptographic property that cryptographers can study and analyze with a well-studied classical analog.

3 Technical Overview

3.1 How to Not Build Quantum Money from Lattices

We first describe a natural attempt to construct quantum money from lattices, which was folklore but first outlined by Zhandry [Zha19]. The public key will contain a random tall matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}, m \gg n$. To mint a banknote, first generate a superposition $|\psi\rangle = \sum_{\mathbf{y}} \alpha_{\mathbf{y}} |\mathbf{y}\rangle$ of short vectors $\mathbf{y} \in \mathbb{Z}^m$, such that $|\mathbf{y}| \ll q$. A natural $|\psi\rangle$ is the discrete-Gaussian-weighted state, where $\alpha_{\mathbf{y}} \propto \sqrt{e^{-\pi |\mathbf{y}|^2/\sigma^2}}$ for a width parameter σ. Then compute in superposition and measure the output of the map $\mathbf{y} \mapsto \mathbf{A}^T \cdot \mathbf{y} \bmod q$, obtaining $\mathbf{u} \in \mathbb{Z}_q^n$. The state collapses to:

$$|\psi_{\mathbf{u}}\rangle \propto \sum_{\mathbf{y}:\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u}} \alpha_{\mathbf{y}} |\mathbf{y}\rangle \ .$$

This will be the money state, and $\mathbf{u}$ will be the serial number. This state can presumably not be copied: if one could construct two copies of $|\psi_{\mathbf{u}}\rangle$, then one could measure both, obtaining two short vectors $\mathbf{y}, \mathbf{y}'$ with the same coset $\mathbf{u}$. As $|\psi_{\mathbf{u}}\rangle$ is a superposition of many vectors (since $m \gg n$), with high probability $\mathbf{y} \neq \mathbf{y}'$. Subtracting gives a short vector $\mathbf{y} - \mathbf{y}'$ such that $\mathbf{A}^T \cdot (\mathbf{y} - \mathbf{y}') = 0$, solving the Short Integer Solution (SIS) problem. SIS is presumably hard, and this hardness can be justified based on the hardness of worst-case lattice problems such as the approximate Shortest Vector Problem (SVP).

The challenge is: how to verify $|\psi_{\mathbf{u}}\rangle$? Certainly, one can verify that the support of a state is only short vectors $\mathbf{y}$ such that $\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u}$. But this alone is not

[10] [FGH+12] did not analyze correctness of their knot-based proposal, nor analyze the states accepted by their verifier and formalize the property needed for a security proof. [LAF+10] had informal correctness analysis on their proposal, but also did not analyze the security property needed.

enough: one can fool such a verification by any *classical* $\mathbf{y}$ in the support of $|\psi_{\mathbf{u}}\rangle$. To forge then, an adversary simply measures $|\psi_{\mathbf{u}}\rangle$ to obtain $\mathbf{y}$, and then copies $\mathbf{y}$ as many times as it likes.

To get the scheme to work, then, one needs a verifier that can distinguish classical $\mathbf{y}$ from superpositions. This is a typical challenge in designing publicly verifiable money schemes. A typical approach is to perform the quantum Fourier transform (QFT): the QFT of $\mathbf{y}$ will result in a uniform string, whereas the QFT of $|\psi_{\mathbf{u}}\rangle$ will presumably have structure. Indeed, if $|\psi_{\mathbf{u}}\rangle$ is the Gaussian superposition, following ideas of Regev [Reg05], the QFT of $|\psi_{\mathbf{u}}\rangle$ will be statistically close to a superposition of samples $\mathbf{A}\cdot\mathbf{r}+\mathbf{e}$, where $\mathbf{r}$ is uniform in $\mathbb{Z}_q^n$, and $\mathbf{e}\in\mathbb{Z}_q^m$ is another discrete Gaussian of width q/σ. The goal then is to distinguish such samples from uniform.

Unfortunately, such distinguishing is likely hard, as this task is the famous (decisional) Learning with Errors (LWE) problem. LWE is presumably hard, which can be justified based on the hardness of the same worst-case lattice problems as with SIS, namely SVP. So either LWE is hard, or the quantum money scheme is insecure in the first place.

Nevertheless, this leaves open a number of possible strategies for designing quantum money from lattices, including:

1. What if non-Gaussian $|\psi\rangle$ is chosen?
2. What if distinguishing is not done via the QFT but some other quantum process?
3. What if we somehow make LWE easy?

The first significant barrier beyond the hardness of LWE is due to Liu and Zhandry [LZ19]. They show that, if the modulus q is super-polynomial, then the map $\mathbf{y}\mapsto\mathbf{A}^T\cdot\mathbf{y}$ for a random $\mathbf{A}$ is *collapsing* [Unr16]: that is, for *any* starting state $|\psi_{\mathbf{u}}\rangle$ of short vectors, distinguishing $|\psi_{\mathbf{u}}\rangle$ from $\mathbf{y}$ is infeasible for *any* efficient verification process. Collapsing is the preferred notion of post-quantum security for hash functions, as it is known that collision resistance is often not sufficient for applications when quantum adversaries are considered.

The result of [LZ19] follows from the hardness of LWE (which is quantumly equivalent to SIS [Reg05]), albeit with a noise rate super-polynomially smaller than q/σ which is a stronger assumption than the hardness with rate q/σ. Moreover, their result requires q to be super-polynomially larger than σ. In practice, one usually wants q to be polynomial, and the result of [LZ19] leaves open the possibility of building quantum money in such a setting.

What about making LWE easy (while SIS remains hard)? The usual approach in the lattice literature to making decisional LWE easy is to output a short vector $\mathbf{s}$ in the kernel of $\mathbf{A}^T$. If $|\mathbf{s}|\ll(q/\sigma)$, this allows for distinguishing LWE samples from uniform, since $\mathbf{s}\cdot(\mathbf{A}\cdot\mathbf{r}+\mathbf{e})=\mathbf{s}\cdot\mathbf{e}$, which will be small relative to q, while $\mathbf{s}\cdot\mathbf{x}$ for uniform $\mathbf{x}$ will be uniform in $\mathbb{Z}_q$. Unfortunately, adding such short vectors breaks the security proof, since $\mathbf{s}$ is a SIS solution, solving SIS is trivially easy by outputting $\mathbf{s}$. To revive the security, one can try reducing to the 1-SIS problem, which is to find a short SIS solution that is linearly independent of $\mathbf{s}$. 1-SIS can be proved hard based on the same worst-case lattice problems as SIS [BF11].

However, in the scheme above, it is not clear if measuring two forgeries and taking the difference should result in a vector linearly independent of $\mathbf{s}$.

The Recent Work of [KLS22]. Very recently, Khesin, Lu, and Shor [KLS22] attempt to provide a quantum money scheme based on lattices. Their scheme has some similarities to the blueprint discussed above, taking advantage of each of the strategies 1, 2 and 3. But there are other differences as well: the state $|\psi\rangle$ is created as a superposition over a lattice rather than the integers, and the measurement of $\mathbf{u}$ is replaced with a move complex general positive operator-value measurement (POVM). [KLS22] claims to prove security under the hardness of finding a second short vector in a random lattice when already given a short vector. This problem is closely related to 1-SIS, and follows also from the hardness of worst-case lattice problems.

Our Results. First, we show an alternative view of [KLS22] which shows that it does, indeed, fall in the above framework. That is, there is a way to view their scheme as starting from $|\psi\rangle$ that is a non-Gaussian superposition of short integer vectors $\mathbf{y}$. The minting process in our alternate view then measures $\mathbf{A}^T \cdot \mathbf{y}$, where $\mathbf{A}$ is part of the public key, and is chosen to be uniform except that it is orthogonal to 3 short vectors $\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2$. These vectors play a role in verification, as they make the QFT non-uniform. Using this alternative view, we also demonstrate a flaw in the security proof of [KLS22], showing that forged money states actually do not yield new short vectors in the lattice. See Section D of the full version for details.

We then go on (Section 5) to show an explicit attack against their money scheme. More generally, we show an attack on a wide class of instantiations of the above framework. Our attack works in two steps:

- First, we extend the collapsing result of [LZ19] to also handle the case of polynomial modulus, and in particular, we only need LWE to be hard for noise rate that is slightly smaller than q/σ. This resolves an important open by showing that SIS is collapsing for all moduli.
  Our proof requires a novel reduction that exploits a more delicate analysis of the quantum states produced in the proof of [LZ19]. We also extend the result in a meaningful way to the case where several short kernel vectors $\mathbf{s}_0, \mathbf{s}_1, \ldots$ are provided. We show that instead of just using $\mathbf{y}$ as a forgery (which can be distinguished using the short vectors $\mathbf{s}_i$), a particular superposition over vectors of the form $\mathbf{y} + \sum_i c_i \mathbf{s}_i$ can fool any efficient verification. Fooling verification requires the hardness a certain "k-LWE" problem, which we show follows from worst-case lattice problems in many settings (see Section E). This requires us to extend the known results on k-LWE hardness, which may be of independent interest.
- Then we show how to construct such a superposition efficiently given only $\mathbf{y}$ and the $\mathbf{s}_i$, in many natural settings. Our settings include as a special case the setting of [KLS22]. Along the way, we explain how to construct Gaussian superpositions over lattices, when given a short basis. The algorithm is a coherent version of the classical discrete Gaussian sampling algorithm [GPV08].

> In general, it is not possible to take a classical distribution and run it on a superposition of random coins to get a superposition with weights determined by the distribution. This is because the random coins themselves will be left behind and entangled with the resulting state. We show how to implement the classical algorithm coherently in a way that does not leave the random coins behind or any other entangled bits. Such an algorithm was previously folklore (e.g. it was claimed to exist without justification by [KLS22]), but we take care to actually write out the algorithm.

After communicating this flaw and attack to the authors of [KLS22], they have retracted their paper.[11]

3.2 Quantum Money from Walkable Invariants

In the second part of the paper, we describe a general framework for instantiating publicly verifiable quantum money from invariants satisfying certain conditions. This framework abstracts the ideas behind the construction of quantum money from knots [FGH+12] and its precedent [LAF+10].

At a high level, we start from a set X, which is partitioned into many disjoint sets $O \subseteq X$. There is a collection of efficiently computable (and efficiently invertible) permutations on X, such that for every permutation in the collection and every O in the partition, the permutation maps elements of O to O. Such a set of permutations allows one to take an element $x \in O$, and perform a walk through O. We additionally assume an invariant $I : X \to Y$ on X, such that I is constant on each element O of the partition. In other words, I is invariant under action by the collection of permutations.

In the case of [FGH+12], X is essentially the set of knot diagrams[12], the permutations are Reidemeister moves, and the invariant is the Alexander polynomial.

An honest quantum money state will essentially be a uniform superposition over O [13]. Such a state is constructed by first constructing the uniform superposition over X, and then measuring the invariant I. Applying a permutation from the collection will not affect such a state. Thus, verification attempts to test whether the state is preserved under action by permutations in the collection by performing an analog of a swap test, and only accepts if the test passes.

In [FGH+12], it is explained why certain attack strategies are likely to be incapable of duplicating banknotes. However, no security proof is given under widely believed hard computational assumptions. To make matters worse, [FGH+12] do not analyze what types of states are accepted by the verifier. It could be, for example, that duplicating a banknote perfectly is computationally infeasible, but

[11] We once again want to emphasize that the authors of [KLS22] were exceptionally helpful and we thank them for their time spent helping us understand their work.

[12] Due to certain concerns about security, [FGH+12] actually sets X to contain extra information beyond a knot diagram.

[13] Technically, it is a uniform superposition over the pre-images of some y in the image of I. If multiple O have the same y, then the superposition will be over all such O.

there are fake banknotes that pass verification that can be duplicated; this is exactly what happens in the lattice-based schemes analyzed above in Section 3.1. Given the complexities of their scheme, there have been limited efforts to understand the security of the scheme. This is problematic, since there have been many candidates for public key quantum money that were later found to be insecure.

Generally, a fundamental issue with public key quantum money schemes is that, while quantum money schemes rely on the no-cloning principle, the no-cloning theorem is information-theoretic, whereas publicly verifiable quantum money is always information-theoretically clonable. So unclonability crucially relies on the adversary being computationally efficient. Such computational unclonability is far less understood than traditional computational tasks. Indeed, while there have been a number of candidate post-quantum hard computational tasks, there are very few quantum money schemes still standing. The challenge is in understanding if and how quantum information combines with computational bounds to give computational unclonability.

To overcome this challenge, the security analysis should be broken into two parts: one part that relies on *information-theoretic* no-cloning, and another part that relies on a computational hardness assumption. Of course, the security of the scheme itself could be such an assumption, so we want to make the assumption have nothing to do with cloning. One way to accomplish this is to have the assumption have classical inputs and outputs (which we will call "classically meaningful"), so that it could in principle be falsified by a classical algorithm, which are obviously not subject to quantum unclonability. Separating out the quantum information from the computational aspects would hopefully give a clearer understanding of why the scheme should be unclonable, hopefully allow for higher confidence in security. Moreover, as essentially all widely studied assumptions are classically meaningful, any attempt to prove security under a widely studied assumptions would have to follow this blueprint, and indeed the proof of quantum money from obfuscation [Zha19] is of this form.

Our Results. In this work, we make progress towards justifying invariant-based quantum money.

- First, we prove that if a random walk induced by the collection of permutations mixes, then we can completely characterize the states accepted by verification. The states are exactly the uniform superpositions over O [14]. Unfortunately, it is unclear if the knot construction actually mixes, and any formal proof of mixing seems likely to advance knot theory[15].
- Second, we provide concrete security properties under which we can prove security. These properties, while still not well-studied, at least have no obvious connection to cloning, and are meaningful even classically. Under

[14] Or more generally, if multiple O have the same y, then accepting states are exactly those that place equal weight on elements of each O, but the weights may be different across different O

[15] Nevertheless we provide a discussion on the knot money instantiation in the knot instantiation section of the full version.

these assumptions, we can even prove that the schemes are in fact quantum lightning, the aforementioned strengthening of quantum money where not even the mint can create two banknotes of the same serial number.

Our Hardness Assumptions We rely on two hardness assumptions in our invariant money scheme for a provably secure: the *path-finding* assumption and *knowledge of path finding* assumption.

Informally speaking, the path-finding assumption states that, given some adversarially sampled x from a set of elements X and given a set of "permutations" Σ, it is hard for any efficient adversary, given a random $z \in X$, where there exists some $\sigma \in \Sigma$ such that $\sigma(x) = z$, to find such a σ. One can observe that it is similar to a "discrete logarithm" style of problem. Even though we cannot use discrete logarithm due to its quantum insecurity, we have similar hard problems in certain isogenies over ellitic curves, abstracted as "group action discrete logarithm" problems [ADMP20].

Our Knowledge of Path Assumptions. The main novel assumption we use is a "knowledge of path" assumption. This roughly says that if an algorithm outputs two elements x, z in the same O, then it must "know" a path between them: a list of permutations from the collection that, when composed, would take x to z. While such a knowledge of path assumption is undoubtedly a strong assumption, it seems plausible in a number of relevant contexts (e.g. elliptic curve isogenies that have no known non-trivial attacks or "generic" group actions).

Formalizing the knowledge of path assumption is non-trivial. The obvious *classical* way to define knowledge of path is to say that for any adversary, there is an extractor that can compute the path between x and z. Importantly, the extractor must be given the same random coins as the adversary, so that it can compute x and z for itself and moreover know what random choices the adversary made that lead to x, z. Essentially, by also giving the random coins, we would be effectively making the adversary deterministic, which is crucial for the extractor's output to be related to the adversary's output.

Unfortunately, quantumly the above argument does not make much sense, as quantum algorithms can have randomness without having explicit random coins. In fact, there are quantum procedures that are *inherently probabilistic*, in the sense that the process is efficient, but there is no way to run the process twice and get the same outcome both times. This is actually crucial to our setting: we are targeting the stronger quantum lightning, which means that even the mint cannot create two banknotes with the same serial number. This means that the minting process is inherently probabilistic. The adversary could, for example, run the minting process, but with its own minting key. Such an adversary would then be inherently probabilistic and we absolutely would need a definition that can handle such adversaries.

Our solution is to exploit the fact that quantum algorithms can always be implemented *reversibly*. We then observe that with a classical reversible adversary, an equivalent way to define knowledge assumptions would be to just feed the entire *final* state of the adversary (including output) into the extractor. By

reversibility, this is equivalent to giving the input, coins included, to the extractor. But this alternate extraction notion actually *does* make sense quantumly. Thus our knowledge of path assumption is defined as giving the extractor the entire final (quantum) state of the reversible adversary, and asking that the extractor can find a path between x and z. This assumption allows us to bypass the issue of inherently probabilistic algorithms, and is sufficient for us to prove security.

Instantiations of Invariant Quantum Money and Lightning After we provide the characterization of security needed for invariant money, we discuss four candidate instantiations[16]:

- We show a construction from structured oracles and generic cryptographic group actions. Notably, while we do not know how to instantiate these oracles, we can prove that this construction is secure assuming the existence of a cryptographic group action and the assumption that the knowledge of path assumption holds over a generic cryptographic group action.[17]
- We explain how re-randomizable functional encryption, a type of functional encryption with special properties that seem reasonable, can be used to build another candidate quantum lightning. We don't currently have a provably secure construction from standard cryptographic assumptions for this special re-randomizable functional encryption, but we provide a candidate construction based on some relatively well-studied primitives.
- Elliptic curve isogenies are our final new candidate instantiation. We outline how, given some assumptions about sampling certain superpositions of elliptic curves, it may be possible to build quantum lightning from isogeny-based assumptions.
- Finally, we analyze the construction of quantum money from knots in [FGH+12] in our framework.

For all these three constructions, we show that their corresponding *path-finding* problem between two elements x, z in the same O is relatively straightforward to study (reducible to reasonably well-founded assumptions). Nevertheless, we need the knowledge of path assumptions to show that we can extract these paths from a (unitary) adversary. We believe that one may show a knowledge-of-path property when replacing some plain model components in the above candidates with (quantum accessible) *classical* oracles, thus giving the possibility for a first quantum lightning scheme relative to only classical oracles and widely studied assumptions.

4 Preliminaries

In this section we explain some background material needed for our work.

[16] Throughout the sections on invariant quantum money framework and construction in the full verison, we will sometimes interchangeably use "money" or "lightning". But in fact the proposed candidates are all candidates for quantum lightning.

[17] This seems like a very plausible assumption to us: classically, the knowledge of exponent would almost trivially hold over generic groups.

For quantum notations, we denote $|\cdot\rangle$ as the notation for a pure state and $|\cdot\rangle\langle\cdot|$ for its density matrix. ρ denotes a general mixed state.

We will go over some fundamental lattice facts and then move to quantum money definitions. Due to the restriction of space, we leave some additional lattice basics, hardness theorems and necessary quantum background (in particular related to lattices) to Appendix preliminaries section of the full version.

4.1 Lattice Basics

We say a distribution $\mathcal{D}$ is (B, δ)-bounded if the probability that $\mathcal{D}$ outputs a value larger than B is less than δ. We extend this to distributions that output vectors in an entry-by-entry way. Given a set of vectors $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\}$, we define the norm of $\mathbf{B}$, denoted $||\mathbf{B}||$, as the length of the longest vector in $\mathbf{B}$, so $||\mathbf{B}|| = \max_i ||\mathbf{b}_i||$. For any lattice Λ, we define the minimum distance (or first successive minimum) $\lambda_1(\Lambda)$ as the length of the shortest nonzero lattice vector in Λ.

We next define discrete Gaussians formally. Since we later use their lemmas, our definition is loosely based on that of [BLP+13].

Definition 1. *For any $\sigma > 0$, the n-dimensional Gaussian function $\rho_\sigma : \mathbb{R}^n \to [0, 1]$ is defined as*

$$\rho_r(\mathbf{x}) = e^{-\pi \frac{\mathbf{x}^2}{\sigma^2}}$$

We define the discrete Gaussian function *with parameter σ at point $\mathbf{p} \in \mathbb{R}^n$, which we usually denote $\mathcal{D}_{\mathbf{\Psi}_\sigma}$ or just $\mathbf{\Psi}_\sigma$ when the context is clear, as the function over all of the integers $\mathbf{y} \in \mathbb{Z}^n$ such that the probability mass of any $\mathbf{y}$ is proportional to*

$$e^{-\pi \frac{(\mathbf{p}-\mathbf{y})^2}{\sigma^2}}.$$

We can also define more complicated discrete Gaussians over lattices. In this case, let $\mathbf{\Sigma}$ be a matrix in $\mathbb{R}^{n \times n}$. The discrete Gaussian over a lattice Λ with center $\mathbf{p}$ and "skew" parameter Σ is the function over all lattice points *in Λ such that the probability mass of any $\mathbf{y}$ is proportional to*

$$e^{-\pi(\mathbf{p}-\mathbf{y})^T \left(\mathbf{\Sigma}\mathbf{\Sigma}^\mathbf{T}\right)^{-1}(\mathbf{p}-\mathbf{y})},$$

very similar to as before. We usually denote this type of discrete Gaussian as $\mathbf{\Psi}_{\Lambda,\mathbf{\Sigma},\mathbf{p}}$ or $\mathcal{D}_{\mathbf{\Psi}_{\Lambda,\mathbf{\Sigma},\mathbf{p}}}$, where we sometimes substitute σ for $\mathbf{\Sigma}$ when $\mathbf{\Sigma} = \sigma \cdot \mathbf{I}_n$, where $\mathbf{I}_n$ is the $n \times n$ identity matrix. We also sometimes omit parameters when they are obvious (e.g. 0) in context.

We will explain how to efficiently sample discrete Gaussians quantumly in B.3 of the full version.

4.2 General LWE Definition

In this section we define basic LWE with an eye towards eventually defining k-LWE. We note that, while equivalent to the standard definitions, our definitions here are presented a little bit differently than usual in lattice cryptography. This is so that we can keep the notation more consistent with the typical quantum money and quantum algorithms presentation styles. We first provide a properly parameterized definition of the LWE problem [Reg05].

Definition 2. ***Learning with Errors (LWE) Problem:*** *Let n, m, and q be integers, let $\mathcal{D}_{\mathbf{A}}$ and $\mathcal{D}_{\mathbf{r}}$ be distributions over $\mathbb{Z}_q^n$, and let $\mathcal{D}_{\Psi}$ be a distribution over $\mathbb{Z}_q^m$. Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ be a matrix where each row is sampled from $\mathcal{D}_{\mathbf{A}}$, let $\mathbf{r} \in \mathbb{Z}_q^n$ be a vector sampled from $\mathcal{D}_{\mathbf{r}}$, and let $\mathbf{e} \in \mathbb{Z}_q^m$ be a vector sampled from $\mathcal{D}_{\Psi}$. Finally, let $\mathbf{t} \in \mathbb{Z}_q^m$ be a uniformly random vector.*

The $(n, m, q, \mathcal{D}_{\mathbf{A}}, \mathcal{D}_{\mathbf{r}}, \mathcal{D}_{\Psi})$-LWE problem is defined to be distinguishing between the following distributions:

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{r} + \mathbf{e}) \text{ and } (\mathbf{A}, \mathbf{t}).$$

4.3 Quantum Money and Quantum Lightning

Here, we define public key quantum money and quantum lightning. Following Aaronson and Christiano [AC12], we will only consider so-called "mini-schemes", where there is only a single banknote.

Both quantum money and quantum lightning share the same syntax and correctness requirements. There are two quantum polynomial-time algorithms $\mathsf{Gen}, \mathsf{Ver}$ such that:

- $\mathsf{Gen}(1^\lambda)$ samples a classical serial number σ and a quantum state $|\psi\rangle$.
- $\mathsf{Ver}(\sigma, |\psi\rangle)$ outputs a bit 0 or 1.

Correctness. We require that there exists a negligible function negl such that $\Pr[\mathsf{Ver}(\mathsf{Gen}(1^\lambda))] \geq 1 - \mathsf{negl}(\lambda)$.

Security. Where public key quantum money and quantum lightning differ is in security. The differences are analogous to the differences between one-way functions and collision resistance.

Definition 3 (Quantum Money Unforgeability). $(\mathsf{Gen}, \mathsf{Ver})$ *is secure public key quantum money if, for all quantum polynomial-time A, there exists a negligible* negl *such that A wins the following game with probability at most* negl*:*

- *The challenger runs $(\sigma, |\psi\rangle) \leftarrow \mathsf{Gen}(1^\lambda)$, and gives $\sigma, |\psi\rangle$ to A.*
- *A produces a potentially entangled joint state $\rho_{1,2}$ over two quantum registers. Let ρ_1, ρ_2 be the states of the two registers. A sends $\rho_{1,2}$ to the challenger.*
- *The challenger runs $b_1 \leftarrow \mathsf{Ver}(\sigma, \rho_1)$ and $b_2 \leftarrow \mathsf{Ver}(\sigma, \rho_2)$. A wins if $b_1 = b_2 = 1$.*

Definition 4 (Quantum Lightning Unforgeability). $(\mathsf{Gen}, \mathsf{Ver})$ *is secure quantum lightning if, for all quantum polynomial-time* A*, there exists a negligible* negl *such that* A *wins the following game with probability at most* negl*:*

- A*, on input* 1^λ*, produces and sends to the challenger* σ *and* $\rho_{1,2}$*, where* $\rho_{1,2}$ *is a potentially entangled joint state over two quantum registers.*
- *The challenger runs* $b_1 \leftarrow \mathsf{Ver}(\sigma, \rho_1)$ *and* $b_2 \leftarrow \mathsf{Ver}(\sigma, \rho_2)$*.* A *wins if* $b_1 = b_2 = 1$*.*

The difference between quantum lightning and quantum money is therefore that in quantum lightning, unclonability holds, even for adversarially constructed states.

Note that, as with classical collision resistance, quantum lightning does not exist against non-uniform adversaries. Like in the case of collision resistance, we can update the syntax and security definition to utilize a common reference string (crs), which which case non-uniform security can hold. For this paper, to keep the discussion simple, we will largely ignore the issue of non-uniform security.

5 Our General Attack on a Class of Quantum Money

Due to limitation of space, we leave a detailed discussion of the [KLS22] money scheme and its flaw in of the full version.

Now, we show that a natural class of schemes, including the equivalent view on [KLS22] demonstrated in the full version, cannot possibly give secure quantum money schemes, regardless of how the verifier works.

5.1 The General Scheme

Here, we describe a general scheme which captures the alternate view above. Here, we use somewhat more standard notation from the lattice literature. Here we give a table describing how the symbols from section C map to this section:

This Section	Section C.3
q	P
n	1
m	$d' = d + 2$
$\mathbf{A}$	$\mathbf{v}'$ as a column vector
$\|\psi\rangle$	$\|\phi'\rangle$
$\mathbf{u}$	T
W	$k + \sqrt{m} \times \sigma \times \omega(\sqrt{\log(\lambda)}))$

Setup. Let q be a super-polynomial, which may or may not be prime. Sample from some distribution several short vectors $\mathbf{s}_1, \ldots, \mathbf{s}_\ell \in \mathbb{Z}_q^m$ for a constant ℓ, and assemble them as a matrix $\mathbf{S} \in \mathbb{Z}_q^{m \times \ell}$. Then generate a random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ such that $\mathbf{A}^T \cdot \mathbf{S} = 0 \bmod q$.

Minting. Create some superposition $|\psi\rangle$ of vectors in $\mathbf{y} \in \mathbb{Z}_q^m$ such that an all but negligible fraction of the support of $|\psi\rangle$ are on vectors with norm W. Let $\alpha_{\mathbf{y}}$ be the amplitude of $\mathbf{y}$ in $|\psi\rangle$.

Then apply the following map to $|\psi\rangle$:

$$|\mathbf{y}\rangle \to |\mathbf{y}, \mathbf{A}^T \cdot \mathbf{y} \bmod q\rangle$$

Finally, measure the second register to obtain $\mathbf{u} \in \mathbb{Z}_q^n$. This is the serial number, and the note is $|\psi_{\mathbf{u}}\rangle$, whatever remains of the first register, which is a superposition over short vectors $\mathbf{y}$ such that $\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u}$.

Verification. We do not specify verification. Indeed, in the following we will show that the money scheme is insecure, for *any* efficient verification scheme.

5.2 Attacking the General Scheme

We now show how to attack the general scheme. Let $\mathbf{C}$ be a matrix whose columns span the space orthogonal to the columns of $\mathcal{S}$. Let $|\psi'_{\mathbf{u}}\rangle$ be the state sampled from $|\psi_{\mathbf{u}}\rangle$ by measuring $\mathbf{y} \mapsto \mathbf{C}^T \cdot \mathbf{y}$, and letting $|\psi'_{\mathbf{u}}\rangle$ be whatever is left over.

Our attack will consist of two parts:

- Showing that $|\psi'_{\mathbf{u}}\rangle$ is indistinguishable from $|\psi_{\mathbf{u}}\rangle$, for *any* efficient verification procedure. We show (Section 5.3) that this follows from a certain "k-LWE" assumption, which depends on the parameters of the scheme (k, n, m, q, etc). In Section D of the full version, we justify the assumption in certain general cases, based on the assumed hardness of worst-case lattice problems. Note that these lattice problems are essentially (up to small differences in parameters) the same assumptions we would expect are needed to show security for the money scheme in the first place. As such, if k-LWE does not hold for these special cases, most likely the quantum money scheme is insecure anyway. Our cases include the case of [KLS22].
- Showing that $|\psi'_{\mathbf{u}}\rangle$ can be cloned. Our attack first measures $|\psi'_{\mathbf{u}}\rangle$ to obtain a single vector $\mathbf{y}$ in it's support. To complete the attack, it remains to construct $|\psi'_{\mathbf{u}}\rangle$ from $\mathbf{y}$; by repeating such a process many times on the same $\mathbf{y}$, we successfully clone. We show (Section 5.4) that in certain general cases how to perform such a construction. Our cases include the case of [KLS22].

Taken together, our attack shows that not only is [KLS22] insecure, but that it quite unlikely that any tweak to the scheme will fix it.

5.3 Indistinguishability of $|\psi'_{\mathbf{u}}\rangle$

Here, we show that our fake quantum money state $|\psi'_{\mathbf{u}}\rangle$ passes verification, despite being a very different state that $|\psi_{\mathbf{u}}\rangle$. We claim that, from the perspective of any efficient verification algorithm, $|\psi'_{\mathbf{u}}\rangle$ and $|\psi_{\mathbf{u}}\rangle$ are indistinguishable. This would mean our attack succeeds.

Toward this end, let $\mathbf{C} \in \mathbb{Z}_q^{m \times (m-\ell)}$ be a matrix whose rows span the space orthogonal to $\mathbf{S}$: $\mathbf{C}^T \cdot \mathbf{S} = 0$. Notice that the state $|\psi'_{\mathbf{u}}\rangle$ can be equivalently constructed by applying the partial measurement of $\mathbf{C}^T \cdot \mathbf{y}$ to $|\psi_{\mathbf{u}}\rangle$

Consider the following problem, which is closely related to "k-LWE"(definition4 in the full version):

Problem 1. Let n, m, q, Σ be functions of the security parameter, and D a distribution over $\mathbf{S}$. The $(n, m, q, \Sigma, \ell, D)$-LWE problem is to efficiently distinguish the following two distributions:

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{r} + \mathbf{e}) \quad \text{and} \quad (\mathbf{A}, \mathbf{C} \cdot \mathbf{r}' + \mathbf{e}) \ ,$$

Where $\mathbf{r}$ is uniform in $\mathbb{Z}_q^n$, $\mathbf{r}'$ is uniform in $\mathbb{Z}_q^{m-\ell}$, and $\mathbf{e}$ is Gaussian of width Σ. We say the problem is *hard* if, for all polynomial time quantum algorithms, the distinguishing advantage is negligible.

In Section D of the full version, we explain that in many parameter settings, including importantly the setting of [KLS22], that the hardness of Problem 1 is true (assuming standard lattice assumptions).

With the hardness of Problem 1, we can show the following, which is a generalization of a result of [LZ19] that showed that the SIS hash function is collapsing for super-polynomial modulus:

Theorem 1. *Consider sampling $\mathbf{A}, \mathbf{S}$ as above, and consider any efficient algorithm that, given $\mathbf{A}, \mathbf{S}$, samples a $\mathbf{u}$ and a state $|\phi_{\mathbf{u}}\rangle$ with the guarantee that all the support of $|\phi_{\mathbf{u}}\rangle$ is on vectors $\mathbf{y}$ such that (1) $\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u} \bmod q$ and (2) $|\mathbf{y}|_2 \leq W$.*

Now suppose $|\phi_{\mathbf{u}}\rangle$ is sampled according to this process, and then either (A) $|\phi_{\mathbf{u}}\rangle$ is produced, or (B) $|\phi'_{\mathbf{u}}\rangle$ is produced, where $|\phi'_{\mathbf{u}}\rangle$ is the result of applying the partial measurement of $\mathbf{C}^T \cdot \mathbf{y}$ to the state $|\phi_{\mathbf{u}}\rangle$.

Suppose there exists Σ such that $q/W\Sigma = \omega(\sqrt{\log \lambda})$ such that $(n, m, q, \Sigma, \ell, D)$-LWE is hard. Then cases (A) and (B) are computationally indistinguishable.

Note that an interesting consequence of Theorem 1 in the case $\ell = 0$ is that it shows that the SIS hash function is collapsing for any modulus, under an appropriate (plain) LWE distribution. This improves upon [LZ19], who showed the same but only for super-polynomial modulus. We now give the proof of Theorem 1:

Proof. For an integer t, let $\lfloor \cdot \rceil_t$ denote the function that maps a point $x \in \mathbb{Z}_q$ to the $z \in \{0, \lfloor q/t \rfloor, \lfloor 2q/t \rfloor, \cdot, \lfloor (t-1)q/t \rfloor\}$ that minimizes $|z - x|$. Here, $|z - x|$ is the smallest a such that $z = x \pm a \bmod q$. In other words, $\lfloor \cdot \rceil_t$ is a course rounding function that rounds an $x \in \mathbb{Z}_q$ to one of t points that are evenly spread out in $\mathbb{Z}_q$.

Let ρ be a mixed quantum state, whose support is guaranteed to be on $\mathbf{y}$ such that (1) $\mathbf{A}^T \cdot \mathbf{y} = \mathbf{u} \bmod q$ and (2) $|\mathbf{y}|_2 \leq W$. For a quantum process M acting on ρ, let $M(\rho)$ be the mixed state produced by applying M_i to ρ. We will consider a few types of procedures applied to on quantum states.

M_0: Given $\mathbf{A}$, M_0 is just the partial measurement of $\mathbf{y} \mapsto \mathbf{C}^T \cdot \mathbf{y}$.

M_1^t: Given $\mathbf{A}$, to apply this measurement, first sample an LWE sample $\mathbf{b} = \mathbf{A}\cdot\mathbf{r}+\mathbf{e}$. Then apply the measurement $\mathbf{y} \mapsto \lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t$. Discard the measurement outcome, and output the remaining state.

Lemma 1. *For any constants t, d, $M_1^t(\rho)$ is statistically close to $\frac{1}{d} M_1^{t\times d}(\rho) + \left(1 - \frac{1}{d}\right)\rho$*

Note that Lemma 1 means that M_1^t can be realized by the mixture of two measurements: $M_1^{t\times d}$ with probability $1/t^2$, and the identity with probability $\left(1 - \frac{1}{d}\right)$. We now give the proof.

Proof. Consider the action of M_1^t on $|\mathbf{y}\rangle\langle\mathbf{y}'|$, for a constant t. First, an LWE sample $\mathbf{b} = \mathbf{A}\cdot\mathbf{r}+\mathbf{e}$ is chosen. Then conditioned on this sample, if $\lfloor \mathbf{b}\cdot\mathbf{y} \rceil_t = \lfloor \mathbf{b}\cdot\mathbf{y}' \rceil_t$, the output is $|\mathbf{y}\rangle\langle\mathbf{y}'|$. Otherwise the output is 0. Averaging over all $\mathbf{b}$, we have that

$$M_1^t(|\mathbf{y}\rangle\langle\mathbf{y}'|) = \Pr_{\mathbf{b}}[\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t = \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t]$$

where the probability is over $\mathbf{b}$ sampled as $\mathbf{b} = \mathbf{A} \cdot \mathbf{r} + \mathbf{e}$. Recalling that $\mathbf{u} = \mathbf{A}^T \cdot \mathbf{y} = \mathbf{A}^T \cdot \mathbf{y}'$, we have that:

$$\mathbf{b} \cdot \mathbf{y} = \mathbf{r} \cdot \mathbf{u} + \mathbf{e} \cdot \mathbf{y}$$
$$\mathbf{b} \cdot \mathbf{y}' = \mathbf{r} \cdot \mathbf{u} + \mathbf{e} \cdot \mathbf{y}'$$

Now, by our choice of Σ, $|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')| < q/t$ for any constant t, except with negligible probability. We will therefore assume this is the case, incurring only a negligible error.

Note that $\mathbf{z} := \mathbf{r} \cdot \mathbf{u}$ is uniform in $\mathbb{Z}_q$ and independent of $\mathbf{e} \cdot \mathbf{y}, \mathbf{e} \cdot \mathbf{y}'$. So measuring $\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t$ is identical to measuring the result of rounding $\mathbf{e} \cdot \mathbf{y}$, except that the rounding boundaries are rotated by a random $\mathbf{z} \in \mathbb{Z}_q$. Since the rounding boundaries are q/t apart, at most a single rounding boundary can be between $\mathbf{e} \cdot \mathbf{y}$ and $\mathbf{e} \cdot \mathbf{y}'$, where "between" means lying in the shorter of the two intervals (of length $|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')|$) resulting by cutting the circle $\mathbb{Z}_q$ at the points $\mathbf{e} \cdot \mathbf{y}$ and $\mathbf{e} \cdot \mathbf{y}'$. $\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t = \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t$ if and only if no rounding boundary is between them.

Since the cyclic shift $\mathbf{z}$ is uniform each rounding boundary is uniform. Since there are t rounding boundaries and no two of them can between $\mathbf{e} \cdot \mathbf{y}$ and $\mathbf{e} \cdot \mathbf{y}'$, we have that, conditioned on $\mathbf{e}$, the probability $\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t \neq \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t$ is therefore $\frac{t}{q}|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')|$. Averaging over all $\mathbf{e}$, we have that, up to negligible error:

$$M_1^t(|\mathbf{y}\rangle\langle\mathbf{y}'|) = \left(1 - \frac{t}{q} \mathop{\mathbb{E}}_{\mathbf{e}}[|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')|]\right) |\mathbf{y}\rangle\langle\mathbf{y}'|$$

Notice then that $M_1^t(|\mathbf{y}\rangle\langle\mathbf{y}'|) = \frac{1}{d} M_1^{t\times d}(|\mathbf{y}\rangle\langle\mathbf{y}'|) + \left(1 - \frac{1}{d}\right)|\mathbf{y}\rangle\langle\mathbf{y}'|$. By linearity, we therefore prove Lemma 1. □

Note that the proof of Lemma 1 also demonstrates that M_0 and M_1^t commute, since their action on density matrices is just component-wise multiplication by a fixed matrix.

M_2^t: Given $\mathbf{A}$, to apply this measurement, first sample an LWE sample $\mathbf{b} = \mathbf{C} \cdot \mathbf{r}' + \mathbf{e}$. Then apply the measurement $\mathbf{y} \mapsto \lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t$. Let p_t be the probability that $\lfloor x \rceil_t = \lfloor y \rceil_t$ for uniformly random $x, y \in \mathbb{Z}_q$. Note that for any constant t, $p_t \leq t^{-1} + O(q^{-1})$.

Lemma 2. *For any constant t, $M_2^t(\rho)$ is statistically close to $M_0(M_1^t(\rho)) + p_t(\rho - M_0(\rho))$.*

Note that unlike Lemma 1, the expression in Lemma 2 does not correspond to a mixture of measurements applied to ρ. However, we will later see how to combine Lemma 2 with Lemma 1 to obtain such a mixture.

Proof. The proof proceeds similarly to Lemma 1. We consider the action of M_2^t on $|\mathbf{y}\rangle\langle\mathbf{y}'|$, and conclude that

$$M_2^t(|\mathbf{y}\rangle\langle\mathbf{y}'|) = \Pr_{\mathbf{b}}[\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t = \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t]$$

where the probability is over $\mathbf{b} = \mathbf{C} \cdot \mathbf{r}' + \mathbf{e}$. But now we have that

$$\mathbf{b} \cdot \mathbf{y} = \mathbf{r}'^T \cdot \mathbf{C}^T \mathbf{y} + \mathbf{e} \cdot \mathbf{y}$$
$$\mathbf{b} \cdot \mathbf{y}' = \mathbf{r}'^T \cdot \mathbf{C}^T \mathbf{y} + \mathbf{e} \cdot \mathbf{y}'$$

We consider two cases:

- $\mathbf{C}^T \cdot \mathbf{y} = \mathbf{C}^T \cdot \mathbf{y}'$. This case is essentially identical to the proof of Lemma 1, and we conclude that $\Pr_{\mathbf{b}}[\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t = \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t] = 1 - \frac{t}{q} \mathbb{E}_{\mathbf{e}}[|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')|]$. Note that for such $\mathbf{y}, \mathbf{y}'$, we also have

$$M_0(M_1^t(|\mathbf{y}\rangle\langle\mathbf{y}'|)) + p_t(|\mathbf{y}\rangle\langle\mathbf{y}'| - M_0(|\mathbf{y}\rangle\langle\mathbf{y}'|)) = M_1^t(M_0(|\mathbf{y}\rangle\langle\mathbf{y}'|)) + p_t \times 0 = 1 - \frac{t}{q} \mathbb{E}_{\mathbf{e}}[|\mathbf{e} \cdot (\mathbf{y} - \mathbf{y}')|] \ ,$$

  since M_0 is the identity on such $|\mathbf{y}\rangle\langle\mathbf{y}'|$. Thus, we have the desired equality for $\rho = |\mathbf{y}\rangle\langle\mathbf{y}'|$.
- $\mathbf{C}^T \cdot \mathbf{y} \neq \mathbf{C}_T \cdot \mathbf{y}'$. In this case, $\mathbf{b} \cdot \mathbf{y}$ and $\mathbf{b} \cdot \mathbf{y}'$ are independent and uniform over $\mathbb{Z}_p$. Therefore, $\Pr_{\mathbf{b}}[\lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t = \lfloor \mathbf{b} \cdot \mathbf{y}' \rceil_t] = p_t$. Note that for such $\mathbf{y}, \mathbf{y}'$, we also have

$$M_0(M_1^t(|\mathbf{y}\rangle\langle\mathbf{y}'|)) + p_t(|\mathbf{y}\rangle\langle\mathbf{y}'| - M_0(|\mathbf{y}\rangle\langle\mathbf{y}'|)) = 0 + p_t|\mathbf{y}\rangle\langle\mathbf{y}'| \ ,$$

  since $M_0(|\mathbf{y}\rangle\langle\mathbf{y}'|) = 0$ in this case.

Thus for each $|\mathbf{y}\rangle\langle\mathbf{y}'|$, we have the desired equality. By linearity, this thus extends to all ρ. □

Combining Lemmas 1 and 2, we obtain:

Corollary 1. *For any constants t, d, $M_2^t(\rho)$ is statistically close to $\frac{1}{d} M_0(M_1^{t \times d}(\rho)) + \left(1 - \frac{1}{d} - p_t\right) M_0(\rho) + p_t \rho$.*

For d such that $1 - \frac{1}{d} - p_t \geq 0$, this represents a mixture of measurements $M_0 \circ M_1^{t \times d}, M_0$, and the identity.

We are now ready to prove Theorem 1. Suppose there is an algorithm A that constructs a mixed state ρ, and then can distinguish ρ from $M_0(\rho)$ with (signed) advantage ϵ. Let d be a positive integer, to be chosen later. Let $\rho_0 = \rho$, and $\rho_i = M_1^{t \times d}(\rho_{i-1})$. Note that for any polynomial i, ρ_i can be efficiently constructed. Let $\epsilon_0 = \epsilon$, and ϵ_i be the (signed) distinguishing advantage of A when given ρ_i vs $M_0(\rho_i)$.

Let δ_i be the (signed) distinguishing advantage of A for $M_2^t(\rho_i)$ and $M_1^t(\rho_i)$. Write $g = 1 - \frac{1}{d} - p_t$. Invoking Lemma 1 and Corollary 1 with d, we have that

$$\delta_i = \frac{1}{d}\epsilon_{i+1} + g\epsilon_i$$

Now, we note that δ_i must be negligible, by the assumed hardness of $(n, m, q, \Sigma, \ell, D)$-LWE. Solving the recursion gives:

$$\epsilon_i(-dg)^{-i} = \epsilon - \frac{1}{d}\sum_{j=0}^{i-1}(-dg)^{-j}\delta_{j+1}$$

Next, assume d is chosen so that dg is a constant greater than 1. Define $T = \sum_{j=0}^{\lambda-1}(dg)^{-j} = \frac{dg}{dg-1} - 2^{-O(\lambda)}$. Consider the adversary A' for $(n, m, q, \Sigma, \ell, D)$-LWE, which does the following:

- On input $\mathbf{A}, \mathbf{S}, \mathbf{b}$, where $\mathbf{b} = \mathbf{A} \cdot \mathbf{r} + \mathbf{e}$ or $\mathbf{b} = \mathbf{C} \cdot \mathbf{r}' + \mathbf{e}$, it chooses $j \in [0, \lambda - 1]$ with probability $(dg)^{-j}/T$
- Then it constructs ρ according to A.
- Next, A' computes ρ_j by applying $M_1^{t \times d}$ to ρ for j times.
- Now A' applies the measurement $\mathbf{y} \mapsto \lfloor \mathbf{b} \cdot \mathbf{y} \rceil_t$ to ρ_j, obtaining ρ'_j.
- A' runs the distinguisher for A, obtaining a bit b
- A' outputs b if j is even, $1 - b$ if j is odd.

Note that if $\mathbf{b}$ is $\mathbf{A} \cdot \mathbf{r} + \mathbf{e}$, then $\rho' = M_1^t(\rho_i)$, and if $\mathbf{b}$ is $\mathbf{C} \cdot \mathbf{r}' + \mathbf{e}$, then $\rho' = M_2^t(\rho_i)$. Therefore, the distinguishing advantage of A' is:

$$\delta = \frac{1}{T}\sum_{j=0}^{\lambda-1}(-dg)^{-j}\delta_{j+1}$$

Thus, we have that

$$\epsilon_\lambda(-dg)^{-\lambda} = \epsilon - \frac{T}{d}\delta \ ,$$

Noting that ϵ_λ must trivially be in $[-1/2, 1/2]$, we have that:

$$|\delta| \geq \frac{d}{T}\left(|\epsilon| - \frac{1}{2}(dg)^{-\lambda}\right) \geq d\left(1 - \frac{1}{dg}\right)|\epsilon| - 2^{-O(\lambda)}$$

Thus, if A has non-negligible distinguishing advantage, so does A', breaking the $(n, m, q, \Sigma, \ell, D)$-LWE assumption. This completes the proof of Theorem 1. $\square$

5.4 Constructing $|\psi'_{\mathbf{u}}\rangle$

Here, we explain how to construct $|\psi'_{\mathbf{u}}\rangle$, given just the vector $\mathbf{y}$ that resulted from measuring it. We first observe that, since $|\psi'_{\mathbf{u}}\rangle$ has support only on vectors that differ from $\mathbf{y}$ by multiples of the columns of $\mathbf{S}$, we can write:

$$|\psi'_{\mathbf{u}}\rangle \propto \sum_{\mathbf{t}} \alpha_{\mathbf{y}+\mathbf{S}\cdot\mathbf{t}} |\mathbf{y} + \mathbf{S}\cdot\mathbf{t}\rangle$$

Where $\alpha_{\mathbf{y}}$ is the amplitude of $\mathbf{y}$ in $|\psi\rangle$. This gives a hint as to how to construct $|\psi'_{\mathbf{u}}\rangle$: create a superposition over short linear combinations of $\mathbf{S}$, and then use linear algebra to transition to a superposition over $\mathbf{y} + \mathbf{S}\cdot\mathbf{t}$, weighted according to α. The problem of course is that α may be arbitrary except for having support only on short vectors. Therefore, we do not expect to be able to construct $|\psi'_{\mathbf{u}}\rangle$ in full generality, and instead focus on special (but natural) cases, which suffice for our use.

Wide Gaussian Distributed. Suppose the initial state $|\psi\rangle$ is the discrete Gaussian over the integers: $|\psi\rangle = |\mathbf{\Psi}_{\mathbb{Z}^m,\Sigma,\mathbf{c}}\rangle$ for some center $\mathbf{c}$ and covariance matrix Σ. Then $|\psi'_{\mathbf{u}}\rangle$ is simply

$$|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\Sigma,\mathbf{c}}\rangle$$

Here, $\mathcal{L}$ is the integer lattice generated by the columns of $\mathbf{S}$, and $\mathcal{L}+\mathbf{y}$ is the lattice $\mathcal{L}$ shifted by $\mathbf{y}$. We can construct the state $|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\Sigma,\mathbf{c}}\rangle$ by first constructing $|\mathbf{\Psi}_{\mathcal{L},\Sigma,\mathbf{c}-\mathbf{y}}\rangle$, and then adding $\mathbf{y}$ to the superposition. Thus, as long as $\mathbf{s}_i^T \cdot \Sigma^{-1} \cdot \mathbf{s}_i \leq 1/\omega(\sqrt{\log\lambda})$ for all i, we can construct the necessary state.

Constant Dimension, Hyper-ellipsoid Bounded. Here, we restrict $\mathcal{L}$ to having a constant number of columns, but greatly generalize the distributions that can be handled.

A hyper-ellipsoid is specified by a positive definite matrix Σ, which defines the set $E_{\Sigma,\mathbf{c}} = \{\mathbf{y} : (\mathbf{y}-\mathbf{c})^T \cdot \mathbf{M} \cdot (\mathbf{y}-\mathbf{c}) \leq 1\}$.

Definition 5 (Good Hyper-ellipsoid). *A* good hyper-ellipsoid *for $|\psi\rangle$ is an $E_{\Sigma,\mathbf{c}}$ such that there exists a function $\eta(\lambda)$ and polynomials $p(\lambda), q(\lambda)$ such that, if $|\psi\rangle$ is measured to get a vector $\mathbf{y}$, then each of the following are true except with negligible probability:*

- *$\mathbf{y} \in E_{\Sigma,\mathbf{c}}$. In other words, $E_{\Sigma,\mathbf{c}}$ contains essentially all the mass of $|\psi\rangle$.*
- *$|\alpha_{\mathbf{x}}|^2 \leq \eta(\lambda)$. In other words, η is an approximate upper bound on $\alpha_{\mathbf{x}}$.*
- *If a random vector $\mathbf{x}$ is chosen from $E_{\Sigma,\mathbf{c}} \cap \{\mathbf{y} + \mathbf{S}\cdot\mathbf{t} : \mathbf{t} \in \mathbb{Z}^\ell\}$, then with probability at least $1/p(\lambda)$, $|\alpha_{\mathbf{x}}|^2 \geq \eta/q(\lambda)$. In other words, $E_{\Sigma,\mathbf{c}}$ doesn't contain too many points with mass too much lower than η.*

Taken together, a good hyper-ellipsoid is one that fits reasonably well around the $|\psi\rangle$. It must contain essentially all the support of $|\psi\rangle$, but can over-approximate it by a polynomial factor.

Lemma 3. *Suppose there is a good hyper-ellipsoid for $|\psi\rangle$, and that $\alpha_{\mathbf{y}}$ can be efficiently computed given any vector $\mathbf{y}$. Then there is a polynomial-time algorithm which constructs $|\psi'_{\mathbf{u}}\rangle$ from $\mathbf{y}$*

Proof. Let $E_{\Sigma,\mathbf{c}}$ be the good hyper-ellipsoid. Let $\mathcal{L}$ be the lattice generated by the columns of $\mathbf{S}$. By assumption, with overwhelming probability if we measure $|\psi\rangle$ to get $\mathbf{y}$, we have $\mathbf{y} \in E_{\Sigma,\mathbf{c}}$. Let $E_{\Sigma',\mathbf{c}'}$ be the ellipsoid that is the intersection of $E_{\Sigma,\mathbf{c}}$ and the affine space $\{\mathbf{y} + \mathbf{S} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{R}^\ell\}$.

Claim. There is PPT algorithm which, given $\mathbf{S}, \Sigma'$, computes $\mathbf{T} = \{\mathbf{r}_1, \cdots, \mathbf{r}_{\ell'}\}$ such that:

- $\mathbf{r}_i^T \cdot (\Sigma')^{-1} \cdot \mathbf{r}_i \leq 2$ for all $i \in [\ell']$, and
- $E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{T} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell'}\} = E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{S} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell}\}$.

Proof. Write $(\Sigma')^{-1}$ as $(\Sigma')^{-1} = \mathbf{U}^T \cdot \mathbf{U}$. Let $\mathbf{S}' = \{\mathbf{s}'_1 = \mathbf{U} \cdot \mathbf{s}_1, \ldots, \mathbf{s}'_\ell = \mathbf{U} \cdot \mathbf{s}_n\}$, and let $\mathcal{L}'$ be the lattice generated by $\mathbf{S}'$. Since ℓ is constant, we can find shortest vectors in $\mathcal{L}'$ in polynomial time. Therefore, compute $\mathbf{r}'_1, \ldots, \mathbf{r}'_\ell$ such that $\mathbf{r}'_i$ is the shortest vector in $\mathcal{L}'$ that is linearly independent from $\{\mathbf{r}'_1, \ldots, \mathbf{r}'_{i-1}\}$. Then let ℓ' be such that $|\mathbf{r}'_{\ell'}|^2 \leq 2$, but $|\mathbf{r}'_{\ell'+1}|^2 > 2$, or $\ell' = \ell$ if no such ℓ' exists.

Finally, let $\mathbf{r}_i = \mathbf{U}^{-1} \cdot \mathbf{r}'_i$. Clearly, we have that $\mathbf{r}_i^T \cdot (\Sigma')^{-1} \cdot \mathbf{r}_i \leq 2$. It remains to show that $E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{T} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell'}\} = E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{S} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell}\}$. First, we notice that the lattice $\mathcal{L}(\mathbf{T})$ spanned by $\mathbf{T}$ is a sub-lattice of $\mathcal{L}(\mathbf{S})$ spanned by $\mathbf{S}$. So one containment is trivial. Now assume toward contradiction that there is a $\mathbf{x} \in E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{S} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell}\}$ that is not in $E_{\Sigma',\mathbf{c}'} \cap \{\mathbf{y} + \mathbf{T} \cdot \mathbf{t} : \mathbf{t} \in \mathbb{Z}^{\ell'}\}$. This means $\mathbf{x} - \mathbf{y}$ is in $\mathcal{L}(\mathbf{S})$. We also have that $(\mathbf{y} - \mathbf{c}')^T \cdot (\Sigma')^{-1} \cdot (\mathbf{y} - \mathbf{c}') \leq 1$ (since and $(\mathbf{x} - \mathbf{c}')^T \cdot (\Sigma')^{-1} \cdot (\mathbf{x} - \mathbf{c}') \leq 1$. By the triangle inequality, we have therefore that $(\mathbf{x} - \mathbf{y})^T \cdot (\Sigma')^{-1} \cdot (\mathbf{x} - \mathbf{y}) \leq 2$.

But then we have that $\mathbf{U} \cdot (\mathbf{x} - \mathbf{y})$ has norm at most 2, lies in $\mathcal{L}'$, and is linearly independent of $\{\mathbf{r}'_1, \ldots, \mathbf{r}'_{\ell'}\}$. This contradicts that $\mathbf{r}'_{\ell'+1}$ (which has norm squared strictly greater than 2) is a shortest vector linearly independent of $\{\mathbf{r}'_1, \ldots, \mathbf{r}'_{\ell'}\}$. This completes the proof of the claim. □

We now return to proving Lemma 3. Let $\beta = \omega(\log \lambda)$. We construct $|\psi'_{\mathbf{u}}\rangle$ in three steps:

- We first construct a state negligibly close to $|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\beta\Sigma',\mathbf{c}'}\rangle$, as we did in the Gaussian-distributed case above.
- We then construct the state $|E\rangle$, defined as the uniform superposition over the intersection of $\mathcal{L} + \mathbf{y}$ and $E_{\Sigma',\mathbf{c}'}$. $|E\rangle$ will be obtained from $|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\beta\Sigma',\mathbf{c}'}\rangle$ via a measurement.
- Construct $|\psi'_{\mathbf{u}}\rangle$ from $|E\rangle$. This also will be obtained via a measurement.

We now describe the two measurements. We start from the second. Let η, p, q be the values guaranteed by the goodness of $E_{\Sigma,\mathbf{c}}$. Define $\eta_{\mathbf{x}} = 1/\eta$ if $|\alpha_{\mathbf{x}}|^2 \leq \eta$, and otherwise $\eta_{\mathbf{x}} = 1/|\alpha_{\mathbf{x}}|^2$. To obtain $|\psi'_{\mathbf{u}}\rangle$ from $|E\rangle$, we apply the following map in superposition and measure the second register:

$$|\mathbf{x}\rangle \mapsto |\mathbf{x}\rangle \left(\sqrt{\eta_{\mathbf{x}}}\alpha_{\mathbf{x}}|0\rangle + \sqrt{1 - |\eta_{\mathbf{x}}\alpha_{\mathbf{x}}|^2}|1\rangle\right)$$

Suppose for the moment that $\eta_{\mathbf{x}} = 1/\eta$ for all $\mathbf{x}$. Then conditioned on the measurement outcome being 0, the resulting state is exactly $|\psi'_{\mathbf{u}}\rangle$. By the guarantee that $E_{\Sigma,\mathbf{c}}$ is good, we have that except with negligible probability over the choice of $\mathbf{y}$, all but a negligible fraction of the support of $|\psi'_{\mathbf{u}}\rangle$ satisfies $\eta_{\mathbf{x}} = 1/\eta$. Therefore, we will assume (with negligible error) this is the case. The probability the measurement is 0 (over the choice of $\mathbf{y}$ as well) is $\mathbb{E}_{\mathbf{x}\leftarrow E_{\Sigma',\mathbf{c}'}}[\alpha_{\mathbf{x}}^2/\eta]$, which, with probability at least $1/p$ over the choice of $\mathbf{y}$, is at least $1/q$. Thus, the overall probability of outputting 0 is inverse polynomial, and in this case we produce a state negligibly close to $|\psi'_{\mathbf{u}}\rangle$.

It remains to construct $|E\rangle$ from $|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\beta\Sigma',\mathbf{c}'}\rangle$. This follows a very similar rejection-sampling argument. Let

$$\gamma_{\mathbf{x}} = \begin{cases} e^{-\pi/\beta} \times \sqrt{e^{\pi(\mathbf{x}-\mathbf{c}')^T\cdot(\beta\Sigma')^{-1}\cdot(\mathbf{x}-\mathbf{c}')}} & \text{if } (\mathbf{x}-\mathbf{c}')^T\cdot(\Sigma')^{-1}\cdot(\mathbf{x}-\mathbf{c}') \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Note that $0 \leq \gamma_{\mathbf{x}} \leq 1$. Now apply to $|\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\beta\Sigma',\mathbf{c}'}\rangle$ the map $|\mathbf{x}\rangle \mapsto |\mathbf{x}\rangle(\gamma_{\mathbf{x}}|0\rangle + \sqrt{1-\gamma_{\mathbf{x}}^2}|1\rangle)$, and measure the second coordinate. If the measurement outcome is 0, then the resulting state is exactly $|E\rangle$. For $\mathbf{x} \in E_{\Sigma',\mathbf{c}'}$, we have $\gamma_{\mathbf{x}} \geq e^{-\pi/\beta} \geq 1 - o(1)$. Therefore, the probability the measurement outputs 0 is at least $1-o(1)$ times the probability measuring $\mathbf{\Psi}_{\mathcal{L}+\mathbf{y},\beta\Sigma',\mathbf{c}'}$ produces an $\mathbf{x} \in E_{\Sigma',\mathbf{c}'}$. This latter probability is $O_\ell(\beta^{-\ell/2})$, where the constant hidden by the big O depends on ℓ. Since ℓ is constant and β is polynomial (in fact, sub-polynomial), the overall probability is polynomial. This completes the construction of $|\psi'_{\mathbf{u}}\rangle$ and the proof of Lemma 3. □

Applying to [KLS22]: To avoid confusion, we first refer the readers to our alternate view on [KLS22] scheme in section C.3 and then we will see how to apply our attack onto their scheme in section C.5 of the full version.

6 Invariant Money

From this section on, we discuss our positive results on quantum money/lightning.

We now describe our framework for instantiating quantum money using invariants, or more precisely what we call *walkable* invariants.

Let X, Y be sets, and $I : X \to Y$ an efficiently computable function from X to Y. I will be called the "invariant." We will additionally assume a collection of permutations $\sigma_i : X \to X$ indexed by $i \in [r]$ for some integer r, with the property that the permutations respect the invariant:

$$I(\ \sigma_i(x)\ ) = I(x), \forall i \in [r]$$

In other words, action by each σ_i preserves the value of the invariant. We require that σ_i is efficiently computable given i. r may be polynomial or may be exponential. To make the formalism below simpler, we will be implicitly assuming that there exists a perfect matching between the σ_i such that for any matched

$\sigma_i, \sigma_{i'}$, we have $\sigma_{i'} = \sigma_i^{-1}$. Moreover, i' can be found given i. This can be relaxed somewhat to just requiring that σ_i^{-1} can be efficiently computed given i, but requires a slightly more complicated set of definitions.

Given a point x, the orbit of x, denoted $O_x \subseteq X$, is the set of all z such that there exists a non-negative integer k and $i_1, \ldots, i_k \in [r]$ such that $z = \sigma_{i_k}(\sigma_{i_{k-1}}(\cdots \sigma_{i_1}(x)\cdot))$. In other words, O_x is the set of all z "reachable" from x by applying some sequence of permutations. Note that $I(z) = y$ for any $z \in O_x$. We will therefore somewhat abuse notation, and define $I(O_x) = y$. We also let P_y be the set of pre-images of y: $P_y = \{x \in X : I(x) = y\}$.

We will additionally require a couple properties, which will be necessary for the quantum money scheme to compile:

- **Efficient Generation of Superpositions:** It is possible to construct the uniform superposition over X: $|X\rangle := \frac{1}{\sqrt{|X|}} \sum_{x \in X} |x\rangle$.
- **Mixing Walks:** For an orbit O, with a slight abuse of notation let $\sigma_{O,i}$ be the (possibly exponentially large) permutation matrix associated with the action by σ_i on O. Then let $M_O = \frac{1}{r} \sum_{i \in [r]} \sigma_{O,i}$ be the component-wise average of the matrices. Let $\lambda_1(O), \lambda_2(O)$ be the largest two eigenvalues by absolute value[18], counting multiplicities. Note that $\lambda_1(O) = 1$, with corresponding eigenvector the all-1's vector. We need that there is an inverse polynomial δ such that, for every orbit O, $\lambda_2(O) \leq 1 - \delta$. This is basically just a way of saying that a random walk on the orbit using the σ_i mixes in polynomial time.

We call such a structure above a walkable invariant.

6.1 Quantum Money from Walkable Invariants

We now describe the basic quantum money scheme.

Minting. To mint a note, first construct the uniform superposition $|X\rangle$ over X. Then apply the invariant I in superposition and measure, obtaining a string y, and the state collapsing to:

$$|P_y\rangle := \frac{1}{\sqrt{|P_y|}} \sum_{x \in P_y} |x\rangle$$

This is the quantum money state, with serial number y.

Verification. To verify a supposed quantum money state $|\phi\rangle$ with serial number y, we do the following.

- First check that the support of $|\phi\rangle$ is contained in P_y. This is done by simply applying the invariant I in superposition, and measuring if the output is y. If the check fails immediately reject.

[18] They are real-valued, since M_O is symmetric, owing to the fact that we assumed the σ_i are perfectly matched into pairs that are inverses of each other.

– Then apply the projective measurement given by the projection $\sum_{O \subseteq P_y} |O\rangle\langle O|$, where O ranges over the orbits contained in P_y, and $|O\rangle := \frac{1}{\sqrt{|O|}} \sum_{x \in O} |x\rangle$.
In other words, project onto states where, for each orbit, the weights of x in that orbit are all identical; weights between different orbits are allowed to be different.
We cannot perform this measurement exactly, but we can perform it approximately using the fact that $\lambda_2(O) \leq 1-\delta$. This is described in Section 6.2 below. Outside of Section 6.2, we will assume for simplicity that the measurement is provided exactly.
If the projection rejects, reject the quantum money state. Otherwise accept.

It is hopefully clear that honestly-generated money states pass verification. Certainly their support will be contained in P_y, and they apply equal weight to each element in an orbit (and in fact, equal weight across orbits).

6.2 Approximate Verification

Here, we explain how to approximately perform the verification projection $V = \sum_{O \subseteq P_y} |O\rangle\langle O|$, using the fact that $\lambda_2(0) \leq 1 - \delta$ for all O. The algorithm we provide is an abstraction of the verification procedure of [FGH+12], except that that work presented the algorithm without any analysis. We prove that the algorithm is statistically close to the projection V, provided the mixing condition $\lambda_2(0) \leq 1 - \delta$ is met.

Theorem 2. *Assume $\lambda_2(0) \leq 1 - \delta$ for all O, for some inverse-polynomial δ. Then there is a QPT algorithm $\tilde{V}$ such that, for any state $|\psi\rangle$, if we let $|\psi'\rangle$ be the un-normalized post-measurement state from applying $\tilde{V}$ to $|\psi\rangle$ in the case $\tilde{V}$ accepts, then $|\psi'\rangle$ is negligibly close to $V|\psi\rangle$.*

We refer the readers to section E.1 of the full version for the proof due to restriction on the space.

6.3 Hardness Assumptions

We rely on two hardness assumptions in our inviant money scheme: the *path-finding assumpion* and the *knowledge of path* assumption. Due to space constraints, we refer the readers to E.2 for the presentation on our hardness assumptions needed.

Informally speaking, the path-finding assumption states that, given some adversarially sampled x in a set X, it is hard for any efficient adversary, given a random $x' \in X$ such that there exists some σ such that $\sigma(x) = x'$, to find such a σ.

The knowledge of path assumption can be thought of as a quantum analogue to the (classical) knowledge of exponent assumption. We define two different versions of the knowledge of path assumption to account for the fact that some of our invariants could be invertible.

6.4 Security

Theorem 3. *Assuming the Path-Finding assumption and the Knowledge of Path Assumption , the scheme above is secure quantum lightning. If the invariant is invertible, then assuming the Path-Finding assumption , the Knowledge of Path Assumption for Invertible Invariants , and the Inversion Inverting assumption , the scheme above is secure quantum lightning.*

We refer the readers to E.3 of the full version for the formal statements of the above assumptions and the proof on the above theorem.

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 Annual ACM Symposium on Theory of Computing*, pages 41–60, New York, NY, USA, May 19–22, 2012. ACM Press.

ADMP20. Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part II*, volume 12492 of *Lecture Notes in Computer Science*, pages 411–439, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.

AGKZ20. Ryan Amos, Marios Georgiou, Aggelos Kiayias, and Mark Zhandry. One-shot signatures and applications to hybrid quantum/classical authentication. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 255–268, 2020.

AL21. Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 501–530, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.

ALL+21. Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 526–555, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

BB87. Charles H. Bennett and Gilles Brassard. Quantum public key distribution reinvented. *SIGACT News*, 18(4):51–53, July 1987.

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 Annual Symposium on Foundations of Computer Science*, pages 320–331, Paris, France, October 7–9, 2018. IEEE Computer Society Press.

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.

BF11. Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th International Conference on Theory and Practice of Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 1–16, Taormina, Italy, March 6–9, 2011. Springer, Heidelberg, Germany.

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: 16th Theory of Cryptography Conference, Part II*, volume 11240 of *Lecture Notes in Computer Science*, pages 544–574, Panaji, India, November 11–14, 2018. Springer, Heidelberg, Germany.

BGS13. Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs - (extended abstract). In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 344–360, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

BLP+13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 575–584, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.

CD22. Wouter Castryck and Thomas Decru. An efficient key recovery attack on sidh (preliminary version). *Cryptology ePrint Archive*, 2022.

CLLZ21. Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

Col09. Roger Colbeck. Quantum and relativistic protocols for secure multi-party computation, 2009.

Col19. Andrea Coladangelo. Smart contracts meet quantum cryptography, 2019.

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.

CS20. Andrea Coladangelo and Or Sattath. A quantum money solution to the blockchain scalability problem. *Quantum*, 4:297, 2020.

CY14. Matthew Coudron and Henry Yuen. Infinite randomness expansion with a constant number of devices. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 427–436, New York, NY, USA, May 31 – June 3, 2014. ACM Press.

Dam92. Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages 445–456, Santa Barbara, CA, USA, August 11–15, 1992. Springer, Heidelberg, Germany.

FGH+10. Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, Daniel Nagaj, and Peter Shor. Quantum state restoration and single-copy tomography for ground states of hamiltonians. *Physical review letters*, 105(19):190503, 2010.

FGH+12. Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, and Peter W. Shor. Quantum money from knots. In Shafi Goldwasser, editor, *ITCS 2012: 3rd Innovations in Theoretical Computer Science*, pages 276–289, Cambridge, MA, USA, January 8–10, 2012. Association for Computing Machinery.

GGH15. Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 498–527, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, May 17–20, 2008. ACM Press.

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

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: 1st Innovations in Computer Science*, pages 20–31, Tsinghua University, Beijing, China, January 5–7, 2010. Tsinghua University Press.

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, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 326–355, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.

MM22. Luciano Maino and Chloe Martindale. An attack on sidh with arbitrary starting curve. *Cryptology ePrint Archive*, 2022.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.

Rob21. Bhaskar Roberts. Security analysis of quantum lightning. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 562–567, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.

Rob22. Damien Robert. Breaking sidh in polynomial time. *Cryptology ePrint Archive*, 2022.

RS19. Roy Radian and Sattath. Semi-quantum money. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, AFT '19, page 132–146, New York, NY, USA, 2019. Association for Computing Machinery.

Unr16. Dominique Unruh. Computationally binding quantum commitments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 497–527, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.

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

WW21. Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 127–156, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.

Zha19. Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 408–438, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.