Quantum Money from Hidden Subspaces*

Scott Aaronson† Paul Christiano‡

ABSTRACT

Forty years ago, Wiesner pointed out that quantum mechanics raises the striking possibility of money that cannot be counterfeited according to the laws of physics. We propose the first quantum money scheme that is

(1) *public-key*—meaning that anyone can verify a banknote as genuine, not only the bank that printed it, and

(2) *cryptographically secure*, under a "classical" hardness assumption that has nothing to do with quantum money.

Our scheme is based on *hidden subspaces*, encoded as the zero-sets of random multivariate polynomials. A main technical advance is to show that the "black-box" version of our scheme, where the polynomials are replaced by classical oracles, is *unconditionally* secure. Previously, such a result had only been known relative to a *quantum* oracle (and even there, the proof was never published).

Even in Wiesner's original setting—quantum money that can only be verified by the bank—we are able to use our techniques to patch a major security hole in Wiesner's scheme. We give the first private-key quantum money scheme that allows unlimited verifications and that remains unconditionally secure, even if the counterfeiter can interact adaptively with the bank.

Our money scheme is simpler than previous public-key quantum money schemes, including a knot-based scheme of Farhi et al. The verifier needs to perform only two tests, one in the standard basis and one in the Hadamard basis—matching the original intuition for quantum money, based on the existence of complementary observables.

Our security proofs use a new variant of Ambainis's quantum adversary method, and several other tools that might be of independent interest.

Categories and Subject Descriptors

F.1.2 [**Theory of Computation**]: Computation by Abstract Devices—*Modes of Computation*; E.3 [**Data**]: Data Encryption

General Terms

Theory, Security

1. INTRODUCTION

"Information wants to be free"—this slogan expresses the idea that classical bits, unlike traditional economic goods, can be copied an unlimited number of times. The copyability of classical information is one of the foundations of the digital economy, but it is also a nuisance to governments, publishers, software companies, and others who wish to prevent copying. Today, essentially all electronic commerce involves a trusted third party, such as a credit card company, to mediate transactions. Without such a third party entering at *some* stage, it is impossible to prevent electronic cash from being counterfeited, regardless of what cryptographic assumptions one makes.[1]

Famously, though, quantum bits do *not* "want to be free" in the same sense that classical bits do: in many respects, they behave more like gold, oil, or other traditional economic goods. Indeed, the *No-Cloning Theorem*, which is an immediate consequence of the linearity of quantum mechanics, says that there is no physical procedure that takes as input an unknown[2] quantum pure state $|\psi\rangle$, and that produces as output two unentangled copies of $|\psi\rangle$, or even a close approximation thereof. The No-Cloning Theorem is closely related to the *uncertainty principle*, which says that there exist "complementary" properties of a quantum state (for example, its position and momentum) that cannot both be measured to unlimited accuracy.[3]

*Extended abstract. For the full version, see www.scottaaronson.com/papers/money.pdf

†MIT. Email: aaronson@csail.mit.edu. This material is based upon work supported by the National Science Foundation under Grant No. 0844626. Also supported by a DARPA YFA grant, an NSF STC grant, a TIBCO Chair, and a Sloan Fellowship.

‡MIT. Email: paulfchristiano@gmail.com.

STOC'12, May 19–22, 2012, New York, New York, USA.

[1]The recent Bitcoin system is an interesting illustration of this principle: it gets rid of the centralized third party, but still uses a "third party" distributed over the community of Bitcoin users.

[2]The adjective "unknown" is needed because, if we knew a classical description of a procedure to *prepare* $|\psi\rangle$, then of course we could run that procedure multiple times to prepare multiple copies.

[3]Indeed, if we could copy $|\psi\rangle$, then we could violate the uncertainty principle by measuring one observable (such as

1.1 The History of Quantum Money

But can one actually *exploit* the No-Cloning Theorem to achieve classically-impossible cryptographic tasks? This question was first asked by Wiesner [39], in a remarkable paper written around 1970 (but only published in 1983) that arguably founded quantum information science. In that paper, Wiesner proposed a scheme for *quantum money* that would be physically impossible to clone. In Wiesner's scheme, each "banknote" would consist of a classical serial number s, together with a quantum state $|\psi_s\rangle$ consisting of n unentangled qubits, each one $|0\rangle$, $|1\rangle$, $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$, or $\frac{|0\rangle-|1\rangle}{\sqrt{2}}$ with equal probability. The issuing bank would maintain a giant database, which stored a classical description of $|\psi_s\rangle$ for each serial number s. Whenever someone wanted to *verify* a banknote, he or she would take it back to the bank—whereupon the bank would use its knowledge of how $|\psi_s\rangle$ was prepared to measure each qubit in the appropriate basis, and check that it got the correct outcomes. On the other hand, it can be proved [31] that someone who did *not* know the appropriate bases could copy the banknote with success probability at most $(3/4)^n$.

Though historically revolutionary, Wiesner's money scheme suffered at least three drawbacks:

(1) **The "Verifiability Problem":** The only entity that can verify a banknote is the bank that printed it.

(2) **The "Online Attack Problem":** A counterfeiter able to submit banknotes for verification, and get them back afterward, can easily break Wiesner's scheme ([28, 3]; see also Section 7).

(3) **The "Giant Database Problem":** The bank needs to maintain a database with an entry for every banknote in circulation.

In followup work in 1982, Bennett, Brassard, Breidbart, and Wiesner [14] (henceforth BBBW) at least showed how to eliminate the giant database problem: namely, by generating the state $|\psi_s\rangle = \left|\psi_{f_k(s)}\right\rangle$ using a *pseudorandom function* f_k, with key k known only by the bank. Unlike Wiesner's original scheme, the BBBW scheme is no longer *information-theoretically* secure: a counterfeiter can recover k given exponential computation time. On the other hand, a counterfeiter cannot break the scheme in polynomial time, unless it can *also* distinguish f_k from a random function.

These early ideas about quantum money inspired the field of *quantum cryptography* [13]. But strangely, the subject of quantum money itself lay dormant for more than two decades, even as interest in quantum computing exploded. However, the past few years have witnessed a "quantum money renaissance." Some recent work has offered partial solutions to the verifiability problem: for example, Mosca and Stebila [32] suggested that the bank use a *blind quantum computing* protocol to offload the verification of banknotes to local merchants, while Gavinsky [23] proposed a variant of Wiesner's scheme that requires only *classical* communication between the merchant and bank.

However, most of the focus today is on a more ambitious goal: namely, creating what Aaronson [3] called *public-key quantum money*, or quantum money that *anyone* could authenticate, not just the bank that printed it. As with public-key cryptography in the 1970s, it is far from obvious *a priori* whether public-key quantum money is possible at all. Can a bank publish a description of a quantum circuit that lets people feasibly *recognize* a state $|\psi\rangle$, but does not let them feasibly prepare or even copy $|\psi\rangle$?

Aaronson [3] gave the first formal treatment of public-key quantum money, as well as related notions such as copy-protected quantum software. He proved that there exists a *quantum oracle* relative to which secure public-key quantum money is possible. Unfortunately, that result, though already involved, did not lead in any obvious way to an explicit (or "real-world") quantum money scheme.[4] He raised as an open problem whether secure public-key quantum money is possible relative to a *classical* oracle. In the same paper, Aaronson also proposed an explicit scheme, based on random stabilizer states, but could not offer any evidence for its security. And indeed, the scheme was broken about a year afterward by Lutomirski et al. [30], using an algorithm for finding planted cliques in random graphs due to Alon, Krivelevich, and Sudakov [7].

Recently, Farhi et al. [22] took a completely different approach to public-key quantum money. They proposed a quantum money scheme based on *knot theory*, where each banknote is a superposition over exponentially-many oriented link diagrams. Within a given banknote, all the link diagrams L have the same Alexander polynomial $p(L)$ (a certain knot invariant).[5] This $p(L)$, together with a digital signature of $p(L)$, serves as the banknote's "classical serial number." Besides the unusual mathematics employed, the work of Farhi et al. [22] (building on [30]) also developed an idea that will play a major role in our work. That idea is to construct public-key quantum money schemes by composing two "simpler" ingredients: first, objects that we call *mini-schemes*; and second, classical digital signature schemes.

The main disadvantage of the knot-based scheme, which it shares with every previous scheme, is that no one can say much about its security—other than that it has not yet been broken, and that various known counterfeiting strategies fail. Indeed, even characterizing *which quantum states Farhi et al.'s verification procedure accepts* remains a difficult open problem, on which progress seems likely to require major advances in knot theory! In other words, there might be states that look completely different from "legitimate banknotes," but are still accepted with high probability.

In followup work, Lutomirski [29] proposed an "abstract" version of the knot scheme, which gets rid of the link diagrams and Alexander polynomials, and simply uses a classical oracle to achieve the same purposes. Lutomirski raised the challenge of proving that this *oracle* scheme is secure—in which case, it would have yielded the first public-key

position) on some copies, and a complementary observable (such as momentum) on other copies. Conversely, if we could measure all the properties of $|\psi\rangle$ to unlimited accuracy, then we could use the measurement results to create additional copies of $|\psi\rangle$.

[4] Also, the proof of Aaronson's result never appeared—an inexcusable debt that this paper finally repays, with interest.

[5] Instead of knots, Farhi et al. [22] could also have used, say, superpositions over n-vertex *graphs* having the same eigenvalue spectrum. But in that case, their scheme would have been breakable, the reason being that the *graph isomorphism problem* is easy for random graphs. By contrast, it is not known how to solve *knot isomorphism* efficiently, even with a quantum computer and even for random knots.

quantum money scheme that was proven secure relative to a classical oracle. Unfortunately, proving the security of Lutomirski's scheme remains open, and seems hard.[6]

As alluded to earlier, there is already some research on ways to *break* quantum money schemes. Besides the papers by Lutomirski [28] and Lutomirski et al. [30] mentioned before, let us mention the beautiful work of Farhi et al. on *quantum state restoration* [21]. As we discuss in Section 7, quantum state restoration can be used to break many public-key quantum money schemes: roughly speaking, any scheme where the banknotes contain only limited entanglement, and where verification consists of a rank-1 projective measurement. This fact explains why our scheme, like the knot-based scheme of Farhi et al. [22], will require highly-entangled banknotes.

1.2 The Challenge

Work over the past few years has revealed a surprising richness in the quantum money problem—both in the ideas that have been used to construct public-key quantum money schemes, *and* in the ideas that have been used to break them. Of course, this record also underscores the need for caution! To whatever extent we can, we ought to hold quantum money schemes to modern cryptographic standards, and not be satisfied with "we tried to break it and failed."

It is easy to see that, if public-key quantum money is possible, then it must rely on *some* computational assumption, in addition to the No-Cloning Theorem.[7] The best case would be to show that secure, public-key quantum money is possible, *if* (for example) there exist one-way functions resistant to quantum attack. Unfortunately, we seem a long way from showing anything of the kind. The basic problem is that *uncloneability* is a novel cryptographic requirement: something that would not even make sense in a classical context. Indeed, work by Farhi et al. [21] and Aaronson [3] has shown that it is sometimes possible to copy quantum banknotes, via attacks that do not even *measure* the banknotes in an attempt to learn a classical secret! Rather, these attacks simply perform some unitary transformation on a legitimate banknote $|\$\rangle$ together with an ancilla $|0\rangle$, the end result of which is to produce $|\$\rangle^{\otimes 2}$. Given such a strange attack, how can one deduce the failure of any "standard" cryptographic assumption?

Yet despite the novelty of the quantum money problem—or perhaps because of it—it seems reasonable to want *some* non-tautological evidence that a public-key quantum money scheme is secure. A minimal wish-list might include:

(1) Security under *some* plausible assumption, of a sort cryptographers know how to evaluate. Such an assumption should talk only about computing a classical output from a classical input; it should have nothing to do with cloning of quantum states.

(2) A proof that the money scheme is secure against *black-box counterfeiters*: those that do not exploit the structure of some cryptographic function f used in verifying the banknotes.

(3) A "simple" verification process, which accepts all valid banknotes $|\$\rangle$ with probability 1, and rejects all banknotes that are far from $|\$\rangle$.

1.3 Our Results

Our main contribution is a new public-key quantum money scheme, which achieves all three items in the wish-list above, and which is the first to achieve (1) or (2). Regardless of whether our particular scheme stands or falls, we introduce at least four techniques that should be useful for the design and analysis of *any* public-key quantum money scheme. These are:

- The "inner-product adversary method," a new variant of Ambainis's quantum adversary method [8] that can be used to rule out black-box counterfeiting strategies.
- A notion of *mini-schemes*, and a proof that (together with standard cryptographic assumptions) these objects imply full-fledged quantum money schemes.
- A method to *amplify* weak counterfeiters into strong ones, so that one only needs to rule out the latter to show security.
- A new connection between the security of quantum money schemes and *direct-product assumptions* in cryptography.

A second contribution is to construct the first *private-key* quantum money schemes that remain *unconditionally* secure, even if the counterfeiter can interact adaptively with the bank. This gives the first solution to the "online attack problem," a major security hole in the Wiesner [39] and BBBW [14] schemes pointed out by Lutomirski [28] and Aaronson [3]. These private-key schemes are direct adaptations of our public-key scheme.

In more detail, our quantum money scheme is based on *hidden subspaces* of the vector space $\mathbb{F}_2^n$. Each of our money states is a uniform superposition of the vectors in a random $n/2$-dimensional subspace $A \leq \mathbb{F}_2^n$. We denote this superposition by $|A\rangle$. Crucially, we can recognize the state $|A\rangle$ using only membership oracles for A and for its dual subspace $A^\perp$. To do so, we apply the membership oracle for A, then a Fourier transform, then the membership oracle for $A^\perp$, and then a second Fourier transform to restore the original state. We prove that this operation computes a rank-1 projection onto $|A\rangle$.

Underlying the security of our money schemes is the assertion that the states $|A\rangle$ are difficult to clone, even given membership oracles for A and $A^\perp$. Or more concretely: *any quantum algorithm that maps $|A\rangle$ to $|A\rangle^{\otimes 2}$ must make $2^{\Omega(n)}$ queries to the $A, A^\perp$ oracles.*

In order to prove this statement, we introduce a new method for proving lower bounds on quantum query complexity, which we call the *inner-product adversary method*. This technique considers a single counterfeiting algorithm

[6] One way to understand the difficulty is that any security proof for Lutomirski's scheme would need to contain Aaronson's quantum lower bound for the *collision problem* [1] as a (tiny) special case. The lower bound for the collision problem is proved using the polynomial method of Beals et al. [11]. In this work, by contrast, we will only manage to prove the security of *our* oracle scheme using a specially-designed variant of Ambainis's quantum adversary method [8]. Despite great progress in quantum lower bounds over the past decade, it is still not known (except implicitly) how to prove the collision lower bound using an adversary argument.

[7] This is because a counterfeiter with unlimited time could simply search for a state $|\psi\rangle$ that the (publicly-known) verification procedure accepted.

being run in parallel to clone two distinct states $|A\rangle$ and $|A'\rangle$, with each having access to the membership oracles for $A, A^\perp$ or $A', A'^\perp$, as appropriate. To measure how much progress the algorithm has made, we consider the inner product between the states produced by the parallel executions: because $\langle A|^{\otimes 2}|A'\rangle^{\otimes 2} < \langle A|A'\rangle$ for many pairs of subspaces A, A', in order to succeed a counterfeiter will have to reduce this inner product substantially. We prove that when averaged over a suitable distribution of pairs A, A', the *expected inner product* between the two states produced by the counterfeiter cannot decrease too much with a single query to the membership oracles. We conclude that in order to produce $|A\rangle^{\otimes 2}$ given $|A\rangle$ and membership oracles for $A, A^\perp$, a counterfeiter must use exponentially many queries.

Having ruled out the possibility of nearly perfect cloning, we introduce a new amplification protocol, which allows us to transform a counterfeiter who succeeds with $\Omega\left(1/\operatorname{poly}(n)\right)$ success probability into a counterfeiter who succeeds with probability arbitrarily close to 1. This technique is based on combining standard Grover search with a monotonic state amplification protocol of Tulsi, Grover, and Patel [38], to obtain monotonic convergence with the quadratic speedup of Grover search.[8] Combining this amplification with the inner-product adversary method, and applying a random linear transformation to convert the counterfeiter's worst case to its average case, we conclude that no counterfeiting algorithm can succeed with any non-negligible probability on a non-negligible fraction of states $|A\rangle$.

Using these results, how do we produce a secure quantum money scheme? We now need to step back, and discuss some general constructions of this paper that have nothing to do with hidden subspaces in particular. Before constructing a quantum money scheme, we first introduce the notion of a *quantum money mini-scheme*, a formalization of the setting in which the bank issues only a single money state and maintains no secret information. Formally, a mini-scheme is a protocol Bank for outputting pairs (s, ρ_s) and a verification procedure Ver_s for identifying ρ_s. We say a mini-scheme is *complete* if the state ρ_s passes the verification Ver_s with high probability, and we say the scheme is *secure* if furthermore no counterfeiter can take a single state ρ_s, and produce two (possibly-entangled) states ρ_1 and ρ_2 which simultaneously pass the verification procedure with non-negligible probability.

In the case of hidden subspace money, for example, we can use our uncloneability result to produce a secure mini-scheme relative to a classical oracle. The algorithm Bank queries the classical oracle to obtain a serial number s and the description of a subspace A. Using this description it prepares $|A\rangle$, and publishes $(s, |A\rangle)$. The verification procedure uses the serial number s as an index into another classical oracle, which allows it to test membership in A and $A^\perp$. We prove that the uncloneability of the states $|A\rangle$ implies that this mini-scheme is secure.

Crucially, we also give a general reduction from quantum money schemes to mini-schemes, based on combining a mini-scheme with a secure signature scheme. The bank maintains a secret key for the signature scheme, and to issue a banknote, it runs Bank to produce a pair (s, ρ_s), then digitally signs the serial number s. Special cases of this reduction appeared in [22, 30], but we provide the first rigorous security proof.

By combining this reduction with our mini-scheme, we are able to obtain a "black-box" public key quantum money scheme relative to a classical oracle, which is unconditionally secure:

Theorem (Security of Hidden Subspace Money)**.** *Relative to some (classical) oracle A, there exists a secure public-key quantum money scheme.*

More precisely, there is an algorithm KeyGen^A which outputs pairs $(k_{\text{private}}, k_{\text{public}})$, an algorithm $\mathsf{Bank}^A(k_{\text{private}})$ which generates a "quantum banknote" $|\$\rangle$, and a verification algorithm $\mathsf{Ver}^A(k_{\text{public}}, |\$\rangle)$ which tests the authenticity of a purported banknote. These algorithms have the following properties:

Completeness: *If $(k_{\text{private}}, k_{\text{public}})$ is produced by KeyGen^A, then $\mathsf{Ver}^A\left(k_{\text{public}}, \mathsf{Bank}^A(k_{\text{private}})\right)$ accepts with certainty.*

Soundness: *Suppose a would-be counterfeiter with access to A and k_{public} is given q valid banknotes. If this counterfeiter outputs any number of (possibly-entangled) quantum states, there is at most an exponentially-small probability that Ver^A will accept more than q of them.*

By adapting these ideas to the private-key setting, we are also able to provide the first *private*-key quantum money scheme that is unconditionally secure, even if the counterfeiter is able to interact adaptively with the bank. This patches a security hole in Wiesner's original scheme which was observed in [28, 3], but which has not previously been addressed in a provably-secure way.

Finally, we provide a candidate cryptographic protocol for obfuscating the indicator functions of subspaces $A \leq \mathbb{F}_2^n$. In order to obfuscate a membership oracle for A, we provide a random system of polynomials $p_1, \ldots, p_m$ that vanish on A. Membership in A can be tested by evaluating the p_i's, but given only the p_i's, we conjecture that it is difficult to recover A. Combining this protocol with the black-box money scheme, we obtain an *explicit* quantum money scheme. This scheme is also the first public-key quantum money scheme whose security can be based on a plausible "classical" cryptographic assumption. Here is the assumption:

Conjecture (*)**.** *Suppose A is a uniformly-random $n/2$-dimensional subspace of $\mathbb{F}_2^n$, and that $\{p_i\}_{1\leq i\leq 2n}, \{q_i\}_{1\leq i\leq 2n}$ are systems of degree-d polynomials from $\mathbb{F}_2^n$ to $\mathbb{F}_2$, which vanish on A and $A^\perp$ respectively but are otherwise uniformly-random. Then for large enough constant d, there is no polynomial-time quantum algorithm that takes as input descriptions of the p_i's and q_i's, and that outputs a basis for A with success probability $\Omega\left(2^{-n/2}\right)$.*

Note that we can trivially guess a *single* nonzero A element with success probability $2^{-n/2}$, but guessing a whole *basis* for A would succeed with probability only $2^{-\Omega(n^2)}$. Conjecture (*) asserts that it is harder to find many elements of A than to find just one element.

The following theorem says that, *if* a counterfeiter could break our quantum money scheme, then with nontrivial success probability, it could *also* recover a description of A from the p_i's and q_i's alone—even without having access to a bank that provides a valid money state $|A\rangle$.

Theorem. *Assuming Conjecture (*), there exists a public-key quantum money scheme with perfect completeness and*

[8] Although the "quadratic speedup" part is not strictly necessary for us, it improves our lower bound on the number of queries the counterfeiter needs to make—to the tight one, in fact—and might be of independent interest.

exponentially-small soundness error. That is, the verifier always accepts valid banknotes, and a would-be counterfeiter succeeds only with exponentially-small probability.[9]

The problem of recovering a subspace A, given a system of equations that vanish on A, is closely related to *algebraic cryptanalysis*, and in particular to the so-called *polynomial isomorphism problem*. In the latter problem, we are given as input two polynomials $p, q : \mathbb{F}^n \rightarrow \mathbb{F}$ related by an unknown linear change of basis L; the challenge is to find L. When $\deg(p) = \deg(q) = 3$, the best known algorithms for the polynomial isomorphism problem require exponential time [35, 24, 17]. An attacker *might* be able to use known techniques to effectively reduce the degree of the polynomials in our scheme by 1, at the expense of an exponentially reduced success probability [17]. Provided the degree is at least 4, however, recovering A seems to be well beyond existing techniques.

1.4 Motivation

Unlike the closely-related task of *quantum key distribution* [13] (which is already practical), quantum money currently seems to be a long way off. The basic difficulty is how to maintain the *coherence* of a quantum money state for an appreciable length of time. *All* money eventually loses its value unless it is spent, but money that decohered on a scale of microseconds would be an extreme example!

So one might wonder: why develop rigorous foundations for a cryptographic functionality that seems so far from being practical? One answer is that, just as quantum key distribution uses many of the same ideas as private-key quantum money, but without requiring long-lasting coherence, so it is not hard to imagine protocols that would use many of the same ideas as *public*-key quantum money without requiring long-lasting coherence. Indeed, depending on the problem, rapid decoherence might be a *feature* rather than a bug!

As one example, public-key quantum money that decohered quickly could be used to create **non-interactive uncloneable signatures**. These are n-qubit quantum states $|\psi\rangle$ that an agent can efficiently prepare using a private key, then freely hand out to passersby. By feeding $|\psi\rangle$, together with the agent's *public* key, into suitable measuring equipment, anyone can verify on the spot that the agent is who she says she is and not an impostor. Compared with *classical* identification protocols, the novel feature here is that the agent does not need to respond to a *challenge*—for example, digitally signing a random string—but can instead just hand out a fixed $|\psi\rangle$ non-interactively. Furthermore, because $|\psi\rangle$ decoheres in a matter of seconds, and recovering a classical *description* of $|\psi\rangle$ from measurements on it is computationally intractable, someone who is given $|\psi\rangle$ cannot use it later to impersonate the agent.

Of course, if an attacker managed to solve the technological problem of keeping $|\psi\rangle$ coherent for very long times, then he could break this system, by collecting one or more copies of $|\psi\rangle$ that an agent had handed out, and using them to impersonate the agent. But in that case, whatever method the attacker was using to keep the states coherent could *also*—once discovered—be used to create a secure public-key quantum money scheme!

However, we believe the "real" reason to study quantum money is basically the same as the "real" reason to study quantum computing as a whole—or for that matter, to study the many interesting aspects of *classical* cryptography that are equally far from application. As theoretical computer scientists, we are in the business of mapping out the inherent capabilities and limits of information processing.

In our case, what quantum money provides is a near-ideal playground for understanding the implications of the uncertainty principle and the No-Cloning Theorem. In the early days of quantum mechanics, Bohr [15] and others argued that the uncertainty principle requires us to change our conception of science itself—their basic argument being that, in physics, predictions are only ever as good as our knowledge of a system's initial state $|\psi\rangle$, but the uncertainty principle might mean that the initial state is unknowable even with arbitrarily-precise measurements.

But does this argument have any "teeth"? In other words: among the properties of a quantum state $|\psi\rangle$ that make the state impossible to learn precisely or to duplicate, can any of those properties ever *matter empirically*? To us, quantum money is interesting precisely because it gives one of the clearest examples where the answer to that question is yes.

2. PRELIMINARIES

To begin, we fix some notation. Let $[N] = \{1, \ldots, N\}$. Given a subspace S of a vector space V, let $S^{\perp}$ be the orthogonal complement of S (that is, the set of $y \in V$ such that $x \cdot y = 0$ for all $x \in S$). We call a function $\delta(n)$ *negligible* if $\delta(n) = o(1/p(n))$ for every polynomial p.

By a *classical oracle*, we will mean a unitary transformation of the form $|x\rangle \rightarrow (-1)^{f(x)}|x\rangle$, for some Boolean function $f : \{0,1\}^* \rightarrow \{0,1\}$. Note that, unless specified otherwise, even a classical oracle can be queried in quantum superposition. A *quantum oracle*, by contrast, is an arbitrary n-qubit unitary transformation U (or rather, a collection of such U's, one for each n) that a quantum algorithm can apply in a black-box fashion. Quantum oracles were defined and studied by Aaronson and Kuperberg [5].

2.1 Cryptography

Before we construct quantum money schemes, it will be helpful to have some "conventional" cryptographic primitives in our toolbox. Foremost among these is a *digital signature scheme secure against quantum chosen-message attacks*. We now define digital signature schemes—both for completeness, and to fix the quantum attack model that is relevant for us.

Definition 1 (Digital Signature Schemes)**.** *A (classical, public-key)* ***digital signature scheme*** $\mathcal{D}$ *consists of three probabilistic polynomial-time classical algorithms:*

- KeyGen*, which takes as input a security parameter* 0^n*, and generates a* ***key pair*** $(k_{\text{private}}, k_{\text{public}})$.
- Sign*, which takes as input* k_{private} *and a message* x*, and generates a* ***signature*** $\mathsf{Sign}(k_{\text{private}}, x)$.[10]

[9] This theorem remains true even if the statement of Conjecture (*) is weakened by adding random noise to the p_i's and q_i's, so that only a constant fraction of them vanish on A or $A^{\perp}$. The presence of noise interferes substantially with known techniques for solving systems of equations, though an attacker who was able to recover A from a *single* polynomial would of course not be hindered by such noise.

[10] We indulge in slight abuse of notation, since if Sign is ran-

- Ver, *which takes as input* k_{public}*, a message* x*, and a claimed signature* w*, and either accepts or rejects.*

We say $\mathcal{D}$ *has* ***completeness error*** ε *if*

$$\mathsf{Ver}\left(k_{\text{public}}, x, \mathsf{Sign}\left(x, k_{\text{private}}\right)\right)$$

accepts with probability at least $1-\varepsilon$ *for all messages* x *and key pairs* $(k_{\text{private}}, k_{\text{public}})$*. Here the probability is over the behavior of* Ver *and* Sign*.*

Let C *(the* ***counterfeiter****) be a quantum circuit of size* $\text{poly}(n)$ *that takes* k_{public} *as input and does the following:*

(1) Probabilistically generates a classical list of messages $x_1, \ldots, x_m$*, and submits them to a* ***signing oracle*** $\mathcal{O}$*.*

(2) Gets back independently-generated signatures $w_1, \ldots, w_m$*, where* $w_i := \mathsf{Sign}\left(k_{\text{private}}, x_i\right)$*.*

(3) Outputs a pair (x, w)*.*

We say C ***succeeds*** *if* $x \notin \{x_1, \ldots, x_m\}$ *and* $\mathsf{Ver}\left(k_{\text{public}}, x, w\right)$ *accepts. We say* $\mathcal{D}$ *has* ***soundness error*** δ *if every counterfeiter* C *succeeds with probability at most* δ*. Here the probability is over the key pair* $(k_{\text{private}}, k_{\text{public}})$ *and the behavior of* C*,* Sign*, and* Ver*.*

We call $\mathcal{D}$ ***secure against nonadaptive quantum chosen-message attacks*** *if it has completeness error* $\leq 1/3$ *and negligible soundness error.*

Intuitively, we call a signature scheme "secure" if no *quantum* counterfeiter with *nonadaptive, classical* access to a signing oracle $\mathcal{O}$ can forge a signature for any message that it did not submit to $\mathcal{O}$. Depending on the application, one might want to change Definition 1 in various ways: for example, by giving the counterfeiter *adaptive* or *quantum* access to $\mathcal{O}$, or by letting KeyGen, Sign, and Ver be quantum algorithms themselves. For this paper, however, Definition 1 provides all we need.

Do signature schemes secure against quantum attack exist? Naturally, signature schemes based on RSA or other number-theoretic problems can all be broken by a quantum computer. However, building on earlier work by Naor and Yung [33] (among many others), Rompel [37] showed that a secure public-key signature scheme can be constructed from *any* one-way function—not necessarily a trapdoor function. Furthermore, Rompel's security reduction, from breaking the signature scheme to inverting the one-way function, is *black-box*: in particular, nothing in it depends on the assumption that the adversary is classical rather than quantum. We therefore get the following consequence:

Theorem 2 (Quantum-Secure Signature Schemes [37])**.** *If there exists a (classical) one-way function* f *secure against quantum attack, then there also exists a digital signature scheme secure against quantum chosen-message attacks.*

Recently, Boneh et al. [16] proved several results similar to Theorem 2, and they needed nontrivial work to do so. However, a crucial difference is that Boneh et al. were (justifiably) concerned with quantum adversaries who can make *quantum* queries to the signing oracle $\mathcal{O}$. By contrast, as mentioned earlier, for our application it suffices to consider adversaries who query $\mathcal{O}$ *classically*—and in that case, the standard security reductions go through essentially without change.

Let us state another consequence of Theorem 2, which will be useful for our oracle construction in Section 5.

Theorem 3 (Relativized Quantum-Secure Signatures)**.** *Relative to a suitable oracle* A*, there exists a digital signature scheme secure against quantum chosen-message attacks.*

Proof Sketch. It is easy to give an oracle $A : \{0,1\}^* \rightarrow \{0,1\}$ relative to which there exists a one-way function $f_n : \{0,1\}^n \rightarrow \{0,1\}^{p(n)}$ secure against quantum adversaries. Now, the security reduction of Rompel [37] is not only black-box but *relativizing*: that is, it goes through if all legitimate and malicious parties have access to the same oracle A. So by Theorem 2, starting from $\{f_n\}$ one can construct a digital signature scheme relative to the same oracle A, which is secure against quantum chosen-message attacks. Further details are given in the full version. ■

2.2 Quantum Information

Let us collect a few facts about quantum pure and mixed states that are used in the paper. We assume basic familiarity with the formalism of bras, kets, density matrices, etc.; see Nielsen and Chuang [34] for a good overview.

Given two mixed states ρ and σ, their *trace distance* is defined as $D(\rho,\sigma) := \frac{1}{2}\sum_{i=1}^{N} |\lambda_i|$, where $\lambda_1, \ldots, \lambda_N$ are the eigenvalues of $\rho - \sigma$. Trace distance is a metric and satisfies $0 \leq D(\rho,\sigma) \leq 1$. Also, the *fidelity* $0 \leq F(\rho,\sigma) \leq 1$ is defined, in this paper, as the maximum of $|\langle\psi|\varphi\rangle|$ over all purifications $|\psi\rangle$ of ρ and $|\varphi\rangle$ of σ.[11] By extension, given a subspace S, we let $F(\rho, S)$ be the maximum of $|\langle\psi|\varphi\rangle|$ over all purifications $|\psi\rangle$ of ρ and all unit vectors $|\varphi\rangle \in S$. While fidelity is *not* a metric, it does satisfy the following inequality, which will be helpful in Section 5.

Lemma 4 ("Triangle Inequality" for Fidelity)**.** *Suppose*

$$\min\left\{\langle\psi|\,\rho\,|\psi\rangle, \langle\varphi|\,\sigma\,|\varphi\rangle\right\} \geq 1-\varepsilon.$$

Then $F(\rho,\sigma) \leq |\langle\psi|\varphi\rangle| + 2\varepsilon^{1/4}$*.*

Proof. Deferred to the full version. ■

Finally, the following lemma of Aaronson [2] will imply that, as long as a quantum money scheme has small *completeness error* (i.e., small probability of rejecting a valid banknote), the banknotes can be reused many times.

Lemma 5 ("Almost As Good As New Lemma" [2])**.** *Suppose a measurement on a mixed state* ρ *yields a particular outcome with probability* $1-\varepsilon$*. Then after the measurement, one can recover a state* $\widetilde{\rho}$ *such that* $\|\widetilde{\rho} - \rho\|_{\text{tr}} \leq \sqrt{\varepsilon}$*.*

2.3 Quantum Search

In our security proof for quantum money, a crucial step will be to *amplify* a counterfeiter who copies a banknote \$ with any non-negligible fidelity to a counterfeiter who copies \$ almost perfectly. Taking the contrapositive, this will imply that to rule out the former sort of counterfeiter, it suffices to rule out the latter.

In this section, we first review two variants of Grover's search algorithm [25] that are useful for amplifying the fidelity of quantum states. We then introduce a new variant that combines the advantages of both.

domized then the signature need not be a function of k_{private} and x.

[11] Some authors instead define "fidelity" as the maximum of $|\langle\psi|\varphi\rangle|^2$.

Throughout, assume we are given a pure initial state $|\text{Init}\rangle$, in some Hilbert space $\mathcal{H}$. Our goal is to map $|\text{Init}\rangle$ to a final state $|\Psi\rangle$ that lies in (or close to) a "good subspace" $G \leq \mathcal{H}$. We have oracle access to two unitary transformations:

- U_{Init}, which maps $|\text{Init}\rangle$ to $-|\text{Init}\rangle$, and acts as the identity on all $|v\rangle$ orthogonal to $|\text{Init}\rangle$.
- U_G, which maps $|v\rangle$ to $-|v\rangle$ for all $|v\rangle \in G$, and acts as the identity on all $|v\rangle$ orthogonal to G.

We are promised that $F(|\text{Init}\rangle, G) = \max_{|\psi\rangle \in G} \langle \text{Init}|\psi\rangle$, the fidelity of the initial state with G, is at least some $\varepsilon > 0$. In this scenario, *provided* $F(|\text{Init}\rangle, G)$ *is known*, the amplitude amplification framework of Brassard, Høyer, Mosca, and Tapp [18] lets us prepare a state close to G using $\Theta(1/\varepsilon)$ iterations:

Lemma 6 (Amplitude Amplification [18]). *Write* $|\text{Init}\rangle$ *as* $\sin\theta\,|\text{Good}\rangle + \cos\theta\,|\text{Bad}\rangle$, *where* $|\text{Good}\rangle$ *is the unit vector in* G *closest to* $|\text{Init}\rangle$ *and* $|\text{Bad}\rangle$ *is orthogonal to* $|\text{Good}\rangle$. *Then by using* $O(T)$ *oracle calls to* U_{Init} *and* U_G, *we can prepare the state*

$$|\Phi_T\rangle := \sin[(2T+1)\theta]\,|\text{Good}\rangle + \cos[(2T+1)\theta]\,|\text{Bad}\rangle$$

Note that Grover's algorithm is simply a special case of Lemma 6, where $|\text{Init}\rangle$ is the uniform superposition over N basis states $|1\rangle, \ldots, |N\rangle$, and G is the subspace spanned by "marked" states.

However, Lemma 6 has an annoying drawback, which it shares with ordinary Grover search. Namely, the algorithm does *not* converge monotonically toward the target subspace G, but could instead "wildly overshoot it," cycling around the 2-dimensional subspace spanned by $|\text{Bad}\rangle$ and $|\text{Good}\rangle$. If we know the fidelity $F(|\text{Init}\rangle, G)$ in advance (rather than just a lower bound on the fidelity), *or* if we can prepare new copies of $|\text{Init}\rangle$ "free of charge" in case of failure, then this overshooting is not a serious problem. Alas, neither of those conditions will hold in our application.

Fortunately, for independent reasons, in 2005 Tulsi, Grover, and Patel [38] introduced a new quantum search algorithm that *does* guarantee monotonic convergence toward G, by alternating unitary transformations with measurements. (Their algorithm was later simplified and improved by Chakraborty, Radhakrishnan, and Raghunathan [19].)

Lemma 7 (Fixed-Point Quantum Search [38, 19]). *By using* T *oracle calls to* U_{Init} *and* U_G, *we can prepare a state* $|\Psi\rangle$ *such that* $F(|\Psi\rangle, G) \geq 1 - \exp(-T\varepsilon^2)$.

Rearranging, Lemma 7 lets us prepare a state $|\Psi\rangle$ such that $F(|\Psi\rangle, G) \geq 1-\delta$ using $T = O\left(\frac{1}{\varepsilon^2}\ln\frac{1}{\delta}\right)$ iterations. On the positive side, the dependence on $1/\delta$ in this bound is logarithmic: we get not only monotonic convergence toward G, but *exponentially-fast* convergence. On the negative side, notice that the dependence on ε has worsened from $1/\varepsilon$ to $1/\varepsilon^2$—negating the quadratic speedup that was the original point of quantum search!

Thus, in the full version, we give a "hybrid" quantum search algorithm that *combines* the advantages of Lemmas 6 and 7—i.e., it converges monotonically toward the target subspace G (rather than "overshooting" G), but *also* achieves a quadratic speedup. In the context of our security proof for quantum money, the hybrid algorithm leads to a quadratically-better lower bound on the number of queries that a counterfeiter needs to make, compared to what we would get from using Lemma 7 by itself. While this quadratic improvement is perhaps only of moderate interest, we hope it might find other applications.

Theorem 8 (Faster Fixed-Point Search). *Let* $\delta \geq 2\varepsilon$. *Then by using* $O\left(\frac{\log 1/\delta}{\varepsilon\delta^2}\right)$ *oracle calls to* U_{Init} *and* U_G, *we can prepare a state* ρ *such that* $F(\rho, G) \geq 1-\delta$.

Proof. Deferred to the full version. ■

Note that Theorem 8 loses the *exponentially-fast* convergence toward the target subspace G, but that property will not be important for us anyway. We leave as an open problem whether there exists a hybrid algorithm with exponentially-fast convergence.

3. FORMALIZING QUANTUM MONEY

In this section, we first give a formal cryptographic definition of *public-key quantum money schemes*. Our definition is similar to that of Aaronson [3]. However, departing from [3], we next define the new notion of a *quantum money mini-scheme*, which is easier to construct and analyze than a full-blown quantum money scheme. A mini-scheme is basically a quantum money scheme where only *one* banknote ever needs to printed, not many banknotes; and where the procedure for verifying that banknote is treated as given (rather than something that *itself* needs to be authenticated using the bank's public key). We then prove two basic results: the amplification of weak counterfeiters into strong ones (Theorem 13), and the construction of full-blown quantum money schemes from mini-schemes together with quantumly-secure digital signature schemes (Theorem 14).

3.1 Quantum Money Schemes

Intuitively, a *public-key quantum money scheme* is a scheme by which

(1) a trusted "bank" can feasibly generate an unlimited number of quantum banknotes,

(2) anyone can feasibly verify a valid banknote as having come from the bank, but

(3) no one besides the bank can feasibly map $q = \text{poly}(n)$ banknotes to $r > q$ banknotes with any non-negligible success probability.[12]

We now make the notion more formal.

Definition 9 (Quantum Money Schemes). *A* ***public-key quantum money scheme*** $\mathcal{S}$ *consists of three polynomial-time quantum algorithms:*

- KeyGen, *which takes as input a security parameter* 0^n, *and probabilistically generates a key pair* $(k_{\text{private}}, k_{\text{public}})$.
- Bank, *which takes as input* k_{private}, *and probabilistically generates a quantum state* \$ *called a* ***banknote***. *(Usually* \$ *will be an ordered pair* (s, ρ_s), *consisting of a classical* ***serial number*** s *and a* ***quantum money state*** ρ_s, *but this is not strictly necessary.)*

[12] Previously, Aaronson [3] required only that no polynomial-time counterfeiter could increase its *expected* number of valid banknotes. However, the stronger condition required here is both achievable, and seemingly more natural from the standpoint of security proofs.

- Ver, *which takes as input* k_{public} *and an alleged banknote* $\$$*, and either accepts or rejects.*

We say $\mathcal{S}$ *has* ***completeness error*** ε *if* $\mathsf{Ver}(k_{\text{public}}, \$)$ *accepts with probability at least* $1-\varepsilon$ *for all public keys* k_{public} *and valid banknotes* $\$$*. If* $\varepsilon = 0$ *then* $\mathcal{S}$ *has* ***perfect completeness****.*

Let Count *(the* ***money counter****) take as input* k_{public} *as well as a collection of (possibly-entangled) alleged banknotes* $\$_1, \ldots, \$_r$*, and output the number of indices* $i \in [r]$ *such that* $\mathsf{Ver}(k_{\text{public}}, \$_i)$ *accepts. Then we say* $\mathcal{S}$ *has* ***soundness error*** δ *if, given any quantum circuit* $C(k_{\text{public}}, \$_1, \ldots, \$_q)$ *of size* $\text{poly}(n)$ *(called the* ***counterfeiter****), which maps* $q = \text{poly}(n)$ *valid banknotes* $\$_1, \ldots, \$_q$ *to* $r = \text{poly}(n)$ *(possibly-entangled) alleged banknotes* $\$_1, \ldots, \$_r$*,*

$$\Pr\left[\mathsf{Count}(k_{\text{public}}, C(k_{\text{public}}, \$_1, \ldots, \$_q)) > q\right] \leq \delta.$$

Here the probability is over the key pair $(k_{\text{private}}, k_{\text{public}})$*, valid banknotes* $\$_1, \ldots, \$_q$ *generated by* $\mathsf{Bank}(k_{\text{private}})$*, and the behavior of* Count *and* C*.*

We call $\mathcal{S}$ ***secure*** *if it has completeness error* $\leq 1/3$ *and negligible soundness error.*

In the full version, we show that the completeness error in any quantum money scheme can be amplified to $1/2^{\text{poly}(n)}$, at the cost of only a small increase in the soundness error. Note that, by Lemma 5 (the "Almost As Good As New Lemma"), once we make the completeness error exponentially small, we can also give our scheme the property that *any banknote* $\$$ *can be verified exponentially many times*, before $\$$ gets "worn out" by repeated measurements. This observation is part of what justifies our use of the term "money."[13]

In this paper, we will often consider **relativized** quantum money schemes, which simply means that the three procedures KeyGen, Bank, Ver—as well as the counterfeiter C—all get access to exactly the same oracle $A : \{0,1\}^* \rightarrow \{0,1\}$. We will also consider relativized digital signature schemes, etc., which are defined analogously.

A **private-key quantum money scheme** is the same as a public-key scheme, except that the counterfeiter C no longer gets access to k_{public}. (Thus, we might as well set $k := k_{\text{public}} = k_{\text{private}}$, since the public and private keys no longer play separate roles.) We call a private-key scheme **query-secure**—a notion "intermediate" between private-key and public-key—if the counterfeiter C is allowed to interact repeatedly with the bank. Given any alleged banknote σ, the bank runs the verification procedure $\mathsf{Ver}(k, \sigma)$, then returns to C both the classical result (i.e., accept or reject) *and* the post-measurement quantum state $\widetilde{\sigma}$.

3.2 Mini-Schemes

While Definition 9 captures our intuitive requirements for a public-key quantum money scheme, experience has shown that it is cumbersome to work with in practice. So in this section, we introduce a simpler primitive called *mini-schemes*, which require only *one* uncopyable banknote. We also prove an amplification theorem for a large class of mini-schemes. Then, in Section 3.3, we will show how mini-schemes can be generically combined with conventional digital signature schemes to create full public-key quantum money schemes.

Definition 10 (Mini-Schemes)**.** *A (public-key)* ***mini-scheme*** $\mathcal{M}$ *consists of two polynomial-time quantum algorithms:*

- Bank*, which takes as input a security parameter* 0^n*, and probabilistically generates a banknote* $\$ = (s, \rho_s)$*, where* s *is a classical* ***serial number****, and* ρ_s *is a quantum money state.*
- Ver*, which takes as input an alleged banknote* $\$$*, and either accepts or rejects.*

We say $\mathcal{M}$ *has* ***completeness error*** ε *if* $\mathsf{Ver}(\$)$ *accepts with probability at least* $1-\varepsilon$ *for all valid banknotes* $\$$*. If* $\varepsilon = 0$ *then* $\mathcal{M}$ *has perfect completeness. If, furthermore,* $\rho_s = |\psi_s\rangle\langle\psi_s|$ *is always a pure state, and* Ver *simply consists of a projective measurement onto the rank-1 subspace spanned by* $|\psi_s\rangle$*, then we say* $\mathcal{M}$ *is* ***projective****.*[14]

Let Ver_2 *(the* ***double verifier****) take as input a single serial number* s *as well as two (possibly-entangled) states* σ_1 *and* σ_2*, and accept if and only* $\mathsf{Ver}(s, \sigma_1)$ *and* $\mathsf{Ver}(s, \sigma_2)$ *both accept. We say* $\mathcal{M}$ *has* ***soundness error*** δ *if, given any quantum circuit* C *of size* $\text{poly}(n)$ *(the* ***counterfeiter****),* $\mathsf{Ver}_2(s, C(\$))$ *accepts with probability at most* δ*. Here the probability is over the banknote* $\$$ *output by* $\mathsf{Bank}(0^n)$*, as well as the behavior of* Ver_2 *and* C*.*

We call $\mathcal{M}$ ***secure*** *if it has completeness error* $\leq 1/3$ *and negligible soundness error.*

We observe a simple relationship between Definitions 9 and 10:

Proposition 11. *If there exists a secure public-key money scheme* $\mathcal{S} = (\mathsf{KeyGen}_{\mathcal{S}}, \mathsf{Bank}_{\mathcal{S}}, \mathsf{Ver}_{\mathcal{S}})$*, then there also exists a secure mini-scheme* $\mathcal{M} = (\mathsf{Bank}_{\mathcal{M}}, \mathsf{Ver}_{\mathcal{M}})$*.*

PROOF. Each banknote output by $\mathsf{Bank}_{\mathcal{M}}(0^n)$ will have the form $(k_{\text{public}}, \mathsf{Bank}_{\mathcal{S}}(k_{\text{private}}))$, where $(k_{\text{private}}, k_{\text{public}})$ is a key pair output by $\mathsf{KeyGen}_{\mathcal{S}}(0^n)$. Then $\mathsf{Ver}_{\mathcal{M}}(s, \rho_s)$ will accept if and only if $\mathsf{Ver}_{\mathcal{S}}(s, \rho_s)$ does. Any counterfeiter $C_{\mathcal{M}}$ against $\mathcal{M}$ can be converted directly into a counterfeiter $C_{\mathcal{S}}$ against $\mathcal{S}$. ■

Call a mini-scheme $\mathcal{M} = (\mathsf{Bank}, \mathsf{Ver})$ **secret-based** if Bank works by first generating a uniformly-random classical string r, and then generating a banknote $\$_r := (s_r, \rho_r)$. Intuitively, in a secret-based scheme, the bank can generate many identical banknotes by simply reusing r, while in a non-secret-based scheme, *not even the bank* might be able to generate two identical banknotes. Here is an interesting observation:

Proposition 12. *If there exists a secure, secret-based mini-scheme, then there also exists a one-way function secure against quantum attack.*

PROOF. The desired OWF is $\mathsf{SerialNum}(r) := s_r$. If there existed a polynomial-time quantum algorithm to recover r given s_r, then we could use that algorithm to produce an unlimited number of additional banknotes $\$_r$. ■

All of the mini-schemes developed in this paper will be secret-based. By contrast, the earlier schemes of Lutomirski

[13] By contrast, BBBW [14] introduced the term "subway tokens" for quantum money states that get destroyed immediately upon verification.

[14] We similarly call a full quantum money scheme projective, if $\mathsf{Ver}(\$)$ consists of a measurement on one part of $\$$ in the computational basis, followed by a rank-1 projective measurement on the remaining part.

et al. [30] and Farhi et al. [22] are non-secret-based, since the serial number s is only obtained as the outcome of a quantum measurement.

The following result is one of the most useful in the paper. Intuitively, it says that in projective mini-schemes, a counterfeiter that copies a banknote with *any* non-negligible fidelity can be "amplified" to a counterfeiter that copies the banknote almost *perfectly*—or conversely, that to rule out the former sort of counterfeiter, it suffices to rule out the latter. The proof makes essential use of the amplitude amplification results from Section 2.3.

Theorem 13 (Amplification of Counterfeiters)**.** *Let $\mathcal{M} = (\mathsf{Bank}, \mathsf{Ver})$ be a projective mini-scheme, and let $\$ = (s, \rho)$ be a valid banknote in $\mathcal{M}$. Suppose there exists a counterfeiter C that copies $\$$ with probability $\varepsilon > 0$: that is,*

$$\Pr\left[\mathsf{Ver}_2\left(s, C\left(\$\right)\right) \ \textit{accepts}\right] \geq \varepsilon.$$

Then for all $\delta > 0$, there is also a modified counterfeiter C' (depending only on ε and δ, not $\$$), which makes

$$O\left(\frac{\log 1/\delta}{\sqrt{\varepsilon}\left(\sqrt{\varepsilon}+\delta^2\right)}\right)$$

queries to C, C^{-1}, and Ver and which satisfies

$$\Pr\left[\mathsf{Ver}_2\left(s, C'\left(\$\right)\right) \ \textit{accepts}\right] \geq 1-\delta.$$

Proof. Write $\$$ as a mixture of pure states:

$$\$ = \sum p_i \left|\psi_i\right\rangle \left\langle\psi_i\right|.$$

By linearity, clearly it suffices to show that

$$\Pr\left[\mathsf{Ver}_2\left(s, C'\left(\left|\psi_i\right\rangle\right)\right) \ \text{accepts}\right] \geq 1-\delta$$

for all i such that $p_i > 0$. We focus on $\left|\psi\right\rangle := \left|\psi_1\right\rangle$ without loss of generality.

By assumption, there exists a subspace S such that

$$\Pr\left[\mathsf{Ver}\left(\rho\right) \ \text{accepts}\right] = F\left(\rho, S\right)$$

for all ρ. Then $F\left(\$, S\right) = F\left(\left|\psi\right\rangle, S\right) = 1$.

Now, just as Ver is simply a projector onto S, so Ver_2 is a projector onto $S^{\otimes 2}$. Thus

$$F\left(C\left(\left|\psi\right\rangle\right), S^{\otimes 2}\right) \geq \sqrt{\varepsilon}.$$

So consider performing a fixed-point Grover search, with $C\left(\left|\psi\right\rangle\right)$ as the initial state and $S^{\otimes 2}$ as the target subspace. By Lemma 7, this will produce a state ρ such that $F\left(\rho, S^{\otimes 2}\right) \geq 1-\delta$ using $O\left(\frac{1}{\varepsilon}\log\frac{1}{\delta}\right)$ Grover iterations. Each iteration requires a reflection about $C\left(\left|\psi\right\rangle\right)$ and a reflection about $S^{\otimes 2}$, which can be implemented using $O\left(1\right)$ queries to C, C^{-1} and Ver respectively. Therefore the number of queries to C, C^{-1} and Ver is $O\left(\frac{1}{\varepsilon}\log\frac{1}{\delta}\right)$ as well.

If δ is large compared to ε, then we can instead use Theorem 8, which produces a state ρ such that $F\left(\rho, S^{\otimes 2}\right) \geq 1-\delta$ using $O\left(\frac{1}{\sqrt{\varepsilon}\delta^2}\log\frac{1}{\delta}\right)$ iterations. Taking the minimum of the two bounds gives us the claimed bound on query complexity. ■

Theorem 13 is unlikely to hold for *arbitrary* (non-projective) mini-schemes, for the simple reason that we can always create a mini-scheme where Ver accepts *any* state with some small nonzero probability ε. We leave it as an open problem to find the largest class of mini-schemes for which Theorem 13 holds.

3.3 The Standard Construction

We are now ready to define the "standard construction" of public-key quantum money schemes from mini-schemes and digital signature schemes, and to prove this construction's security.

Theorem 14 (Standard Construction of Public-Key Quantum Money)**.** *Let $\mathcal{M}$ be any secure mini-scheme, and let $\mathcal{D}$ be any digital signature scheme secure against quantum chosen-message attacks. By combining $\mathcal{M}$ and $\mathcal{D}$, we can create a secure public-key quantum money scheme $\mathcal{S}$.*

Proof. Given

$$\mathcal{M} = \left(\mathsf{Bank}_{\mathcal{M}}, \mathsf{Ver}_{\mathcal{M}}\right),$$
$$\mathcal{D} = \left(\mathsf{KeyGen}_{\mathcal{D}}, \mathsf{Sign}_{\mathcal{D}}, \mathsf{Ver}_{\mathcal{D}}\right),$$

our quantum money scheme $\mathcal{S} = \left(\mathsf{KeyGen}_{\mathcal{S}}, \mathsf{Bank}_{\mathcal{S}}, \mathsf{Ver}_{\mathcal{S}}\right)$ is defined as follows:

$\mathsf{KeyGen}_{\mathcal{S}}$ is simply $\mathsf{KeyGen}_{\mathcal{D}}$ from the digital signature scheme.

$\mathsf{Bank}_{\mathcal{S}}$ first calls $\mathsf{Bank}_{\mathcal{M}}$ from the mini-scheme to obtain a banknote (s, ρ). It then outputs (s, ρ) together with a digital signature of the serial number s:

$$\mathsf{Bank}_{\mathcal{S}}\left(k_{\text{private}}\right) := \left(s, \mathsf{Sign}_{\mathcal{D}}\left(k_{\text{private}}, s\right), \rho\right).$$

$\mathsf{Ver}_{\mathcal{S}}$ accepts an alleged banknote (s, w, σ), if and only if $\mathsf{Ver}_{\mathcal{M}}\left(s, \sigma\right)$ and $\mathsf{Ver}_{\mathcal{D}}\left(k_{\text{public}}, s, w\right)$ both accept.

Now, suppose there exists a counterfeiter $C_{\mathcal{S}}$ against $\mathcal{S}$: that is, a polynomial-time quantum algorithm such that

$$\Pr\left[\mathsf{Count}\left(k_{\text{public}}, C_{\mathcal{S}}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right)\right) > q\right] \geq \frac{1}{p\left(n\right)}.$$

Here $\$_i := \left(s_i, w_i, \rho_i\right)$ is a valid banknote, Count is the money counter from Definition 9, and p is some polynomial. Also, the probability is over the key pair $\left(k_{\text{private}}, k_{\text{public}}\right)$, the valid banknotes $\$_1, \ldots, \$_q$, and the behavior of Count and $C_{\mathcal{S}}$. Suppose further that $\mathcal{D}$ is secure. Then it suffices to show that, by using $C_{\mathcal{S}}$, we can construct a counterfeiter $C_{\mathcal{M}}$ against the underlying mini-scheme $\mathcal{M}$. Let $\mathsf{New}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right)$ be an algorithm that does the following:

(1) Records the serial numbers $s_1, \ldots, s_q$ of $\$_1, \ldots, \$_q$, and lets $U := \left\{s_1, \ldots, s_q\right\}$.

(2) Runs $C_{\mathcal{S}}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right)$, and examines the output states $\$'_1, \ldots, \$'_r$.

(3) Returns the number of $i \in [r]$ such that $\mathsf{Ver}_{\mathcal{S}}\left(\$'_i\right)$ accepts, *and* $\$'_i$'s serial number s'_i does not belong to U.

Then we claim that $\Pr\left[\mathsf{New}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right) > 0\right]$ is negligibly small, where the probability is over the same variables as before. The proof is simply that, if this were not so, then we could easily create a counterfeiter $C_{\mathcal{D}}$ against the digital signature scheme $\mathcal{D}$. With non-negligible probability, $C_{\mathcal{D}}$ would generate a valid signature $\mathsf{Sign}_{\mathcal{D}}\left(k_{\text{private}}, s'_i\right)$, for a message s'_i for which $C_{\mathcal{D}}$ had never before seen a valid signature, by running $C_{\mathcal{S}}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right)$, then measuring $\$'_i = \left(s'_i, w'_i, \rho'_i\right)$ for a uniformly random $i \in [r]$. (Note that $C_{\mathcal{D}}$ can generate q money states $\$_1, \ldots, \$_q$, without knowledge of k_{private}, by generating the s_i's and ρ_i's on its own, then calling the signing oracle $\mathcal{O}$ to get the w_i's.)

But now we can define a counterfeiter $C_{\mathcal{M}}$ against the mini-scheme $\mathcal{M}$, which works as follows. First, $C_{\mathcal{M}}$ runs

$\mathsf{KeyGen}_{\mathcal{D}}(0^n)$ to generate a *new* key pair $\left(k'_{\text{private}}, k'_{\text{public}}\right)$. Next, it labels the banknote to copied (s_ℓ, ρ_ℓ), for some $\ell \in [q]$ chosen uniformly at random. It repeatedly calls $\mathsf{Bank}_{\mathcal{M}}(0^n)$ to generate $q-1$ serial numbers and quantum money states, labeled (s_i, ρ_i) for all $i \in [q] \setminus \{\ell\}$. Let $U := \{s_1, \ldots, s_q\}$. Then $C_{\mathcal{M}}$ generates a digital signature $w_i := \mathsf{Sign}_{\mathcal{D}}\left(k'_{\text{private}}, s_i\right)$ for each $i \in [q]$. Let $\$_i := (s_i, w_i, \rho_i)$. Next, $C_{\mathcal{M}}$ runs the counterfeiter $C_{\mathcal{S}}\left(k_{\text{public}}, \$_1, \ldots, \$_q\right)$, to obtain $r > q$ alleged banknotes $\not\$_1, \ldots, \not\$_r$ where $\not\$_j = \left(s'_j, w'_j, \rho'_j\right)$. Finally, $C_{\mathcal{M}}$ chooses $j, k \in [r]$ uniformly at random without replacement, and output $\left(\rho'_j, \rho'_k\right)$ as a candidate for two copies of ρ_ℓ.

Suppose that $\mathsf{Count} > q$, as happens with probability at least $\frac{1}{p(n)}$. Also suppose that $\mathsf{New} = 0$, as happens all but a negligible fraction of the time. Then by the pigeonhole principle, there must exist indices $j \neq k$ such that $s'_j = s'_k$. With probability at least $1/\binom{r}{2}$, the counterfeiter $C_{\mathcal{M}}$ will find such a (j,k) pair. Therefore it succeeds with overall probability $\Omega(1/\operatorname{poly}(n))$. ■

Theorem 14 reduces the construction of a public-key quantum money scheme to two "smaller" problems: constructing a mini-scheme, and constructing a signature scheme secure against quantum attacks.

In practice, however, the situation is even better, since in this paper, all of our constructions of mini-schemes will *also* yield signature schemes "free of charge"! The following proposition explains why:

Proposition 15. *If there exists a secure, secret-based mini-scheme $\mathcal{M}$, then there also exists a secure public-key quantum money scheme $\mathcal{S}$.*

PROOF. Starting from $\mathcal{M}$, we can get a one-way function secure against quantum attack from Proposition 12, and hence a digital signature scheme $\mathcal{D}$ secure against quantum chosen-message attack from Theorem 2. Combining $\mathcal{M}$ and $\mathcal{D}$ now yields $\mathcal{S}$ by Theorem 14. ■

Finally, let us make explicit what Theorem 14 means for oracle construction.

Corollary 16. *Suppose there exists a mini-scheme $\mathcal{M}$ that is provably secure relative to some oracle $A_{\mathcal{M}}$ (i.e., any counterfeiter $C_{\mathcal{M}}$ against $\mathcal{M}$ must make superpolynomially many queries to $A_{\mathcal{M}}$). Then there exists a public-key quantum money scheme $\mathcal{S}$ that is provably secure relative to some other oracle $A_{\mathcal{S}}$.*

PROOF. By Theorem 3, relative to a suitable oracle $A_{\mathcal{D}}$ (in fact, a *random* oracle suffices), there exists a signature scheme $\mathcal{D}$, such that any quantum chosen-message attack against $\mathcal{D}$ must make superpolynomially many queries to $A_{\mathcal{D}}$.

The oracle $A_{\mathcal{S}}$ will simply be a concatenation of $A_{\mathcal{M}}$ with $A_{\mathcal{D}}$. Relative to $A_{\mathcal{S}}$, we claim that the mini-scheme $\mathcal{M}$ and signature scheme $\mathcal{D}$ are *both* secure—and therefore, by Theorem 14, we can construct a secure public-key quantum money scheme $\mathcal{S}$.

The only worry is that a counterfeiter $C_{\mathcal{M}}$ against $\mathcal{M}$ might gain some advantage by querying $A_{\mathcal{D}}$; or conversely, a counterfeiter $C_{\mathcal{D}}$ against $\mathcal{D}$ might gain some advantage by querying $A_{\mathcal{M}}$. However, this worry is illusory, for the simple reason that the oracles $A_{\mathcal{D}}$ and $A_{\mathcal{M}}$ are generated independently. Thus, if $C_{\mathcal{M}}$ can break $\mathcal{M}$ by querying $A_{\mathcal{D}}$, then it can *also* break $\mathcal{M}$ by querying a randomly-generated "mock-up" $A'_{\mathcal{D}}$ of $A_{\mathcal{D}}$; and conversely, if $C_{\mathcal{D}}$ can break $\mathcal{D}$ by querying $A_{\mathcal{M}}$, then it can also break $\mathcal{D}$ by querying a randomly-generated mock-up $A'_{\mathcal{M}}$ of $A_{\mathcal{M}}$. Regardless of the *computational* cost of generating these mock-ups, they give us a break against $\mathcal{D}$ or $\mathcal{M}$ that makes only $\operatorname{poly}(n)$ oracle queries, thereby giving the desired contradiction. ■

4. INNER-PRODUCT ADVERSARY METHOD

At least in the black-box setting, our goal is to create quantum money (mini-)schemes that we can *prove* are secure—by showing that any counterfeiter would need to make exponentially many queries to some oracle. Proving security results of this kind turns out to require interesting quantum lower bound machinery. In this section, we introduce the *inner-product adversary method*, a new variant of Ambainis's quantum adversary method [8] that is well-adapted to proving the security of quantum money schemes, and that seems likely to find other applications.

Let us explain the difficulty we need to overcome. In a public-key quantum money scheme, a counterfeiter C has *two* powerful resources available:

(1) One or more copies of a "legitimate" quantum money state $|\psi\rangle$.

(2) Access to a *verification procedure* V, which accepts $|\psi\rangle$ and rejects every state orthogonal to $|\psi\rangle$.

Indeed, for us, the situation is even better for C (i.e., worse for us!), since C can query not only the verification procedure V itself, but also an underlying *classical* oracle U that the legitimate buyers and sellers use to implement V. But let us ignore that issue for now.

As a first step, of course, we should understand how to rule out counterfeiting given (1) or (2) separately. If C has a copy of $|\psi\rangle$, but no oracle access to V, then the impossibility of preparing $|\psi\rangle|\psi\rangle$ essentially amounts to the No-Cloning Theorem. Conversely, if C has oracle access to V, but no copy of $|\psi\rangle$, then given unlimited time, C *can* prepare as many copies of $|\psi\rangle$ as it wants, by using Grover's algorithm to search for a quantum state that V accepts. The problem is "merely" that, if $|\psi\rangle$ has n qubits, then Grover's algorithm requires $\Theta\left(2^{n/2}\right)$ iterations, and the BBBV hybrid argument [12] shows that Grover's algorithm is optimal.

What we need, then, is a theorem showing that any counterfeiter needs exponentially many queries to V to prepare $|\psi\rangle|\psi\rangle$, *even if* the counterfeiter has a copy of $|\psi\rangle$ to start with. Such a theorem would contain both the No-Cloning Theorem and the BBBV hybrid argument as special cases. Aaronson [3] called the desired generalization the *Complexity-Theoretic No-Cloning Theorem*, and sketched a proof of it using Ambainis's adversary method. Based on that result, Aaronson also argued that there exists a *quantum oracle* (i.e., a black-box unitary transformation V) relative to which secure public-key quantum money is possible. However, the details were never published.

In this section, we prove a result—Theorem 18—that is much more general than Aaronson's previous Complexity-Theoretic No-Cloning Theorem [3]. Then, in Section 5, we apply Theorem 18 to prove the security of public-key quantum money relative to a *classical* oracle. In the full version, we also apply Theorem 18 to prove the "original" Complexity-Theoretic No-Cloning Theorem [3], which in-

volves Haar-random n-qubit states $|\psi\rangle$, rather than superpositions $|A\rangle$ over subspaces $A \leq \mathbb{F}_2^n$.[15]

4.1 Idea of Method

So, what *is* the inner-product adversary method? In Ambainis's adversary method [8]—like in the BBBV hybrid argument [12] from which it evolved—the basic idea is to upper-bound how much "progress" a quantum algorithm Q can make at distinguishing pairs of oracles, as the result of a single query. Let $\left|\Psi_t^U\right\rangle$ be Q's state after t queries, assuming that the oracle is U. Then *normally*, before any queries have been made, we can assume that $\left|\Psi_0^U\right\rangle = \left|\Psi_0^V\right\rangle$ for all oracles U and V. By contrast, after the final query T, for all oracle pairs (U,V) that Q is trying to distinguish, we must have (say) $\left|\left\langle \Psi_T^U | \Psi_T^V \right\rangle\right| \leq 1/2$. Thus, *if* we can show that the inner product $\left|\left\langle \Psi_t^U | \Psi_t^V \right\rangle\right|$ can decrease by at most ε as the result of a single query, then it follows that Q must make $\Omega\left(1/\varepsilon\right)$ queries.

But when we try to apply the above framework to quantum money, we run into serious difficulties. Most obviously, it is no longer true that $\left|\Psi_0^U\right\rangle = \left|\Psi_0^V\right\rangle$ for all oracles U, V. Indeed, before Q makes even a *single* query to its oracle V, it already has a great deal of information about V, in the form of a legitimate money state $|\psi\rangle$ that V accepts. The task is "merely" to prepare a second copy of a state that Q already has! Worse yet, once we fix two oracles U and V, we find that Q generally *can* exploit the "head start" provided by its initial state to decrease the inner product $\left|\left\langle \Psi_t^U | \Psi_t^V \right\rangle\right|$ by a constant amount, by making just a single query to U or V respectively.

Our solution is as follows. We first carefully choose a distribution $\mathcal{D}$ over oracle pairs (U,V). We then analyze how much the *expected* inner product

$$\mathop{\mathrm{E}}_{(U,V)\sim\mathcal{D}}\left[\left|\left\langle \Psi_t^U | \Psi_t^V \right\rangle\right|\right]$$

can decrease as the result of a single query to U or V. We will find that, even if Q can substantially increase the angle between $\left|\Psi_t^U\right\rangle$ and $\left|\Psi_t^V\right\rangle$ for *some* (U,V) pairs by making a single query, it cannot do so for *most* pairs.

To illustrate, let $|\psi\rangle$ and $|\varphi\rangle$ be two possible quantum money states, which satisfy (say) $\langle\psi|\varphi\rangle = 1/2$. Then if a counterfeiting algorithm succeeds perfectly, it must map $|\psi\rangle$ to $|\psi\rangle^{\otimes 2}$, and $|\varphi\rangle$ to $|\varphi\rangle^{\otimes 2}$. Since

$$\langle\psi|^{\otimes 2}\,|\varphi\rangle^{\otimes 2} = \left(\langle\psi|\varphi\rangle\right)^2 = \frac{1}{4},$$

this means that the counterfeiter must *decrease the corresponding inner product* by at least $1/4$. However, we will show that the *average* inner product can decrease by at most $1/\exp\left(n\right)$ as the result of a single query. From this it will follow that the counterfeiter needs to make $2^{\Omega(n)}$ queries.

Let us mention that today, there are several "sophisticated" versions of the quantum adversary method [9, 27], which *can* yield lower bounds for quantum state generation tasks not unlike the ones we consider. However, a drawback of these methods is that they are extremely hard to apply to concrete problems: doing so typically requires eigenvalue bounds, and often the use of representation theory. For this reason, even if one of the "sophisticated" adversary methods (or a variant thereof) could be applied to the quantum money problem, our approach might still be preferable.

[15] For whatever it is worth, we get a lower bound of $\Omega\left(2^{n/2}\right)$ on the number of queries needed to copy a Haar-random state, which is quadratically better than the $\Omega\left(2^{n/4}\right)$ that we get for subspace states.

4.2 The Method

We now introduce the inner-product adversary method. Let $\mathcal{O}$ be a set of quantum oracles acting on n qubits each. For each $U \in \mathcal{O}$, assume there exists a subspace $S_U \leq \mathbb{C}^{2^n}$ such that

(i) $U|\psi\rangle = -|\psi\rangle$ for all $|\psi\rangle \in S_U$, and

(ii) $U|\eta\rangle = |\eta\rangle$ for all $|\eta\rangle \in S_U^{\perp}$.

Let $R \subset \mathcal{O}\times\mathcal{O}$ be a symmetric binary relation on $\mathcal{O}$, with the properties that

(i) $(U,U) \notin R$ for all $U \in \mathcal{O}$, and

(ii) for every $U \in \mathcal{O}$ there exists a $V \in \mathcal{O}$ such that $(U,V) \in R$.

Suppose that for all $U \in \mathcal{O}$ and all $|\eta\rangle \in S_U^{\perp}$, we have

$$\mathop{\mathrm{E}}_{V\,:\,(U,V)\in R}\left[F\left(|\eta\rangle, S_V\right)^2\right] \leq \varepsilon,$$

where $F\left(|\eta\rangle, S_V\right) = \max_{|\psi\rangle\in S_V}|\langle\eta|\psi\rangle|$ is the fidelity between $|\eta\rangle$ and S_V. Let Q be a quantum oracle algorithm, and let Q^U denote Q run with the oracle $U \in \mathcal{O}$. Suppose Q^U begins in the state $\left|\Psi_0^U\right\rangle$ (possibly already dependent on U). Let $\left|\Psi_t^U\right\rangle$ denote the state of Q^U immediately after the t^{th} query. Also, define a *progress measure* p_t by

$$p_t := \mathop{\mathrm{E}}_{U,V\,:\,(U,V)\in R}\left[\left|\left\langle \Psi_t^U | \Psi_t^V \right\rangle\right|\right].$$

The following lemma bounds how much p_t can decrease as the result of a single query.

Lemma 17 (Bound on Progress Rate)**.**

$$p_t \geq p_{t-1} - 4\sqrt{\varepsilon}.$$

Proof. Let $\left|\Phi_t^U\right\rangle$ denote the state of Q^U immediately *before* the t^{th} query. Then for all t, it is clear that $\left\langle \Phi_t^U | \Phi_t^V \right\rangle = \left\langle \Psi_{t-1}^U | \Psi_{t-1}^V \right\rangle$: in other words, the unitary transformations that Q performs in between query steps have no effect on the inner products. So to prove the lemma, it suffices to show the following inequality:

$$\mathop{\mathrm{E}}_{U,V\,:\,(U,V)\in R}\left[\left|\left\langle \Phi_t^U | \Phi_t^V \right\rangle\right|\right] - \mathop{\mathrm{E}}_{U,V\,:\,(U,V)\in R}\left[\left|\left\langle \Psi_t^U | \Psi_t^V \right\rangle\right|\right] \leq 4\sqrt{\varepsilon}. \quad (*)$$

Let $\{|i\rangle\}_{i\in[B]}$ be an arbitrary orthonormal basis for Q's workspace register. Then we can write

$$\begin{aligned}\left|\Phi_t^U\right\rangle &= \sum_{i\in[B]} \alpha_{t,i}^U\,|i\rangle\left|\Phi_{t,i}^U\right\rangle \\ &= \sum_{i\in[B]} |i\rangle\left(\beta_{t,i}^U\left|\eta_{t,i}^U\right\rangle + \gamma_{t,i}^U\left|\psi_{t,i}^U\right\rangle\right),\end{aligned}$$

where $\left|\eta_{t,i}^U\right\rangle \in S_U^{\perp}$ and $\left|\psi_{t,i}^U\right\rangle \in S_U$. (By normalization, $\left|\beta_{t,i}^U\right|^2 + \left|\gamma_{t,i}^U\right|^2 = \left|\alpha_{t,i}^U\right|^2$.) A query transforms the above state to

$$\left|\Psi_t^U\right\rangle = \sum_{i\in[B]} |i\rangle\left(\beta_{t,i}^U\left|\eta_{t,i}^U\right\rangle - \gamma_{t,i}^U\left|\psi_{t,i}^U\right\rangle\right).$$

So for all $U, V \in \mathcal{O}$,

$$\begin{aligned}
&\left\langle \Phi_t^U | \Phi_t^V \right\rangle - \left\langle \Psi_t^U | \Psi_t^V \right\rangle \\
=&\sum_{i \in [B]} \left(\overline{\beta}_{t,i}^U \left\langle \eta_{t,i}^U \right| + \overline{\gamma}_{t,i}^U \left\langle \psi_{t,i}^U \right| \right) \left(\beta_{t,i}^V \left| \eta_{t,i}^V \right\rangle + \gamma_{t,i}^V \left| \psi_{t,i}^V \right\rangle \right) \\
&- \sum_{i \in [B]} \left(\overline{\beta}_{t,i}^U \left\langle \eta_{t,i}^U \right| - \overline{\gamma}_{t,i}^U \left\langle \psi_{t,i}^U \right| \right) \left(\beta_{t,i}^V \left| \eta_{t,i}^V \right\rangle - \gamma_{t,i}^V \left| \psi_{t,i}^V \right\rangle \right) \\
=&2 \sum_{i \in [B]} \left(\overline{\beta}_{t,i}^U \gamma_{t,i}^V \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle + \overline{\gamma}_{t,i}^U \beta_{t,i}^V \left\langle \psi_{t,i}^U | \eta_{t,i}^V \right\rangle \right).
\end{aligned}$$

By Cauchy-Schwarz, the above implies that

$$\begin{aligned}
\left| \left\langle \Phi_t^U | \Phi_t^V \right\rangle \right| - \left| \left\langle \Psi_t^U | \Psi_t^V \right\rangle \right| \leq & \, 2 \max_{i \in [B]} \left| \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle \right| \\
& + 2 \max_{i \in [B]} \left| \left\langle \psi_{t,i}^U | \eta_{t,i}^V \right\rangle \right|.
\end{aligned}$$

Now fix $U \in \mathcal{O}$ and $i \in [B]$. Then again applying Cauchy-Schwarz,

$$\begin{aligned}
\mathop{\mathrm{E}}_{V \,:\, (U,V) \in R} \left[\left| \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle \right| \right] &\leq \sqrt{\mathop{\mathrm{E}}_{V \,:\, (U,V) \in R} \left[\left| \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle \right|^2 \right]} \\
&\leq \sqrt{\mathop{\mathrm{E}}_{V \,:\, (U,V) \in R} \left[\max_{|\psi\rangle \in S_V} \left| \left\langle \eta_{t,i}^U | \psi \right\rangle \right|^2 \right]} \\
&\leq \sqrt{\varepsilon}.
\end{aligned}$$

Hence

$$\mathop{\mathrm{E}}_{U,V \,:\, (U,V) \in R} \left[\left| \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle \right| \right] \leq \sqrt{\varepsilon}$$

as well, and likewise

$$\mathop{\mathrm{E}}_{U,V \,:\, (U,V) \in R} \left[\left| \left\langle \psi_{t,i}^U | \eta_{t,i}^V \right\rangle \right| \right] \leq \sqrt{\varepsilon}$$

by symmetry. Putting everything together,

$$\begin{aligned}
p_{t-1} - p_t &= \mathop{\mathrm{E}}_{U,V \,:\, (U,V) \in R} \left[\left| \left\langle \Phi_t^U | \Phi_t^V \right\rangle \right| - \left| \left\langle \Psi_t^U | \Psi_t^V \right\rangle \right| \right] \\
&\leq 2 \mathop{\mathrm{E}}_{U,V \,:\, (U,V) \in R} \left[\max_{i \in [B]} \left| \left\langle \eta_{t,i}^U | \psi_{t,i}^V \right\rangle \right| \right] \\
&\quad + 2 \mathop{\mathrm{E}}_{U,V \,:\, (U,V) \in R} \left[\max_{i \in [B]} \left| \left\langle \psi_{t,i}^U | \eta_{t,i}^V \right\rangle \right| \right] \\
&\leq 4\sqrt{\varepsilon}.
\end{aligned}$$

This proves inequality (*) and hence the lemma. ■

From Lemma 17 we immediately deduce the following.

Theorem 18 (Inner-Product Adversary Method)**.** *Suppose that initially $\left| \left\langle \Psi_0^U | \Psi_0^V \right\rangle \right| \geq c$ for all $(U,V) \in R$, whereas by the end we need $\left| \left\langle \Psi_T^U | \Psi_T^V \right\rangle \right| \leq d$ for all $(U,V) \in R$. Then Q must make $T = \Omega\left(\frac{c-d}{\sqrt{\varepsilon}} \right)$ oracle queries.*

5. CLASSICAL ORACLE SCHEME

In this section, we construct a mini-scheme, called the *Hidden Subspace Mini-Scheme*, that requires only a classical oracle. We then use the inner-product adversary method from Section 4 to show that our mini-scheme is secure—indeed, that any counterfeiter must make $\Omega\left(2^{n/4}\right)$ queries to copy a banknote. By the results of Sections 3.3 and 2.1, our mini-scheme will automatically imply a full-blown public-key quantum money scheme, which requires only a classical oracle and is unconditionally secure.

5.1 The Hidden Subspace Mini-Scheme

We identify n-bit strings $x \in \{0,1\}^n$ with elements of the vector space $\mathbb{F}_2^n$ in the standard way. Then in our mini-scheme, each n-qubit money state will have the form

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

where A is some randomly-chosen subspace of $\mathbb{F}_2^n$ (i.e., a set of codewords of a linear code), with $\dim A = n/2$. Let $A^\perp$ be the orthogonal complement of A, so that $\dim A^\perp = n/2$ as well. Notice that we can transform $|A\rangle$ to $\left|A^\perp\right\rangle$ and vice versa by simply applying $H_2^{\otimes n}$: a Hadamard gate on each of the n qubits, or equivalently a quantum Fourier transform over $\mathbb{F}_2^n$.

The basic idea of the mini-scheme is as follows: the bank can easily prepare the quantum money state $|A\rangle$, starting from a classical description $\langle A \rangle$ of A (e.g., a list of $n/2$ generators). The bank distributes the state $|A\rangle$, but keeps the classical description $\langle A \rangle$ secret. Along with $|A\rangle$ itself, the bank also publishes details of how to *verify* $|A\rangle$ by querying two classical oracles, U_A and $U_{A^\perp}$. The first oracle, U_A, decides membership in A: for all n-qubit basis states $|x\rangle$,

$$U_A |x\rangle = \begin{cases} -|x\rangle & \text{if } x \in A \\ |x\rangle & \text{otherwise} \end{cases}$$

The second oracle, $U_{A^\perp}$, decides membership in $A^\perp$ in the same way.

Using U_A, it is easy to implement a projector $\mathbb{P}_A$ onto the set of basis states in A. To do so, simply initialize a control qubit to $|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}$, then apply U_A conditioned on the control qubit being in state $|1\rangle$, then measure the control qubit in the $\{|+\rangle, |-\rangle\}$ basis, and postselect on getting the outcome $|-\rangle$. Likewise, using $U_{A^\perp}$, it is easy to implement a projector $\mathbb{P}_{A^\perp}$ onto the set of basis states in $A^\perp$. Then V_A, the public verification algorithm for the money state $|A\rangle$, will simply consist of $\mathbb{P}_A$, then a Fourier transform, then $\mathbb{P}_{A^\perp}$, and finally a second Fourier transform to return the legitimate money state back to $|A\rangle$:

$$V_A := H_2^{\otimes n} \mathbb{P}_{A^\perp} H_2^{\otimes n} \mathbb{P}_A.$$

We show in Lemma 19 that $V_A = |A\rangle \langle A|$ is just a projector onto $|A\rangle$. This means, in particular, that $V_A |A\rangle = |A\rangle$, and that V_A accepts an arbitrary state $|\psi\rangle$ with probability $|\langle \psi | A \rangle|^2$. Thus, our mini-scheme is *projective* and has *perfect completeness*.

But what about security? Intuitively, a counterfeiter could query U_A or $U_{A^\perp}$ to find a generating set for A or $A^\perp$—but that would require an exponentially-long Grover search, since $|A| = \left|A^\perp\right| = 2^{n/2} \ll 2^n$. Alternatively, the counterfeiter could measure $|A\rangle$ in the standard or Hadamard bases—but that would reveal just *one* random element of A or $A^\perp$. Neither ability seems useful for *copying* $|A\rangle$, let alone recovering a full classical description of A.[16]

[16] Obviously, if the counterfeiter had $\Omega(n)$ copies of $|A\rangle$, then it *could* recover a generating set for A, by simply measuring each copy independently in the standard basis. That is why, in our full quantum money scheme, the counterfeiter will *not* have $\Omega(n)$ copies of $|A\rangle$. Instead, each banknote will involve a completely different subspace $A_s \leq \mathbb{F}_2^n$ (parameterized by its unique serial number s), so that measuring one banknote reveals nothing about the others.

And indeed, using the inner-product adversary method plus some other tools, we will prove the following tight lower bound (Theorem 23): even if given a single copy of $|A\rangle$, as well as oracle access to U_A and $U_{A^\perp}$, a counterfeiter still needs $\Omega\left(\epsilon 2^{n/4}\right)$ queries to prepare a state that has fidelity ϵ with $|A\rangle^{\otimes 2}$. This will imply that our mini-scheme has *exponentially-small soundness error.*

5.2 Formal Specification

We are not quite done, since we never explained how the bank provides access to U_A and $U_{A^\perp}$. Thus, in our "final" mini-scheme $\mathcal{M} = (\mathsf{Bank}_{\mathcal{M}}, \mathsf{Ver}_{\mathcal{M}})$, the bank, verifier, and counterfeiter will all have access to a *single* classical oracle U, which consists of four components:

A **banknote generator** $\mathcal{G}(r)$, which takes as input a random string $r \in \{0,1\}^n$, and outputs a set of linearly independent generators $\langle A_r \rangle = \{x_1, \ldots, x_{n/2}\}$ for a subspace $A_r \leq \mathbb{F}_2^n$, as well as a unique $3n$-bit *serial number* $s_r \in \{0,1\}^{3n}$. The function $\mathcal{G}$ is chosen uniformly at random, subject to the constraint that the serial numbers are all distinct.[17]

A **serial number checker** $\mathcal{H}(s)$, which outputs 1 if $s = s_r$ is a valid serial number for some $\langle A_r \rangle$, and 0 otherwise.

A **primal subspace tester** $\mathcal{T}_{\text{primal}}$, which takes an input of the form $|s\rangle|x\rangle$, applies U_{A_r} to $|x\rangle$ if $s = s_r$ is a valid serial number for some $\langle A_r \rangle$, and does nothing otherwise.

A **dual subspace tester** $\mathcal{T}_{\text{dual}}$, identical to $\mathcal{T}_{\text{primal}}$ except that it applies $U_{A_r^\perp}$ instead of U_{A_r}.

Then $\mathcal{M} = (\mathsf{Bank}_{\mathcal{M}}, \mathsf{Ver}_{\mathcal{M}})$ is defined as follows:

- $\mathsf{Bank}_{\mathcal{M}}(0^n)$ chooses $r \in \{0,1\}^n$ uniformly at random. It then looks up $\mathcal{G}(r) = (s_r, \langle A_r \rangle)$, and outputs the banknote $|\$_r\rangle = |s_r\rangle|A_r\rangle$.
- $\mathsf{Ver}_{\mathcal{M}}(\phi)$ first uses $\mathcal{H}$ to check that ϕ has the form (s, ρ), where $s = s_r$ is a valid serial number. If so, then it uses $\mathcal{T}_{\text{primal}}$ and $\mathcal{T}_{\text{dual}}$ to apply $V_{A_r} = H_2^{\otimes n}\mathbb{P}_{A_r^\perp}H_2^{\otimes n}\mathbb{P}_{A_r}$, and accepts if and only if $V_{A_r}(\rho)$ accepts.

5.3 Analysis

We now analyze the mini-scheme defined in Sections 5.1 and 5.2. For convenience, we assume for most of the proof that the subspace $A \leq \mathbb{F}_2^n$ is *fixed*, and that the counterfeiter (who does not know A) only has access to the oracles U_A and $U_{A^\perp}$. Then, at the end, we will explain how to generalize the conclusions to the "final" mini-scheme $\mathcal{M}$.

It will be convenient to consider the subset $A^* \subset \{0,1\}^{n+1}$, defined by $A^* := (0, A) \cup (1, A^\perp)$. Let S_{A^*} be the subspace of $\mathbb{C}^{2^{n+1}}$ that is spanned by basis states $|x\rangle$ such that $x \in A^*$. Then we can think of the pair of oracles $(U_A, U_{A^\perp})$ as being a *single* oracle U_{A^*}, which satisfies $U_{A^*}|\psi\rangle = -|\psi\rangle$ for all $|\psi\rangle \in S_{A^*}$, and $U_{A^*}|\eta\rangle = |\eta\rangle$ for all $|\eta\rangle \in S_{A^*}^\perp$.

Recall the definition of the verifier V_A:

$$V_A := H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}\mathbb{P}_A,$$

where $\mathbb{P}_A$ and $\mathbb{P}_{A^\perp}$ denote projective measurements that accept a basis state $|x\rangle$ if and only if x belongs to A or $A^\perp$ respectively. The following lemma shows that V_A "works," and indeed that it gives us a projective mini-scheme.

[17] Note that one can implement $\mathcal{G}$ using an ordinary random oracle. In that case, the requirement that the serial numbers are distinct will be satisfied with probability $1 - O\left(2^{-n}\right)$.

Lemma 19. *$V_A = |A\rangle\langle A|$ is simply a projector onto $|A\rangle$. So in particular, $\Pr[V_A(|\psi\rangle)$ accepts$] = |\langle\psi|A\rangle|^2$.*

Proof. It suffices to show that $V_A|A\rangle = |A\rangle$ and that $V_A|\psi\rangle = 0$ for all $|\psi\rangle$ orthogonal to $|A\rangle$. First,

$$\begin{aligned} V_A|A\rangle &= H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}\mathbb{P}_A|A\rangle \\ &= H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}|A\rangle \\ &= H_2^{\otimes n}\mathbb{P}_{A^\perp}|A^\perp\rangle \\ &= H_2^{\otimes n}|A^\perp\rangle \\ &= |A\rangle. \end{aligned}$$

Second, if $\langle\psi|A\rangle = 0$ then we can write

$$|\psi\rangle = \sum_{x \in 2^n} c_x |x\rangle$$

where $\sum_{x\in A} c_x = 0$. Then

$$\begin{aligned} V_A|\psi\rangle &= H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}\mathbb{P}_A \sum_{x\in 2^n} c_x|x\rangle \\ &= H_2^{\otimes n}\mathbb{P}_{A^\perp}H_2^{\otimes n}\sum_{x\in A} c_x|x\rangle \\ &= \frac{1}{\sqrt{2^n}}H_2^{\otimes n}\mathbb{P}_{A^\perp}\sum_{x\in A} c_x \sum_{y\perp x}|y\rangle \\ &= \frac{1}{\sqrt{2^n}}H_2^{\otimes n}\sum_{y\in A^\perp}|y\rangle\sum_{x\in A}c_x \\ &= 0. \end{aligned}$$

■

We now show that perfect counterfeiting requires exponentially many queries to U_{A^*}.

Theorem 20 (Lower Bound for Perfect Counterfeiting)**.** *Given one copy of $|A\rangle$, as well as oracle access to U_{A^*}, a counterfeiter needs $\Omega\left(2^{n/4}\right)$ queries to prepare $|A\rangle^{\otimes 2}$ with certainty (for a worst-case $|A\rangle$).*

Proof. We will apply Theorem 18. Let the set $\mathcal{O}$ contain U_{A^*} for every possible subspace $A \leq \mathbb{F}_2^n$ with $\dim A = n/2$. Also, put $(U_{A^*}, U_{B^*}) \in R$ if and only if $\dim(A \cap B) = n/2 - 1$. Then given $U_{A^*} \in \mathcal{O}$ and $|\eta\rangle \in S_{A^*}^\perp$, let

$$|\eta\rangle = \sum_{x\in\{0,1\}^{n+1}\setminus A^*} \alpha_x|x\rangle.$$

We have

$$\begin{aligned} &\mathop{\mathrm{E}}_{U_{B^*}\,:\,(U_{A^*},U_{B^*})\in R}\left[F(|\eta\rangle, S_{B^*})^2\right] \\ &= \mathop{\mathrm{E}}_{B\,:\,\dim(B)=n/2,\dim(A\cap B)=n/2-1}\left[\sum_{x\in B^*\setminus A^*}|\alpha_x|^2\right] \\ &\leq \max_{x\in\{0,1\}^{n+1}\setminus A^*}\left(\Pr_{B\,:\,\dim(B)=n/2,\dim(A\cap B)=n/2-1}[x\in B^*]\right) \\ &= \max_{x\in\{0,1\}^n\setminus A}\left(\Pr_{B\,:\,\dim(B)=n/2,\dim(A\cap B)=n/2-1}[x\in B]\right) \\ &= \frac{|B\setminus A|}{|\{0,1\}^n\setminus A|} \quad (\text{for } \dim(B)=n/2,\ \dim(A\cap B)=n/2-1) \\ &= \frac{2^{n/2-1}}{2^n - 2^{n/2}} \\ &\leq \frac{1}{2^{n/2}}. \end{aligned}$$

Here the first line uses the definition of fidelity, the second line uses the easy direction of the minimax theorem, the third line uses the symmetry between A and $A^{\perp}$, and the fourth line uses the symmetry among all $2^n - 2^{n/2}$ strings $x \in \{0,1\}^n \backslash A$. The conclusion is that we can set $\varepsilon := 2^{-n/2}$.

Fix $(U_{A^*}, U_{B^*}) \in R$. Then $|\langle A|B\rangle| = 1/2$. On the other hand, if the counterfeiter succeeds, it must map $|A\rangle$ to some state $|f_A\rangle := |A\rangle |A\rangle |\text{garbage}_A\rangle$, and $|B\rangle$ to some state $|f_B\rangle := |B\rangle |B\rangle |\text{garbage}_B\rangle$. Therefore $|\langle f_A|f_B\rangle| \leq 1/4$. So setting $c = 1/2$ and $d = 1/4$, Theorem 18 tells us that the counterfeiter must make

$$\Omega\left(\frac{c-d}{\sqrt{\varepsilon}}\right) = \Omega\left(2^{n/4}\right)$$

queries to U_{A^*}. ■

A simple modification to the proof of Theorem 20 shows that even to counterfeit money *almost* perfectly, one still needs exponentially many queries to U_{A^*}.

Corollary 21 (Lower Bound for Small-Error Counterfeiting)**.** *Given one copy of $|A\rangle$, as well as oracle access to U_{A^*}, a counterfeiter needs $\Omega\left(2^{n/4}\right)$ queries to prepare a state ρ such that $\langle A|^{\otimes 2} \rho |A\rangle^{\otimes 2} \geq 0.9999$ (for a worst-case $|A\rangle$).*

Proof. Let $|\langle A|B\rangle| = c$, and let $\epsilon = 0.00001$. If the counterfeiter succeeds, it must map $|A\rangle$ to some state ρ_A, and $|B\rangle$ to some state ρ_B, such that $\langle A|^{\otimes 2} \rho_A |A\rangle^{\otimes 2}$ and $\langle B|^{\otimes 2} \rho_B |B\rangle^{\otimes 2}$ are both at least $1-\epsilon$. So letting $|f_A\rangle$ and $|f_B\rangle$ be purifications of ρ_A and ρ_B respectively, we have

$$\begin{aligned} |\langle f_A|f_B\rangle| &\leq F(\rho_A, \rho_B) \\ &\leq \left|\langle A|^{\otimes 2} |B\rangle^{\otimes 2}\right| + 2\epsilon^{1/4} \\ &= c^2 + 2\epsilon^{1/4} \end{aligned}$$

where the second line follows from Lemma 4. So setting $d := c^2 + 2\epsilon^{1/4}$, Theorem 18 tells us that the counterfeiter must make

$$\Omega\left(\frac{c - c^2 - 2\epsilon^{1/4}}{\sqrt{2^{-n/2}}}\right)$$

queries to U_{A^*}. Fixing $c := 1/2$, the above is $\Omega\left(2^{n/4}\right)$. ■

Since the verifier V_A is projective, we can now combine Corollary 21 with Theorem 13 to obtain the following "amplified" lower bound.

Corollary 22 (Lower Bound for High-Error Counterfeiting)**.** *Let $1/\varepsilon = o\left(2^{n/2}\right)$. Given one copy of $|A\rangle$, as well as oracle access to U_{A^*}, a counterfeiter needs $\Omega\left(\sqrt{\varepsilon} 2^{n/4}\right)$ queries to prepare a state ρ such that $\langle A|^{\otimes 2} \rho |A\rangle^{\otimes 2} \geq \varepsilon$ (for a worst-case $|A\rangle$).*

Proof. Suppose we have a counterfeiter C that makes $o\left(\sqrt{\varepsilon} 2^{n/4}\right)$ queries to U_{A^*}, and prepares a state σ such that $\langle A|^{\otimes 2} \sigma |A\rangle^{\otimes 2} \geq \varepsilon$. Let $\delta := 0.00001$. Then by Theorem 13, there exists an amplified counterfeiter C' that makes

$$O\left(\frac{\log 1/\delta}{\sqrt{\varepsilon}\left(\sqrt{\varepsilon} + \delta^2\right)}\right) = O\left(\frac{1}{\sqrt{\varepsilon}}\right)$$

calls to C and V_A, and that prepares a state ρ such that $\langle A|^{\otimes 2} \rho |A\rangle^{\otimes 2} \geq 1 - \delta$. Now, counting the $o\left(\sqrt{\varepsilon} 2^{n/4}\right)$ queries from each C invocation and $O(1)$ queries from each V_A invocation, the total number of queries that C' makes to U_{A^*} is

$$\left[o\left(\sqrt{\varepsilon} 2^{n/4}\right) + O(1)\right] \cdot O\left(\frac{1}{\sqrt{\varepsilon}}\right) = o\left(2^{n/4}\right).$$

But this contradicts Corollary 21. ■

So far, we have only made statements about the *worst* case for a would-be counterfeiter. But such guarantees are clearly not enough: it could be that *most* money states $|A\rangle$ are easy to duplicate, without contradicting any of the results we have seen so far.

We will show that the problem faced by a counterfeiter is *random self-reducible*: if a counterfeiter could duplicate a uniformly-random money state $|A\rangle$, then it could duplicate *any* $|A\rangle$. Thus the bank can ensure security by creating uniformly-random money states.

In what follows, let $\mathcal{S}$ be the set of all subspaces $A \leq \mathbb{F}_2^n$ such that $\dim A = n/2$. Also, let $V_A^{\otimes 2} = (|A\rangle \langle A|)^{\otimes 2}$ be the projector onto $|A\rangle^{\otimes 2}$.

Theorem 23 (Lower Bound for Average-Case Counterfeiting)**.** *Let $A \leq \mathbb{F}_2^n$ be a uniformly-random element of $\mathcal{S}$. Then given one copy of $|A\rangle$, as well as oracle access to U_{A^*}, a counterfeiter C needs $\Omega\left(\sqrt{\varepsilon} 2^{n/4}\right)$ queries to prepare a $2n$-qubit state ρ that $V_A^{\otimes 2}$ accepts with probability at least ε, for all $1/\varepsilon = o\left(2^{n/2}\right)$. Here the probability is taken over the choice of $A \in \mathcal{S}$, as well as the behavior of C and $V_A^{\otimes 2}$.*

Proof. Suppose we had a counterfeiter C that violated the above. Using C as a black box, we will show how to construct a new counterfeiter C' that violates Corollary 22.

Given a (deterministically-chosen) money state $|A\rangle$ and oracle access to U_{A^*}, first choose an invertible linear map $f : \mathbb{F}_2^n \to \mathbb{F}_2^n$ uniformly at random. Then $f(A)$, the image of A under f, is a uniformly-random element of $\mathcal{S}$. Furthermore, the state $|A\rangle$ can be transformed into $|f(A)\rangle$ straightforwardly, and the oracle $U_{f(A)^*}$ can be simulated by composing f with U_{A^*}. So by using the counterfeiter C for uniformly-random states, we can produce a state ρ_f that $V_{f(A)}^{\otimes 2}$ accepts with probability at least ε. By applying f^{-1} to both registers of ρ_f, we can then obtain a state ρ that $V_A^{\otimes 2}$ accepts with probability at least ε, thereby contradicting Corollary 22. ■

We are now ready to prove security for the "final" mini-scheme $\mathcal{M}$ defined in Section 5.2.

Theorem 24 (Security of Mini-Scheme)**.** *The mini-scheme $\mathcal{M} = (\mathsf{Bank}_{\mathcal{M}}, \mathsf{Ver}_{\mathcal{M}})$, which is defined relative to the classical oracle U, has perfect completeness and exponentially-small soundness error.*

Proof Sketch. That $\mathcal{M}$ has perfect completeness follows from its definition and from Lemma 19. That $\mathcal{M}$ has exponentially-small soundness error *essentially* follows from Theorem 23. We only need to show that, given a banknote of the form $|\$_r\rangle = |s_r\rangle |A_r\rangle$, a polynomial-time counterfeiter C can gain no additional advantage by querying the "full" oracles $\mathcal{G}, \mathcal{H}, \mathcal{T}_{\text{primal}}, \mathcal{T}_{\text{dual}}$, beyond what it gains from querying $U_{A_r^*} = \left(U_{A_r}, U_{A_r^{\perp}}\right)$. This follows from some simple observations: first, let $r \in \{0,1\}^n$ be the random string chosen by the bank, so that $\mathcal{G}(r) = (s_r, \langle A_r\rangle)$. Then assuming C does not know r, the BBBV hybrid argument [12] tells us that C can gain nothing by querying $\mathcal{G}$: indeed, if we

randomly *change* the value of $\mathcal{G}(r)$, it will affect C's output by at most an exponentially small amount. However, once we make that change, an adversary trying to counterfeit $|A\rangle$ given U_A and $U_{A^\perp}$ can easily "mock up" a serial number s, as well as the oracles $\mathcal{G}, \mathcal{H}, \mathcal{T}_{\text{primal}}$ and $\mathcal{T}_{\text{dual}}$, for itself. Just like in Corollary 16, since our security guarantees are query complexity bounds, we do not care about the *computational* complexity of creating the mock-ups. By using the mock-ups, one can convert any successful attack on $\mathcal{M}$ into successful counterfeiting of $|A\rangle$, given oracle access to U_A and $U_{A^\perp}$ only. But the latter contradicts Theorem 23. Further details of the argument are deferred to the full version. ■

Finally, using Theorem 24 together with Corollary 16, we can obtain a secure public-key quantum money scheme, relative to a classical oracle.

Theorem 25 (Security of Hidden Subspace Money)**.** *By combining the mini-scheme $\mathcal{M}$ with a digital signature scheme, it is possible to construct a public-key quantum money scheme $\mathcal{S} = (\mathsf{KeyGen}_{\mathcal{S}}, \mathsf{Bank}_{\mathcal{S}}, \mathsf{Ver}_{\mathcal{S}})$, defined relative to some classical oracle U', which has perfect completeness and exponentially-small soundness error.*

6. EXPLICIT QUANTUM MONEY SCHEME

We have shown how to construct a provably-secure public-key quantum money scheme, when an appropriate classical oracle is available. In this section, we propose a way to obtain the same functionality without an oracle. The key challenge is this:

> *Given a subspace $A \leq \mathbb{F}_2^n$, how can a bank distribute an "obfuscated program" P_A, which legitimate buyers and sellers can use to decide membership in both A and $A^\perp$, but which does not reveal anything else about A that might facilitate counterfeiting?*

Note that, aside from the detail that we need security against quantum adversaries, the above challenge is purely "classical"; it and its variants seem interesting even apart from our quantum money application.

We will suggest a candidate protocol to achieve the challenge, based on *multivariate polynomial cryptography*. Given a collection $p_1, \ldots, p_m : \mathbb{F}_2^n \to \mathbb{F}_2$ of multivariate polynomials over $\mathbb{F}_2$, it is generally hard to find a point $v \in \mathbb{F}_2^n$ on which all of the p_i's vanish. On the other hand, it is easy to check whether a *particular* point v has that property. To "hide" a subspace A, we will provide uniformly-random low-degree polynomials $p_1, \ldots, p_m$ that vanish on each point of A. This information is sufficient to decide membership in A. On the other hand, there is no known efficient algorithm to *find* A given the polynomials, and current techniques seem unlikely to yield even a quantum algorithm.

We can also introduce a constant fraction of *noise* into our scheme without interfering with its completeness. In other words, if only $(1-\epsilon)m$ of the polynomials $p_1, \ldots, p_m$ are chosen to vanish on A, and the remaining ϵm are random, then counting the number of p_i's that vanish at a point v still suffices to determine whether $v \in A$. Although we know of no attack even against our noise-free scheme, adding noise in this way might improve security.

Crucially, we will state a "classical" conjecture about the security of multivariate polynomial cryptography, and show that the conjecture *implies* the security of our explicit money scheme. For the benefit of cryptographers, let us now state an "abstract" version of our conjecture, which implies what we need, and which might hold even if our concrete conjecture about multivariate polynomials fails.

Conjecture 26 (Subspace-Hiding Conjecture, Sufficient for Quantum Money)**.** *There exists a polynomial-time algorithm that takes as input a description of a uniformly-random subspace $A \leq \mathbb{F}_2^n$ with $\dim(A) = n/2$, and that outputs circuits C_A and $C_{A^\perp}$, such that the following holds.*

(i) $C_A(v)$ decides whether $v \in A$, and $C_{A^\perp}(v)$ decides whether $v \in A^\perp$, for all $v \in \mathbb{F}_2^n$.

(ii) Given descriptions of C_A and $C_{A^\perp}$, no polynomial-time quantum algorithm can find a generating set for A with success probability $\Omega\left(2^{-n/2}\right)$.

Later, Conjecture 32 will specialize Conjecture 26 to the setting of multivariate polynomials.

6.1 Useful Facts About Polynomials

By viewing elements of $\mathbb{F}_2^n$ as n-tuples $(x_1, \ldots, x_n)$, we can evaluate a polynomial $p(x_1, \ldots, x_n)$ on points of $\mathbb{F}_2^n$.

Given a subspace $A \leq \mathbb{F}_2^n$ and a positive integer d, let $\mathcal{I}_{d,A}$ be the set of degree-d polynomials (not necessarily homogeneous) that vanish on A. Since we are working over $\mathbb{F}_2$, note that $x_i^2 = x_i$, so it suffices to consider *multilinear* polynomials (in which no x_i is ever raised to a higher power than 1).

Before presenting our scheme, we need to establish some basic properties of polynomials over $\mathbb{F}_2^n$. First, we observe that the set of polynomials does not depend on the choice of basis.

Proposition 27. *Let L be any invertible linear transformation on $\mathbb{F}_2^n$. Then the map $p(v) \mapsto p(Lv)$ defines a permutation on the set of degree-d polynomials, which maps $\mathcal{I}_{d,A}$ to $\mathcal{I}_{d,L^{-1}A}$.*

Implementing our scheme will require sampling uniformly from $\mathcal{I}_{d,A}$, which the next lemma shows is possible.

Lemma 28. *It is possible to sample a uniformly-random element of $\mathcal{I}_{d,A}$ in time $O(n^d)$.*

PROOF. By Proposition 27, we can instead sample from the space of polynomials which vanish on span $\left(x_1, \ldots, x_{n/2}\right)$, and then apply an appropriate change of basis to obtain a sample from $\mathcal{I}_{d,A}$. So assume without loss of generality that $A = \text{span}\left(x_1, \ldots, x_{n/2}\right)$.

We claim that a polynomial p vanishes on A if and only if every monomial of p intersects $\left\{x_{n/2+1}, \ldots, x_n\right\}$. This will immediately give an $O\left(n^d\right)$-time sampling algorithm, because we can consider each of the $O\left(n^d\right)$ degree-d monomials in turn, and include each one independently with probability 1/2 if it intersects $\left\{x_{n/2+1}, \ldots, x_n\right\}$.

To prove the claim: first, if every monomial intersects $\left\{x_{n/2+1}, \ldots, x_n\right\}$, then clearly p vanishes on A. Otherwise, let m be a minimal monomial that does *not* intersect $\left\{x_{n/2+1}, \ldots, x_n\right\}$. Consider the vector $v = (v_1, \ldots, v_n)$ with $v_i = 1$ if and only if $x_i \in m$. Since m does not intersect $\left\{x_{n/2+1}, \ldots, x_n\right\}$, clearly $v \in A$. Also, since m is minimal, every other monomial must evaluate to 0 on v. Thus $p(v) = m(v) = 1$, so p is not identically zero on A. ■

In addition to sampling polynomials that vanish on A, we would like to guarantee that a sufficiently large system of such polynomials uniquely determines the space A, so that such a system can be effectively used as a membership oracle.

Lemma 29. *Fix $A \leq \mathbb{F}_2^n$ and $\beta > 1$, and choose βn polynomials $p_1, \ldots, p_{\beta n}$ uniformly and independently from $\mathcal{I}_{d,A}$. Let Z be the set of $v \in \mathbb{F}_2^n$ such that $p_i(v) = 0$ for all $i \in [\beta n]$. Then $A \subseteq Z$, and $\Pr[A = Z] = 1 - 2^{-\Omega(n)}$.*

PROOF. $A \subseteq Z$ is clear. For the second part, for each $v \notin A$, there must be some $w \in A^\perp$ such that $w \cdot v = 1$. Then the map $p(v) \mapsto p(v) + w \cdot v$ defines an involution of $\mathcal{I}_{d,A}$, such that exactly one of $p(v)$ and $p(v) + w \cdot v$ is zero. This means that exactly half of the polynomials in $\mathcal{I}_{d,A}$ vanish at v. Hence $\Pr[p_1(v) = \cdots = p_{\beta n}(v) = 0] = 2^{-\beta n}$. By the union bound, the probability that there are *any* shared zeroes $v \notin A$ is therefore at most $2^n 2^{-\beta n} = 2^{-\Theta(n)}$. ■

As mentioned earlier, we would also like to allow sampling from *noisy* systems of equations, defined as follows: let $\mathcal{R}_{d,A,m,\epsilon}$ be the probability distribution over m-tuples $(p_1, \ldots, p_m)$ that sets exactly $(1-\epsilon)m$ of the polynomials p_i (chosen uniformly at random) to be uniformly-random samples from $\mathcal{I}_{d,A}$, and that sets the remaining ϵm of the polynomials p_i to be uniformly-random samples from $\mathcal{I}_{d,A'}$, for a uniformly-random subspace $A' \leq \mathbb{F}_2^n$ of dimension $\dim(A)$. (Note that a *different* A' is chosen for every such p_i.) Then using a Chernoff bound, it is not hard to show that, provided m is large enough compared to n, a sample from $\mathcal{R}_{d,A,m,\epsilon}$ *also* uniquely defines the subspace A with overwhelming probability.

Lemma 30. *Fix $A \leq \mathbb{F}_2^n$ and $\epsilon < 1/2$, let $\beta \geq \frac{3}{(1-2\epsilon)^2}$, and choose polynomials $p_1, \ldots, p_{\beta n}$ from $\mathcal{R}_{d,A,\beta n,\epsilon}$. Let $w(v) := \sum_{i=1}^{\beta n} p_i(v)$, and let Z be the set of $v \in \mathbb{F}_2^n$ such that $w(v) \leq \epsilon \beta n$. Then $A \subseteq Z$, and $\Pr[A = Z] = 1 - 2^{-\Omega(n)}$.*

PROOF. Deferred to the full version. ■

6.2 Explicit Hidden-Subspace Mini-Scheme

In our explicit mini-scheme, the bank chooses a subspace A randomly and publishes sets of polynomials drawn from $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^\perp,\beta n,\epsilon}$, along with the quantum money state $|A\rangle$. By Lemma 30, a user can use these polynomials to test membership in A and $A^\perp$, and can therefore implement the oracle mini-scheme in Section 5.1.

Formally, the mini-scheme $\mathcal{E}$ is defined as follows. Parameters $\epsilon \in [0, 1/2)$, $\beta \geq \frac{3}{(1-2\epsilon)^2}$, and $d \geq 4$ are fixed. The complexity of the verification procedure will grow like $O\left(\beta n^{d+1}\right)$, but security might also improve for larger ϵ and d. Then:

- $\mathsf{Bank}(0^n)$ selects an $n/2$-dimensional subspace $A \leq \mathbb{F}_2^n$ uniformly at random, say by selecting $n/2$ random linearly-independent generators. It then sets $s := (s_A, s_{A^\perp})$, where s_A and $s_{A^\perp}$ are lists of polynomials drawn from $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^\perp,\beta n,\epsilon}$ respectively. It prepares the money state $|A\rangle$ and outputs the banknote $|\$_s\rangle := |s\rangle|A\rangle$.
- $\mathsf{Ver}(\phi)$ first checks that ϕ has the form $(s_A, s_{A^\perp}, \rho)$ where $s_A = (p_1, \ldots, p_{\beta n})$ and $s_{A^\perp} = (q_1, \ldots, q_{\beta n})$ are lists of βn polynomials over $\mathbb{F}_2^n$. If not, it rejects. If so, then it lets Z and $Z^\perp$ be the sets of points $v \in \mathbb{F}_2^n$ such that $\sum_{i=1}^{\beta n} p_i(v) \leq \epsilon \beta n$ and $\sum_{i=1}^{\beta n} q_i(v) \leq \epsilon \beta n$ respectively. (Recall that with overwhelming probability, $Z = A$ and $Z^\perp = A^\perp$.) It then applies the operation $V_Z := H_2^{\otimes n} \mathbb{P}_{Z^\perp} H_2^{\otimes n} \mathbb{P}_Z$ to ρ, and accepts ϕ if and only if $V_Z(\rho)$ accepts.

6.3 Analysis

We first observe that the mini-scheme $\mathcal{E}$ has perfect completeness.

Theorem 31. *$\mathcal{E}$ has perfect completeness.*

PROOF. This follows from Lemmas 29 and 30, and particularly from the fact that $A \subseteq Z$ and $A^\perp \subseteq Z^\perp$ with certainty. From this it follows that $V_Z := H_2^{\otimes n} \mathbb{P}_{Z^\perp} H_2^{\otimes n} \mathbb{P}_Z$ accepts the state $|A\rangle$ with probability 1. ■

Let us remark that, if the fraction ϵ of bad polynomials is in $[1/2, 1)$, one can still define a variant of our scheme—for example, where Ver guesses that $v \in A$ if $\sum_{i=1}^{\beta n} p_i(v) \leq \frac{1}{4}(1+\epsilon)\beta n$, and that $v \notin A$ otherwise. The only disadvantage is that we now lose the property of perfect completeness, and can only guarantee a completeness error that is exponentially small.

We now wish to argue about $\mathcal{E}$'s soundness. Naturally, we can only hope to prove soundness assuming some computational hardness conjecture. What is nice, though, is that we can base $\mathcal{E}$'s soundness on a conjecture that talks only about the hardness of a "classical" cryptographic problem (i.e., a problem with classical inputs and outputs). Let us now state that conjecture, which is simply the abstract Conjecture 26 specialized to the setting of multivariate polynomials.

Conjecture 32 (Direct Product for Finding Subspace Elements)**.** *Let $\epsilon < 1/2$ and $\beta := \frac{3}{(1-2\epsilon)^2}$. Given samples from $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^\perp,\beta n,\epsilon}$, no polynomial-time quantum algorithm can find a complete list of generators for A with success probability $\Omega\left(2^{-n/2}\right)$.*

Note that it is easy to find *one* nonzero element of A with success probability $2^{-n/2}$, by choosing $x \in \mathbb{F}_2^n$ randomly. Conjecture 32 asserts both that it is impossible to do too much better using $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^\perp,\beta n,\epsilon}$, and that finding multiple elements of A is significantly harder than finding one element.

The security of mini-scheme $\mathcal{E}$ follows easily from Conjecture 32, *despite* the fact that a would-be counterfeiter has access to a valid quantum banknote, whereas Conjecture 32 involves no such assumption.

Theorem 33 (Security Reduction for Explicit Mini-Scheme)**.** *If Conjecture 32 holds, then $\mathcal{E}$ is secure.*

PROOF. Let $C_{\mathcal{E}}$ be a counterfeiter against $\mathcal{E}$. Then we need to show that, using $C_{\mathcal{E}}$, we can find a complete list of generators for A with $\Omega\left(2^{-n/2}\right)$ success probability.

Given $A \leq \mathbb{F}_2^n$ with $\dim(A) = n/2$, let $s := (s_A, s_{A^\perp})$ where s_A and $s_{A^\perp}$ are samples from $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^\perp,\beta n,\epsilon}$ respectively. Recall from Lemma 30 that $\Pr[A = Z] = 1 - 2^{-\Omega(n)}$ and $\Pr\left[A^\perp = Z^\perp\right] = 1 - 2^{-\Omega(n)}$. Provided both of these events occur, we can use s to decide membership in A, and can therefore apply the projective measurement $\mathbb{P}_A$. So let us prepare the uniform superposition over all 2^n elements of $\mathbb{F}_2^n$, and then apply $\mathbb{P}_A$ to it. With probability $2^{-n/2}$, this produces the state $|A\rangle$.

Once we have s and $|A\rangle$, we can then form the banknote $|\$\rangle := |s\rangle |A\rangle$, and provide this banknote to the counterfeiter $C_{\mathcal{E}}$. By hypothesis, $C_{\mathcal{E}}$ outputs a (possibly-entangled) state ρ on two registers, such that $\langle A|^{\otimes 2} \rho |A\rangle^{\otimes 2} \geq \Delta$ for some $\Delta = \Omega(1/\operatorname{poly}(n))$. But now, because the mini-scheme $\mathcal{E}$ is projective, Theorem 13 applies, and we can *amplify* ρ to increase its fidelity with $|A\rangle^{\otimes 2}$. After $O\left(\frac{1}{\Delta^2} \log n\right)$ calls to $C_{\mathcal{E}}$, this gives us a state σ such that

$$\langle A|^{\otimes 2} \sigma |A\rangle^{\otimes 2} \geq 1 - \frac{1}{n^2}.$$

More generally, by alternating counterfeiting steps and amplification steps, we can produce as many registers as we like that each have large overlap with $|A\rangle$. In particular, we can produce a state ξ such that

$$\langle A|^{\otimes n} \xi |A\rangle^{\otimes n} \geq 1 - o(1).$$

If we now run Ver on each of the registers of ξ, the probability that every invocation accepts is $1 - o(1)$. Furthermore, supposing that happens, the state we are left with is simply $|A\rangle^{\otimes n}$. Finally, we measure each register of $|A\rangle^{\otimes n}$ in the standard basis. This gives us n elements $x_1, \ldots, x_n \in A$, which are independent and uniformly random. So by standard estimates, the probability that $x_1, \ldots, x_n$ do *not* contain a complete generating set for A is $1/\exp(n)$.

Overall, the procedure above succeeded with probability $2^{-n/2}(1 - o(1))$, thereby giving us the desired contradiction with Conjecture 32. ■

Using the standard construction of quantum money schemes, we can now produce a complete explicit money scheme, whose security follows from Conjecture 32.

Theorem 34 (Security Reduction for Explicit Scheme)**.** *Assuming Conjecture 32, there exists a public-key quantum money scheme with perfect completeness and soundness error* $2^{-\Omega(n)}$.

Proof. We apply the standard construction of Theorem 14 with the mini-scheme $\mathcal{E}$, whose completeness and soundness follow from Theorems 31 and 33 respectively, assuming Conjecture 32. ■

6.4 Justifying Our Hardness Assumption

Though our hardness assumption is new, it is closely related to standard assumptions in *multivariate polynomial cryptography*. Given a system of multivariate quadratics over $\mathbb{F}_2$, finding a common zero is known to be NP-hard; moreover, it is strongly believed that the problem remains hard even for *random* systems of multivariate polynomials, and cryptosystems based on this hardness assumption are considered promising candidates for post-quantum cryptography [20]. Therefore, if Conjecture 32 fails, it will almost certainly be because some additional structure in this problem facilitates a new attack.

There are several ways in which Conjecture 32 is stronger than the assumption that solving random systems of multivariate polynomials is hard. First, our systems have large, well-structured solution spaces A and $A^{\perp}$. Systems with many solutions are not normally considered in the literature, and while there seem to be no known attacks that exploit this structure, the possibility is not ruled out. Second, we provide two related systems, one with zeroes in A and one with zeroes in $A^{\perp}$. Again, this is a very specific structural property which has not been considered, and there might be unexpected attacks exploiting it. Third, Conjecture 32 asserts that no adversary can succeed with probability $2^{-n/2}$, which seems significantly easier than succeeding with non-negligible probability.

On the other hand, Conjecture 32 is *weaker* than typical assumptions in multivariate polynomial cryptography in at least one respect: a would-be counterfeiter needs to solve a system of polynomial equations with a constant fraction of noise. Solving noisy systems of *linear* equations over $\mathbb{F}_2$ is called the *learning parity with noise* problem, and is generally believed to be hard even for quantum computers [36]. If true, this suggests that Gaussian elimination is fundamentally hard to adapt to the presence of noise. But computing a Gröbner basis is a strict generalization of Gaussian elimination to higher degree, and involves a nearly identical process of elimination. It therefore seems unlikely that these approaches can be efficiently adapted to the setting with noise. The problem of solving polynomials with noise has been studied recently, and the best-known approaches involve performing an exponential time search to determine which equations are noisy [6].

But if solving linear systems with noise is already hard, why do we even use higher-degree polynomials in our scheme? The reason is that, alas, the "dual" structure of our money scheme facilitates a simple attack in the case $d = 1$.

Claim 35. *For all $\epsilon < 1$, there exists a β such that one can recover A efficiently given samples from $\mathcal{R}_{d,A,\beta n,\epsilon}$ and $\mathcal{R}_{d,A^{\perp},\beta n,\epsilon}$.*

Proof. Let $p_1, \ldots, p_m$ and $q_1, \ldots, q_m$ be homogeneous linear polynomials, of which a $1 - \epsilon$ fraction vanish on A and $A^{\perp}$ respectively. Then the key observation is that each p_i vanishes on A if and only if it has the form $p_i(v) = u_i \cdot v$ for some $u_i \in A^{\perp}$, while each q_i vanishes on $A^{\perp}$ if and only if it has the form $q_i(v) = w_i \cdot v$ for some $w_i \in A$. But by Lemma 30, if $\beta > \frac{3}{(1-2\epsilon)^2}$, then for each $i \in [m]$, we can efficiently *decide* whether $u_i \in A^{\perp}$ by counting the number of j's for which $q_j(u_i) = 0$, and can likewise decide whether $w_i \in A$ by counting the number of j's for which $p_j(w_i) = 0$.[18] Thus we can learn $\Theta(n)$ random elements of A or $A^{\perp}$, and thereby recover a basis for A. ■

There *might* be a more sophisticated attack for higher degrees, but this is suggested only weakly by the existence of an attack in the linear case. Indeed, the relation between the complementary linear subspaces A and $A^{\perp}$ is precisely the sort of structure that should be preserved by linear maps, but *not* by higher-degree polynomials!

For degree-2 polynomials, it is possible to obtain a similar attack which recovers A from only a *single* sample. This attack relies on the observation that quadratics have an easily-computed canonical form [17], from which a basis for A can be extracted in polynomial time. The essential problem is that quadratic polynomials are very closely related to bilinear forms, and that powerful methods from linear algebra can therefore be applied to them.

Fortunately, the linear structure seems to be computationally obscured when $d \geq 3$. This phenomenon is related to the sharp discontinuity in the difficulty of tensor problems

[18] If $\epsilon \geq 1/2$, then we can use a variant of Lemma 30 (not stated in this extended abstract), for which it suffices to take $\beta \geq \frac{12}{(1-\epsilon)^2}$. We lose perfect completeness, but that is not important here.

with order 3 and higher. More concretely, the coefficients of a degree-d polynomial can be viewed as the entries of an order-d tensor, and the existence of an attack in the degree $d = 2$ case corresponds to the possibility of efficient operations on order-2 tensors. Basic operations on order-3 tensors are NP-hard [26], however, and this suggests that analogous attacks might not exist against degree-3 polynomials.

This state of affairs is reflected in existing attacks on a standard cryptographic assumption called *polynomial isomorphism with one secret*. Here we are given two polynomials p, q which are related by an unknown linear change of coordinates L, and the task is to find such an L. For degree-2 polynomials, this problem can be easily solved in polynomial time [17], but already for degree-3 polynomials the best known attacks take exponential time [35, 24, 17]. However, if an attacker is given n bits of partial information about the linear transformation, then even in the $d = 3$ case, it becomes possible to find the linear transformation that relates the polynomials [17]. This does *not* directly facilitate an attack on our assumption, but it suggests that a similar attack *might* be possible when $d = 3$, since an attacker is only required to succeed with $2^{-n/2}$ probability. Fortunately, this attack seems to rely on the particular structure of degree 2 and 3 polynomials. Of course it is possible that similar algorithms may be discovered for higher-degree polynomials, but this would represent an advance in algebraic cryptanalysis.

7. PRIVATE-KEY QUANTUM MONEY

Recall that a *private-key* quantum money scheme is one where only the bank itself is able to verify banknotes, using an n-bit key $k = k_{\text{private}} = k_{\text{public}}$ that it keeps a closely-guarded secret. Compensating for this disadvantage, private-key schemes are known with much stronger security guarantees than seem possible for public-key schemes.

In particular, as mentioned in Section 1.1, already forty years ago Wiesner [39] described how to create private-key quantum money that is *information-theoretically secure*. In Wiesner's scheme, each banknote consists of n unentangled qubits together with a classical serial number s. Wiesner's scheme also requires a giant database of serial numbers maintained by the bank, or in our setting, access to a random oracle R. But in followup work, BBBW [14] pointed out that we can simply replace R by any *pseudorandom function family* $\{f_k\}_k$, to obtain a private-key quantum money scheme that is computationally secure, *unless* a polynomial-time counterfeiter can distinguish the f_k's from random functions.

Strangely, we are unaware of any rigorous proof of the security of Wiesner's scheme until recently. However, answering a question by one of us,[19] Molina, Vidick and Watrous [31] have now supplied the key ingredient for a security proof. Specifically they show that, if a counterfeiter tries to copy an n-qubit banknote $|\$\rangle$ in Wiesner's scheme, then the output can have squared fidelity at most $(3/4)^n$ with $|\$\rangle^{\otimes 2}$. (They also show that this is tight: there *exists* a non-obvious counterfeiting strategy that succeeds with $(3/4)^n$ probability.)

To complete the security proof, one needs to show that, even given q banknotes $|\$_1\rangle, \ldots, |\$_q\rangle$, a counterfeiter cannot prepare an additional banknote with non-negligible probability (even with a new serial number). In a forthcoming paper [4], we will show how to adapt the methods of Section 3 to prove that claim. Briefly, one can first define a notion of *private-key mini-schemes*, in close analogy to public-key mini-schemes. The work of Molina et al. [31] then directly implies the security of what we call the "Wiesner mini-scheme." Next, one can give a general reduction, showing how to construct a full-blown private-key quantum money scheme $\mathcal{S}$ starting from

(1) any private-key *mini*-scheme $\mathcal{M}$, and

(2) any random or pseudorandom function family R.

Though the details turn out to be more complicated in the private-key case, the proof of correctness for this reduction is conceptually similar to the proof of Theorem 14. Namely, one shows that any counterfeiter would yield *either* a break of the underlying mini-scheme $\mathcal{M}$, or *else* a way to distinguish R from a random function. Notice that the analysis is completely unified: if R is a "true" random oracle, then we get information-theoretic security (as in Wiesner's scheme), while if R is pseudorandom, then we get computational security (as in the BBBW scheme).

Unfortunately, as pointed out by Lutomirski [28] and Aaronson [3], the Wiesner and BBBW schemes both have a serious security hole. Namely, suppose a counterfeiter C can repeatedly submit alleged banknotes to a "naïve and trusting bank" for verification. Given a quantum state σ, such a bank not only tells C whether the verification procedure accepted or rejected, but also, in either case, *gives the post-measurement state $\widetilde{\sigma}$ back to C*. Then starting from a single valid banknote $|\$\rangle$, we claim that C can recover a complete classical description of $|\$\rangle$, using $O(n \log n)$ queries to the bank. Once it has such a description, C can of course prepare as many copies of $|\$\rangle$ as it likes.

The attack is simple: let $|\$\rangle = |\theta_1\rangle \cdots |\theta_n\rangle$ (we omit the classical serial number s, since it plays no role here). Then for each $i \in [n]$, the counterfeiter tries "swapping out" the i^{th} qubit $|\theta_i\rangle$ and replacing it with $|b\rangle$, for each of the four possibilities $|b\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. It then uses $O(\log n)$ queries to the bank, to estimate the probability that the state $|\theta_1\rangle \cdots |\theta_{i-1}\rangle |b\rangle |\theta_{i+1}\rangle \cdots |\theta_n\rangle$ passes the verification test. By doing so, C can learn a correct value of $|\theta_i\rangle$ with success probability $1 - o(1/n)$. The crucial point is that none of these queries damage the qubits *not* being investigated ($|\theta_j\rangle$ for $j \neq i$), since the bank measures those qubits in the correct bases. Therefore C can reuse the same banknote for each query.

More generally, recall from Section 3.1 that we call a private-key quantum money scheme *query-secure*, if it remains secure even assuming the counterfeiter C can make adaptive queries to $\mathsf{Ver}(k, \cdot)$. Then we saw that the Wiesner and BBBW schemes are *not* query-secure. Recently, Farhi et al. [21] proved a much more general "no-go" theorem—which says intuitively that, if we want query-secure quantum money, then the banknotes *must* hide information in the "global correlations" between large numbers of qubits.

On the positive side, any *public-key* quantum money scheme—for example, our multivariate polynomial scheme from Section 6—immediately yields a query-secure scheme with the same security guarantee. This is because a counterfeiter who knows the code of Ver can easily simulate oracle access

[19] See http://theoreticalphysics.stackexchange.com/questions/370/rigorous-security-proof-for-wiesners-quantum-money

to Ver. But can we do any better than that, and construct a query-secure money scheme whose security is *unconditional* (as in Wiesner's scheme), or else based on a pseudorandom function (as in the BBBW scheme)?

In the forthcoming paper [4], we will answer this question in the affirmative, by *directly* adapting the hidden subspace scheme from Section 5 (i.e., the scheme based on a classical oracle). Since the idea is an extremely simple one, let us sketch it here.

Theorem 36 (Query-Secure Variant of Wiesner's Scheme)**.** *Relative to a random oracle R,*[20] *there exists a private-key quantum money scheme, with perfect completeness and $2^{-\Omega(n)}$ soundness error, that is information-theoretically query-secure. One can also replace the random oracle R by a pseudorandom function family $\{f_k\}_k$, to obtain a private-key quantum money scheme, with no oracle, that is query-secure assuming that the f_k's cannot be distinguished from random in quantum polynomial time.*

PROOF SKETCH. For each key k and a serial number s, we will think of the random oracle R as encoding a classical description $R(k,s)$ of a subspace $A_{k,s} \leq \mathbb{F}_2^n$, which is uniformly random subject to $\dim(A_{k,s}) = n/2$. Let $|A_{k,s}\rangle$ be a uniform superposition over $A_{k,s}$. Then the private-key money scheme $\mathcal{S} = (\mathsf{KeyGen}, \mathsf{Bank}, \mathsf{Ver})$ is defined as follows:

- $\mathsf{KeyGen}(0^n)$ generates an n-bit key k uniformly at random.
- $\mathsf{Bank}(k)$ outputs a banknote $|\$_s\rangle := |s\rangle|A_{k,s}\rangle$, for a random serial number $s \in \{0,1\}^n$.
- $\mathsf{Ver}(k,(s,\rho))$ applies a projective measurement that accepts ρ with probability $\langle A_{k,s}|\rho|A_{k,s}\rangle$.

Now, suppose it were possible to break $\mathcal{S}$ (i.e., to counterfeit $|A_{k,s}\rangle$), using $\mathrm{poly}(n)$ adaptive queries to $\mathsf{Ver}(k,\cdot)$. Then we claim that it would *also* be possible to break our *public-key* scheme from Section 5, and thereby contradict the unconditional security proof for the latter! The reason is simply that any query to Ver, of the form $\mathsf{Ver}(k,(s,\rho))$, can easily be simulated using queries to $U_{A_{k,s}}$ and $U_{A_{k,s}^{\perp}}$, the membership oracles for $A_{k,s}$ and $A_{k,s}^{\perp}$ respectively that are available to a counterfeiter against the public-key scheme.

Finally, suppose we replace $R(k,s)$ by a pseudorandom function $f_k(s)$. Then just like with the original BBBW scheme [14], we can argue as follows. Since we already showed that $\mathcal{S}$ is information-theoretically secure when instantiated with a "true" random function, any break of $\mathcal{S}$ in the pseudorandom case would thereby distinguish the function f_k from random. ■

8. OPEN PROBLEMS

The "obvious" problem is to better understand the security of our explicit scheme based on polynomials. Are there nontrivial attacks, for example using Gröbner-basis algorithms? Can we base the security of our scheme—or a related scheme—on some cryptographic assumption that does *not* involve exponentially-small success probabilities? What happens as we change the field size or polynomial degree? Does "hiding" a subspace $A \leq \mathbb{F}_2^n$ in the way we suggest, as the set of common zeroes of multivariate polynomials $p_1, \ldots, p_m : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, have other cryptographic applications, for example to *program obfuscation* [10]?

Of course, there is also tremendous scope for inventing new schemes, which might be based on different assumptions and have different strengths and weaknesses.

Let us move on to some general questions about public-key quantum money. First, is there an unconditionally-secure public-key quantum money scheme relative to a *random* oracle R? (Recall that Wiesner's original scheme [39] was unconditionally-secure and used only a random oracle, but was private-key. Meanwhile, our scheme from Section 5 is unconditionally-secure and public-key, but requires a non-random oracle.) Second, is there a public-key quantum money scheme where the banknotes consist of *single, unentangled qubits*, as in Wiesner's scheme? Note that the results of Farhi et al. [21] imply that, if such a scheme exists, then it cannot be projective. Third, is there a general way to amplify soundness error in quantum money schemes?[21] (We show how to amplify *completeness* error in the full version.)

8.1 Quantum Copy-Protection and More

Quantum money is just *one* novel cryptographic use for the No-Cloning Theorem. Given essentially *any* object of cryptographic interest, one can ask whether quantum mechanics lets us make the object uncloneable. Section 1.4 already discussed one example—*uncloneable signatures*—but there are many others, such as commitments and proofs.[22]

Along those lines, Aaronson [3] proposed a task that, if achievable, would arguably be an even more dramatic application of the No-Cloning Theorem than quantum money: namely, *quantum software copy-protection*. He gave explicit schemes—which have not yet been broken—for copy-protecting a restricted class of functions, namely the *point functions*. In these schemes, given a "password" $s \in \{0,1\}^n$, a software vendor can prepare a quantum state $|\psi_s\rangle$, which allows its holder to *recognize* s: in other words, to decide whether $x = s$ given $x \in \{0,1\}^n$ as input. On the other hand, given $|\psi_s\rangle$, it seems intractable not only to *find* s for oneself, but even to prepare a *second* quantum state with which s can be recognized.

Admittedly, recognizing passwords is an extremely restricted functionality. However, relative to a quantum oracle, Aaronson [3] also described a scheme to quantumly copy-protect *arbitrary* programs, just as well as if the software vendor were able to hand out uncloneable black boxes.[23] In the spirit of this paper, we can now ask: is there likewise a way to quantumly copy-protect arbitrary programs relative to a *classical* oracle? We conjecture that the answer is yes, and in fact we have plausible candidate constructions, which are directly related to the hidden-subspace money scheme of Section 5. However, the security of those constructions seems to hinge on the following conjecture.

Conjecture 37 (Direct Product for Finding Black-Box Subspace Elements)**.** *Let A be a uniformly-random subspace of*

[20] Or alternatively, assuming the bank has access to a giant random number table, as in Wiesner's original setup [39].

[21] Theorem 13 gives *some* soundness amplification for projective schemes: namely, from constant to $1/\mathrm{poly}(n)$. Here we are asking whether one can do anything better.

[22] Even within complexity theory, it would be interesting to study the class QMA (Quantum Merlin-Arthur) subject to the constraint that witnesses must be hard to clone—or alternatively, that witnesses must be *easy* to clone!

[23] As usual, full details have not yet appeared yet.

$\mathbb{F}_2^n$ *satisfying* $\dim(A) = n/2$. *Then given membership oracles for both* A *and* $A^\perp$, *any quantum algorithm needs* $2^{\Omega(n)}$ *queries to find two distinct nonzero elements* $x, y \in A$, *with success probability* $\Omega\left(2^{-n/2}\right)$.

Besides its applications for copy-protection, a proof of Conjecture 37 would be an important piece of formal evidence for Conjecture 32, on which we based the security of our explicit money scheme.

9. ACKNOWLEDGMENTS

We thank Andris Ambainis, Boaz Barak, Dmitry Gavinsky, Daniel Gottesman, Aram Harrow, Yuval Ishai, Shaunak Kishore, Greg Kuperberg, Andy Lutomirski, Abel Molina, Rafi Ostrovsky, Amit Sahai, Peter Shor, and John Watrous for helpful discussions and correspondence; and the anonymous reviewers for their comments.

10. REFERENCES

[1] S. Aaronson. Quantum lower bound for the collision problem. *STOC*, p. 635–642, 2002. quant-ph/0111102.

[2] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. quant-ph/0402095.

[3] S. Aaronson. Quantum copy-protection and quantum money. *IEEE Complexity*, p. 229–242, 2009.

[4] S. Aaronson. On the security of private-key quantum money, 2011. In preparation.

[5] S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. quant-ph/0604056.

[6] M. Albrecht and C. Cid. Cold boot key recovery by solving polynomial systems with noise. In J. Lopez and G. Tsudik, ed., *Applied Cryptography and Network Security 2011*, LNCS vol. 6715, p. 57–72, 2011.

[7] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *SODA*, p. 594–598, 1998.

[8] A. Ambainis. Quantum lower bounds by quantum arguments. *JCSS*, 64:750–767, 2002. quant-ph/0002066.

[9] A. Ambainis, L. Magnin, M. Roetteler, and J. Roland. Symmetry-assisted adversaries for quantum state generation. *IEEE Complexity*, p. 167–177, 2011. arXiv:1012.2112.

[10] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. *CRYPTO*, p. 1–18, 2001. ECCC TR01-057.

[11] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. quant-ph/9802049.

[12] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. quant-ph/9701001.

[13] C. H. Bennett and G. Brassard. Quantum cryptography: public key distribution and coin tossing. *Proc. IEEE Int. Conf. on Computers Systems and Signal Proc.*, p. 175–179, 1984.

[14] C. H. Bennett, G. Brassard, S. Breidbart, and S. Wiesner. Quantum cryptography, or unforgeable subway tokens. *CRYPTO*, p. 267–275, 1982.

[15] N. Bohr. *Atomic Physics and Human Knowledge*. Dover, 2010. First published 1961.

[16] D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. *ASIACRYPT*, p. 41–69, 2011. arXiv:1008.0931.

[17] C. Bouillaguet, J.-C. Faugère, P.-A. Fouque, and L. Perret. Practical cryptanalysis of the identification scheme based on the isomorphism of polynomial with one secret problem. *Public Key Cryptography*, p. 473–493, 2011.

[18] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. In S. J. Lomonaco and H. E. Brandt, ed., *Quantum Computation and Information*, AMS, 2002. quant-ph/0005055.

[19] S. Chakraborty, J. Radhakrishnan, and N. Raghunathan. Bounds for error reduction with few quantum queries. *APPROX-RANDOM*, p. 245–256, 2005.

[20] J. Ding and B.-Y. Yang. Multivariate public key cryptography. In D. J. Bernstein, J. Buchmann, and E. Dahmén, ed., *Post-Quantum Cryptography*, p. 198–242, Springer-Verlag, 2009.

[21] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, D. Nagaj, and P. Shor. Quantum state restoration and single-copy tomography. *Phys. Rev. Lett.*, 105(190503), 2010. arXiv:0912.3823.

[22] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor. Quantum money from knots. *ITCS*, p. 276–289, 2012. arXiv:1004.5127.

[23] D. Gavinsky. Quantum money with classical verification. arXiv:1109.0372, 2011.

[24] W. Geiselmann, W. Meier, and R. Steinwandt. An attack on the isomorphisms of polynomials problem with one secret. *Int. J. Inf. Sec*, 2(1):59–64, 2003.

[25] L. K. Grover. A fast quantum mechanical algorithm for database search. *STOC*, p. 212–219, 1996. quant-ph/9605043.

[26] C. Hillar and L.-H. Lim. Most tensor problems are NP hard, 2009. arXiv:0911.1393.

[27] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. *FOCS*, p. 344–353, 2011. arXiv:1011.3020.

[28] A. Lutomirski. An online attack against Wiesner's quantum money. arXiv:1010.0256, 2010.

[29] A. Lutomirski. Component mixers and a hardness result for counterfeiting quantum money. arXiv:1107.0321, 2011.

[30] A. Lutomirski, S. Aaronson, E. Farhi, D. Gosset, A. Hassidim, J. Kelner, and P. Shor. Breaking and making quantum money: toward a new quantum cryptographic protocol. *ITCS*, p. 20–31, 2010. arXiv:0912.3825.

[31] A. Molina, T. Vidick, and J. Watrous. Optimal counterfeiting attacks and generalizations for Wiesner's quantum money. arXiv:1202.4010, 2012.

[32] M. Mosca and D. Stebila. Quantum coins. In *Error-Correcting Codes, Finite Geometries and Cryptography*, vol. 523, p. 35–47, AMS, 2010. arXiv:0911.1295.

[33] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. *STOC*, p. 33–43, 1989.

[34] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge, 2000.

[35] J. Patarin, L. Goubin, and N. Courtois. Improved algorithms for isomorphisms of polynomials. *EUROCRYPT*, p. 184–200, 1998.

[36] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *STOC*, p. 84–93, 2005.

[37] J. Rompel. One-way functions are necessary and sufficient for secure signatures. *STOC*, p. 387–394, 1990.

[38] T. Tulsi, L. Grover, and A. Patel. A new algorithm for fixed point quantum search. *Quantum Inf. and Comput.*, 6(6):483–494, 2006. quant-ph/0505007.

[39] S. Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983. Original manuscript written circa 1970.