Unconditional Unclonable Encryption

Prabhanjan Ananth* UCSB
Seyoon Ragavan† MIT
Amit Sahai‡ UCLA

Abstract

We give an unconditional construction of information-theoretically secure one-time private-key unclonable encryption for one-bit messages, with efficient encryption and decryption and exponentially small unclonable-indistinguishability advantage. Previous work by Bhattacharyya and Culf (Nature Physics, 2026) and Bhattacharyya, Broadbent, and Culf (arXiv:2603.08916) either only showed inverse-polynomial security loss or required inefficient encryption/decryption operations; we avoid both of these caveats.

Our scheme's key is a uniformly random non-identity phase-free Pauli on n qubits, and bit a is encrypted as a random $(-1)^a$ eigenstate of that Pauli. We prove that the probability that both receivers recover the bit is at most $1/2 + O\left(2^{-n/2}\right)$. By a lower bound due to Broadbent, Culf, and Rochette (IACR Communications in Cryptology, 2026), this is the best probability bound achievable with n-qubit ciphertexts while ensuring perfect correctness (up to the constant hidden in the $O(\cdot)$).

Contents

1 Introduction **2**
1.1 Technical Overview . . . 5

2 Preliminaries **9**
2.1 Paulis and Cliffords . . . 11
2.2 Conditions for PSD 2×2 Block Matrices . . . 13
2.3 Unclonable Encryption and Unclonable Indistinguishability . . . 14

3 Construction **15**
3.1 Explicit Instantiation . . . 16

4 Security **18**
4.1 Part 1: Ricochet Reduction . . . 18
4.2 Part 2: Bounding the Product of the Correlation Operators . . . 20
4.2.1 Proof A: Using Pauli Commutation-Anticommutation Structure . . . 20
4.2.2 Proof B: Using Conditional Overlap and Hilbert–Schmidt Orthogonality . . . 21

*prabhanjan@cs.ucsb.edu
†sragavan@mit.edu
‡sahai@cs.ucla.edu

4.3 Part 3: Accounting for Disagreement 23
4.4 Part 4: Completing the Proof 24

5 Statement on AI Use **26**

6 Acknowledgements **27**

A Glossary of Notation **31**

1 Introduction

Unclonable cryptography is a branch of quantum cryptography that leverages the no-cloning principle of quantum mechanics [WZ82; Die82; BL20] to design primitives that are impossible to achieve classically. No-cloning is merely a hope; it precludes the possibility of an adversary generically creating two copies of a given state, but an unclonable security guarantee must prevent an adversary from creating *any two states* that retain some property of the given state.

Within the rapidly growing landscape of unclonable cryptography, unclonable encryption has emerged as an important primitive. Roughly speaking, unclonable encryption (UE) [Got03; BL20] is an encryption scheme whose ciphertexts cannot be copied in a way that retains information about the plaintext message (even if the decryption key is revealed). In the one-time setting, an adversary receives a single copy of a quantum ciphertext while the key is hidden and then splits it between two non-communicating recipients. The key is then revealed to both the recipients, and the security requires that they cannot both simultaneously recover the hidden message substantially better than a coordinated guess.

In more detail, two notions of security have been studied for unclonable encryption:

- **Search security**: the challenger samples a message m uniformly at random from $\{0,1\}^n$ and sends an encryption of m to $\mathcal{A}$ who then produces a bipartite state shared with $\mathcal{B}$ and $\mathcal{C}$. Upon receiving the decryption key, both $\mathcal{B}$ and $\mathcal{C}$ respectively output $m_{\mathcal{B}}$ and $m_{\mathcal{C}}$. $(\mathcal{A},\mathcal{B},\mathcal{C})$ succeed if $m_{\mathcal{B}} = m$ and $m_{\mathcal{C}} = m$. The work of [BL20] presented a scheme where they showed that the probability of any adversary $(\mathcal{A},\mathcal{B},\mathcal{C})$ succeeding is at most $(\cos^2(\frac{\pi}{8}))^n$.

- **Indistinguishability security**: the challenger samples a bit b at random and then encrypts a message m_b, for two adversarially chosen messages (m_0, m_1), and then sends the ciphertext to $\mathcal{A}$. From here on, the experiment is the same as the search security experiment. Unlike search security, there is always an adversarial strategy that can win with probability $\frac{1}{2}$. For example, $\mathcal{A}$ could forward the ciphertext to $\mathcal{B}$, who can later decrypt and leave $\mathcal{C}$ to guess randomly. Alternatively, $\mathcal{B}$ and $\mathcal{C}$ could just agree to both guess 0 regardless of $\mathcal{A}$'s message. So the natural goal is to design an UE scheme satisfying indistinguishability security where the success probability is as close to $\frac{1}{2}$ as possible.

Prior constructions. Over the past few years, many works have studied indistinguishability security for unclonable encryption. We use indUE to denote unclonable encryption satisfying indistinguishability security.

- **Plain model**: [AKY25] achieved indUE where the decryption keys are quantum states. Recently, there has been progress in designing indUE with classical keys. [Bot+26] proposed a

construction of indUE with n-qubit ciphertexts, conjectured an inverse-polynomial advantage bound of $O(1/\sqrt{n})$, proved this bound for $n \leq 3$, and provided numerical evidence for this bound for $4 \leq n \leq 8$. This was improved by [BC26], who proposed a different construction of indUE for which they proved an inverse-polynomial advantage bound. A follow up work by [BBC26b] presents a construction with the desired negligible advantage bound; however, their construction is not efficient, meaning that the encryption and the decryption algorithms do not run in quantum polynomial time[1].

- **Idealized models**: [Ana+22; AKL23] presented indUE with optimal advantage in the quantum random oracle model. [BG26b] also presented optimal indUE with stronger security guarantees albeit in the Haar random oracle model.

Despite the recent progress, the goal of achieving efficient indUE in the plain model with classical keys and negligible adversary advantage has remained elusive. Although indUE could plausibly exist information-theoretically, we do not even have a construction that leans on plain-model cryptographic assumptions. Moreover, a few works present evidence highlighting the challenges in achieving indistinguishability security. Notably [MST21; Ana+22] showed that indUE with deterministic encryption cannot satisfy negligible security. [CLX26a] also highlighted the difficulties in using BB84 states to achieve indUE.

Our contribution. We address the main goal of achieving unclonable encryption satisfying indistinguishability security. We show:

Theorem 1.1 (Informal)**.** *For every $n \geq 1$, there exists a one-time private-key quantum encryption scheme for one-bit messages with a $2n$-bit classical key and an n-qubit ciphertext. The scheme has perfect correctness, uses only single-qubit Clifford gates for encryption and local Pauli measurements for decryption, and satisfies*

$$\Pr[b_{\mathcal{B}} = b_{\mathcal{C}} = b] \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^n}{4^n - 1}}$$

against every information-theoretic pre-key splitting attack in the one-bit unclonable-indistinguishability experiment.

By a lower bound of Broadbent, Culf, and Rochette [BCR26, Corollary 6], this is the best bound achievable with n-qubit ciphertexts and perfect correctness, up to the constant factor in the $2^{-n/2}$ advantage.

Implications. Quantum state indistinguishability obfuscation (qsiO) [CG24] takes a quantum implementation of a classical function—a quantum state together with an evaluation circuit—and produces another implementation of the same function. Its security guarantee is that obfuscations of any two implementations of the same size and functionality are computationally indistinguishable, even to quantum adversaries.
As a consequence, our unconditional construction of indUE yields the following applications of qsiO. Both applications assume qsiO and post-quantum injective one-way functions; our result eliminates the additional assumption that unclonable encryption exists.

[1] Concretely, their scheme has advantage negligible in n, where encryption and decryption can be implemented in time exponential in n.

- **Copy-protection for search-puncturable programs.** Coladangelo and Gunn [CG24, Theorem 18] show that qsiO copy-protects any family of *search-puncturable* programs, assuming injective one-way functions and unclonable encryption. Informally, such a program can be punctured at an input x while preserving its behavior everywhere else, but the punctured description does not enable an efficient adversary to produce a valid output at x. Combining their reduction with our result therefore gives copy-protection for these families from qsiO and post-quantum injective one-way functions.[2] Given one protected program, no efficient adversary can produce two separated evaluators that both return valid outputs on the same uniformly random challenge input, except with negligible probability.
- **Unclonable puncturable obfuscation.** Ananth and Behera [AB24, Section 6, Corollary 52] show that qsiO, together with injective one-way functions and unclonable-indistinguishable encryption of bits, yields unclonable puncturable obfuscation (UPO). Our construction supplies the encryption primitive unconditionally, giving UPO for polynomial-size puncturable keyed circuits from qsiO and post-quantum injective one-way functions. The guarantee is generalized UPO security for *identical, uniformly random challenges*: after an adversary splits an obfuscation between two non-communicating parties, both parties receive the same random puncture point and must distinguish the original circuit from a circuit whose behavior is modified at that point. Their joint success probability is at most $1/2 + \mathrm{negl}(\lambda)$.

Subsequent work. Following the initial version of this work, Bhattacharyya, Broadbent, and Culf [BBC26a] extended efficient, information-theoretically secure unclonable encryption from a single bit to messages of arbitrary length. Their construction exploits the Clifford structure of the encoding and achieves statistical security in the one-time setting, with encryption and decryption polynomial in the message length and security parameter. Thus, the extension to longer messages can also be achieved without computational assumptions.

Ananth [Ana26] studied a general non-local search-to-decision question over $\mathbb{F}_2$. For a bipartite encoding of a uniformly random string, he showed that a non-negligible advantage in simultaneously predicting its inner product with a common random mask implies a non-negligible probability of simultaneously recovering the entire string by local measurements. This result is information-theoretic and does not provide an efficient extractor. Coladangelo, Liu, and Xie [CLX26b] established a simultaneous Goldreich–Levin reduction for two entangled parties receiving the same random mask. Their reduction generically upgrades search security to unclonable indistinguishability and, in particular, proves security of the natural BB84-based construction with a random inner-product mask. They also give an efficient implementation when a lower bound on the prediction advantage is known.

Concurrent work. Concurrently and independently, Botteron, Designolle, and Fawzi [BDF26] studied a general family of *Pauli encodings*, in which ciphertexts are normalized eigenspace projectors of Pauli strings. Their framework includes the family of all non-identity tensor Paulis considered here. Their work provides a complementary analysis of the possibilities and limitations of Pauli encodings.

[2] Note that our result is unclonable encryption for single bits, whereas the Coladangelo–Gunn reduction requires unclonable encryption for arbitrary-length messages. This can be generically handled via the length-boosting results by either [AB24, Corollary 53] or [MH25] combined with [BG26a]. Alternatively, one could use the follow-up by [BBC26a] on our work which achieves unclonable encryption for arbitrary-length messages without any computational assumptions.

Lean 4 formalization. To help with verifiability, we release Lean 4 code for our main proofs. In more detail, the Lean repository proves all correctness and security claims in Theorem 4.1, but it does not formally show that encryption and decryption are efficient. It also provides a modular interface defining unclonable encryption for a bit, allowing for other users to provide their own constructions and security proofs in Lean 4. It also does not prove the aforementioned implications for other unclonable cryptography primitives since that would require formalizing the results of [AB24] and [CG24]. The Lean code was built using Codex GPT-5.6 Sol and GPT-6 Astra and can be found at this GitHub repository.

1.1 Technical Overview

From BB84 parity to Pauli parity. It is useful to start with the natural BB84-based attempt. For each qubit, the encryptor chooses either the X or the Z observable and prepares one of its two eigenstates. The eigenvalue bits are sampled uniformly subject to their XOR being the message. Once the basis string is revealed, the decryptor measures each qubit and takes the parity of the outcomes. The surprising XOR-repetition result of Coladangelo, Liu, and Xie [CLX26a] exhibits an attack on this construction that has constant winning probability $\cos^2(\pi/8)$.

Our construction makes a very simple tweak on this: instead of choosing between X and Z for each qubit, we choose among the full set of Paulis $\{I, X, Y, Z\}$. In more detail, the key is chosen uniformly from all $4^n - 1$ non-identity Paulis

$$P_{x,z} = \bigotimes_{j=1}^{n} P_{x_j,z_j}, \qquad P_{x_j,z_j} = i^{x_j z_j} X^{x_j} Z^{z_j} \in \{I, X, Y, Z\}.$$

To encrypt a bit m, prepare a uniformly random product eigenstate whose eigenvalue under $P_{x,z}$ is $(-1)^m$. Concretely, choose the eigenvalue signs on the non-identity factors uniformly subject to their product being $(-1)^m$, and prepare a maximally mixed qubit on every identity factor. The resulting ciphertext is the maximally mixed state on the $(-1)^m$ eigenspace:

$$\rho_{\mathsf{ct}}^{(k,m)} = \frac{I + (-1)^m P_{x,z}}{2^n}.$$

Decryption measures the local Paulis and multiplies the signs. The product is exactly $(-1)^m$. Thus both encryption and decryption use only local operations; the challenge is to prove security.

Our construction is quite close to that conceived by [Bot+26]; they take their key space to be a collection of pairwise anticommuting Paulis, whereas we take the full collection of non-identity Paulis. We show that this modification of the construction of [Bot+26] satisfies indistinguishability security with negligible advantage. A modification abandoning pairwise anticommutation is necessary, as the pairwise-anticommuting constraint imposed by [Bot+26] automatically limits the number of keys to at most $\mathsf{poly}(n)$.

Before continuing, we briefly foreshadow that our construction does not actually rely on the key being sampled from the full set of $4^n - 1$ non-identity Paulis. We could take any nonempty subset $\mathcal{K}$ of the non-identity Paulis, and our techniques will show a security bound of $\frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^n}{|\mathcal{K}|}}$. Thus any key space of size $2^{n+\omega(\log n)}$ would suffice for security; the BB84 construction falls short of this because it comprises exactly 2^n keys.

The operators and their meaning. The Choi–Jamiołkowski representation lets us describe the cloner's channel by a fixed state τ on a reference register $\mathbf{A}$ and the receivers' registers $\mathbf{B}, \mathbf{C}$. This is the usual ricochet reduction used in unclonable encryption [BL20]. For each key $k = (x, z)$, let B_k and C_k be Bob's and Charlie's decoding observables: outcome 0 has sign $+1$, and outcome 1 has sign -1. By a standard Naimark dilation argument, we may assume without loss of generality that B_k and C_k are unitary reflections. Writing $K = 4^n - 1$ for the number of keys, define

$$\begin{aligned}
\mathsf{E}_B &= \frac{1}{K}\sum_k P_{x,z}^T \otimes B_k \otimes I_{\mathbf{C}}, \\
\mathsf{E}_C &= \frac{1}{K}\sum_k P_{x,z}^T \otimes I_{\mathbf{B}} \otimes C_k, \\
\mathsf{E}_{BC} &= \frac{1}{K}\sum_k I_{\mathbf{A}} \otimes B_k \otimes C_k, \qquad \mathsf{D} = I_{\mathbf{ABC}} - \mathsf{E}_{BC}.
\end{aligned}$$

These operators have simple interpretations:

$$\begin{aligned}
\mathrm{Tr}(\mathsf{E}_B\tau) &= 2\mathsf{Pr}[\text{Bob guesses correctly}] - 1, \\
\mathrm{Tr}(\mathsf{E}_C\tau) &= 2\mathsf{Pr}[\text{Charlie guesses correctly}] - 1, \\
\mathrm{Tr}(\mathsf{D}\tau) &= 2\mathsf{Pr}[\text{Bob and Charlie disagree}].
\end{aligned}$$

In particular, D measures disagreement, with a factor of two in our normalization.

Expanding the success probability gives

$$\mathsf{Pr}[\mathsf{Win}] = \frac{1}{2}\left(1 + \mathrm{Tr}(\mathbf{G}\tau)\right) \le \frac{1}{2}\left(1 + \lambda_{\max}(\mathbf{G})\right), \qquad \mathbf{G} = \frac{1}{2}(\mathsf{E}_B + \mathsf{E}_C - \mathsf{D}). \tag{1}$$

The subtraction of D has a concrete meaning: the individual success probabilities count cases where only one receiver is correct, and in exactly those cases the receivers disagree. Our task is to show that the largest eigenvalue of $\mathbf{G}$ is small. We do not need to bound its negative eigenvalues: for example, if Bob always answers 0 and Charlie always answers 1, then $\mathbf{G} = -I$ and their joint success probability is zero.

Why the bounds of [Tom+13] and SoS techniques appear insufficient. Up to this point, the manipulations to analyze the UE security game have been standard [BL20], modulo our particular choice of UE construction. At this point, we have reduced our task to bounding the top eigenvalue of some Hermitian operator; the two main pre-existing techniques for proceeding from here are a bound by [Tom+13] and SoS/SDP techniques [NPA08] which were recently leveraged by [Bot+26]. We therefore take this opportunity to highlight the limitations of these pre-existing approaches for our setting, before continuing with the overview of our proof.

First, the bound by [Tom+13] is central to their analysis of the BB84 monogamy-of-entanglement game and the construction by [BL20] of search-secure unclonable encryption. They show, for any Hermitian projectors $\Pi_1, \ldots, \Pi_K$ acting on the same Hilbert space, that:

$$\left\| \frac{1}{K}\sum_{i=1}^{K} \Pi_i \right\|_\infty \le \frac{1}{K}\sum_{j\in[K]} \max_{i\in[K]} \left\| \Pi_i \Pi_{\pi_j(i)} \right\|_\infty, \tag{2}$$

where $\pi_1, \ldots, \pi_K$ are any permutations on $[K]$ such that $\pi_j(i) \neq \pi_{j'}(i)$ whenever $j \neq j'$. To cast our goal as bounding the operator norm of a sum of projectors, note that we can write:

$$\frac{1}{2}(I + \mathbf{G}) = \frac{1}{K} \sum_{k=(x,z)\in\mathcal{K}} \sum_{b\in\{0,1\}} \left(\frac{I + (-1)^b P_{x,z}^T}{2}\right) \otimes \left(\frac{I + (-1)^b B_k}{2}\right) \otimes \left(\frac{I + (-1)^b C_k}{2}\right). \quad (3)$$

Any attempt to bound the norm of this above by $1/2 + \mathrm{negl}(n)$ using (2) will be "bottlenecked" by the norms of products of Pauli projectors; if we were to take $B_k = C_k = I$ for all k, then we have:

$$\frac{1}{2}(I + \mathbf{G}) = \frac{1}{K} \sum_{k=(x,z)\in\mathcal{K}} \left(\frac{I + P_{x,z}^T}{2}\right) \otimes I \otimes I. \quad (4)$$

Now one can check:

$$\left\| \left(\frac{I + P_{x,z}^T}{2}\right) \left(\frac{I + P_{x',z'}^T}{2}\right) \right\|_\infty = \begin{cases} 1, \text{ if } P_{x,z}P_{x',z'} = P_{x',z'}P_{x,z}; \\ \frac{1}{\sqrt{2}}, \text{ if } P_{x,z}P_{x',z'} = -P_{x',z'}P_{x,z}. \end{cases}$$

Either way, this bound is well above $1/2$; in fact, the above illustrates that using (2) can only hope to get us to a bound of $1/\sqrt{2}$ (whereas we want $1/2 + \mathrm{negl}(n)$).

A second approach for analyzing unclonable encryption is based on the NPA sum-of-squares hierarchy for nonlocal games [NPA08]; this was recently applied by [Bot+26] for small values of n. Looking ahead, our particular bounding argument does not admit a very direct translation as a constant-degree sum-of-squares proof because it leverages an infinite power series expansion of a matrix inverse (see (5) which we will introduce soon, and its Taylor expansion in (36) in the body). We emphasize that this intuition does not rule out the possibility of a constant-degree SoS analysis of our scheme obtained via a different route.

A bound from the Pauli structure. The first estimate is

$$\|\mathsf{E}_B \mathsf{E}_C\|_\infty \leq \frac{d}{K}, \qquad d = 2^n.$$

We give two proofs. One uses the signs that arise when we commute two Paulis past each other. These signs form a submatrix of a scaled, column-permuted Hadamard matrix, whose norm controls the sum defining $\mathsf{E}_B\mathsf{E}_C$. The other uses Hilbert–Schmidt orthogonality: after tracing out the reference register, the terms from distinct Paulis vanish. Both proofs apply to arbitrary decoding observables.

It is tempting to interpret this product bound as saying that Bob and Charlie cannot both have much advantage individually. But a product of operator expectations need not equal the expectation of the product. Indeed, consider a cloner who flips a fair coin and forwards the ciphertext to one receiver, leaving the other to guess randomly. Each receiver succeeds with probability $3/4$, even though their joint success probability is only $1/2$. Their errors occur on disjoint events. To bound joint success, we must account for disagreement as well.

Accounting for disagreement. We use the positive operator

$$\mathsf{F}(t) = 2t(2tI + \mathsf{D})^{-1}, \qquad t > 0, \tag{5}$$

where t is a parameter we will select later. On an eigenvector of D with eigenvalue $a \geq 0$, this operator acts by the scalar $2t/(2t + a)$. It therefore gives more weight to directions with less disagreement, which are the directions relevant to a successful attack. The term $2tI$ makes the inverse well-defined even when D has a kernel.

The product bound survives this insertion:

$$\|\mathsf{E}_B\mathsf{F}(t)\mathsf{E}_C\|_\infty \leq \frac{d}{K} \qquad \text{for every } t > 0. \tag{6}$$

To see why, expand the inverse as a convergent geometric series in E_{BC}. Each term involves $\mathsf{E}_B\mathsf{E}_{BC}^{\ell}\mathsf{E}_C$. Operators on different receivers' registers commute, so we can move the inserted Charlie operators to the left and the inserted Bob operators to the right. They are all contractions, and hence each term has norm at most $\|\mathsf{E}_B\mathsf{E}_C\|_\infty$. The coefficients of the series are positive and sum to one.

Completing the argument. The specific choice of F is in order to enable the last step of the proof. The remaining analytical task is to convert a bound on the spectral norm of a *product* of matrices that looks roughly like $\mathsf{E}_B\mathsf{D}^{-1}\mathsf{E}_C$ to a bound on the top eigenvalue of a *sum* of matrices, $\mathsf{E}_B + \mathsf{E}_C - \mathsf{D}$. The solution to this is inspired by the scalar identity:

$$(x + y)^2 = (x - y)^2 + 4xy. \tag{7}$$

In our setting, the relevant matrix identity is:

$$(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B + \mathsf{E}_C) = (\mathsf{E}_B - \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B - \mathsf{E}_C) + 2\mathsf{E}_B\mathsf{F}(t)\mathsf{E}_C + 2\mathsf{E}_C\mathsf{F}(t)\mathsf{E}_B. \tag{8}$$

We will show a bound on the LHS by bounding the RHS. The RHS has three terms, the second and third of which are handled by (6). For the first term on the RHS, we show in Proposition 4.11 that:

$$(\mathsf{E}_B - \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B - \mathsf{E}_C) \preceq 2t\mathsf{D}. \tag{9}$$

The intuition for the above is that $\mathsf{E}_B - \mathsf{E}_C, \mathsf{E}_C - \mathsf{E}_B \preceq \mathsf{D}$; if exactly one of Bob and Charlie succeeds, then they must disagree. Plugging this and (6) into (8) implies that:

$$(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B + \mathsf{E}_C) \preceq 2t\mathsf{D} + \frac{4d}{K}I. \tag{10}$$

It remains to show that (10) implies that $\lambda_{\max}(\mathbf{G}) \leq \mathrm{negl}(n)$. To this end, we will now set $t = \lambda_{\max}(\mathbf{G})$ and assume $t > 0$ (otherwise we are done). Choose a unit eigenvector $|\psi\rangle$ with $\mathbf{G}|\psi\rangle = t|\psi\rangle$. Rearranging this equation gives

$$(\mathsf{E}_B + \mathsf{E}_C)|\psi\rangle = (2tI + \mathsf{D})|\psi\rangle.$$

Thus the inverse in $\mathsf{F}(t)$ cancels the operator on the right. It follows that

$$\langle\psi|(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B + \mathsf{E}_C)|\psi\rangle = 4t^2 + 2t\langle\psi|\mathsf{D}|\psi\rangle. \tag{11}$$

Therefore (10) tells us that:

$$\begin{aligned} 4t^2 + 2t\langle\psi|\mathsf{D}|\psi\rangle &\le \langle\psi|2t\mathsf{D} + \frac{4d}{K}I|\psi\rangle \\ &= 2t\langle\psi|\mathsf{D}|\psi\rangle + \frac{4d}{K} \\ \Leftrightarrow t &\le \sqrt{d/K}. \end{aligned}$$

Together with (1), this bounds the winning probability by $\frac{1}{2} + \frac{1}{2}\sqrt{d/K} = \frac{1}{2} + \frac{1}{2}\sqrt{2^n/(4^n-1)}$, completing the proof.

2 Preliminaries

The statistical security parameter is denoted by $n \in \mathbb{N}$. We write $x \xleftarrow{\$} S$ for uniform sampling from a finite set S. A function $\mu : \mathbb{N} \to \mathbb{R}_{\ge 0}$ is *negligible* if, for every $c > 0$, there is an N_c such that $\mu(n) < 1/n^c$ for every $n \ge N_c$.

A register $\mathbf{R}$ is associated with a finite-dimensional Hilbert space, also denoted $\mathbf{R}$. We write $\mathcal{L}(\mathbf{R})$ for the linear operators on $\mathbf{R}$ and $\mathcal{D}(\mathbf{R})$ for the density operators on $\mathbf{R}$. An operator M is positive/positive semidefinite, written $M \succeq 0$, if $\langle\psi|M|\psi\rangle \ge 0$ for every vector $|\psi\rangle$. For self-adjoint (or Hermitian[3]) operators M, N, the Loewner order $M \preceq N$ means that $N - M \succeq 0$.

A density operator $\rho \in \mathcal{D}(\mathbf{R})$ satisfies $\rho \succeq 0$ and $\mathrm{Tr}(\rho) = 1$. A pure state $|\psi\rangle$ is identified with the rank-one density operator $|\psi\rangle\langle\psi|$. If $\rho_{\mathbf{RS}}$ is a bipartite state, its reduced state on $\mathbf{R}$ is $\rho_{\mathbf{R}} = \mathrm{Tr}_{\mathbf{S}}(\rho_{\mathbf{RS}})$. All partial traces in this note are unnormalized as linear maps on operators.

For an operator M, the operator norm and trace norm are

$$\|M\|_\infty := \sup_{\||\psi\rangle\|=1} \||M|\psi\rangle\|, \qquad \|M\|_1 := \mathrm{Tr}\sqrt{M^\dagger M}.$$

We will repeatedly use the equivalent characterization

$$\|M\|_\infty = \sup_{\substack{\||\xi\rangle\|=1 \\ \||\eta\rangle\|=1}} |\langle\eta|M|\xi\rangle|.$$

The operator norm is submultiplicative: $\|MN\|_\infty \le \|M\|_\infty\|N\|_\infty$. If $M = M^\dagger$, then $\|M\|_\infty$ is the largest absolute value of an eigenvalue of M. We also let $\lambda_{\max}(M)$ denote the largest eigenvalue of M (without taking absolute values). Consequently, for every density operator ρ,

$$\mathrm{Tr}(M\rho) \le \lambda_{\max}(M). \tag{12}$$

A quantum channel $\Phi : \mathcal{L}(\mathbf{R}) \to \mathcal{L}(\mathbf{S})$ is a completely positive trace-preserving linear map. We abuse notation and also use $I_{\mathbf{A}} : \mathcal{L}(\mathbf{A}) \to \mathcal{L}(\mathbf{A})$ to denote the identity channel. A measurement with outcome set $\mathcal{X}$ is a POVM $\{M_x\}_{x\in\mathcal{X}}$ satisfying $M_x \succeq 0$ and $\sum_x M_x = I$. On input state ρ, the Born rule gives $\Pr[x] = \mathrm{Tr}(M_x\rho)$.

A binary POVM $\{M_0, M_1\}$ can equivalently be represented by the self-adjoint contraction

$$B := M_0 - M_1, \qquad -I \preceq B \preceq I.$$

[3]We are only concerned with finite-dimensional Hilbert spaces.

Conversely,

$$M_0 = \frac{I + B}{2}, \qquad M_1 = \frac{I - B}{2}.$$

If outcome 0 is assigned value $+1$ and outcome 1 is assigned value -1, then $\mathrm{Tr}(B\rho) = \mathsf{Pr}[0] - \mathsf{Pr}[1]$ is the signed bias of the outcome.

We also use the following standard facts:

(i) every self-adjoint operator has an orthonormal eigenbasis;

(ii) if $M \succeq rI$ for some $r > 0$, then M is invertible and $0 \prec M^{-1} \preceq r^{-1}I$;

(iii) if two positive operators commute, then their product is positive;

(iv) $\mathrm{Tr}(MN) = \mathrm{Tr}(NM)$ whenever the products are defined, and trace is multiplicative over tensor products;

(v) the ordinary vector Cauchy–Schwarz inequality is $|\langle u \mid v\rangle| \le \|u\|\,\|v\|$.

These conventions are standard; see, for example, Watrous [Wat18]. The normalized Choi–Jamiołkowski identity below is standard [Jam72; Cho75]; we record the normalization explicitly because conventions differ by factors of the input dimension.

Choi–Jamiołkowski identity. The following identity expresses a channel's action in terms of its normalized Choi state:

Theorem 2.1. *Let $\Phi_{\mathrm{E}\to\mathrm{BC}} : \mathcal{L}(\mathrm{E}) \to \mathcal{L}(\mathrm{B}\otimes\mathrm{C})$ be a quantum channel, and fix a copy $\mathrm{A} \simeq \mathrm{E}$ with the corresponding basis. For $M_{\mathrm{A}} \in \mathcal{L}(\mathrm{A})$ and $N_{\mathrm{BC}} \in \mathcal{L}(\mathrm{B}\otimes\mathrm{C})$, use the same matrix M on either A or E. Define the maximally entangled state*

$$|\Omega\rangle_{\mathrm{AE}} = \frac{1}{\sqrt{\dim(\mathrm{A})}} \sum_{j\in[\dim(\mathrm{A})]} |j\rangle_{\mathrm{A}}|j\rangle_{\mathrm{E}}.$$

Then

$$\frac{1}{\dim(\mathrm{A})}\,\mathrm{Tr}\left(N_{\mathrm{BC}}\Phi_{\mathrm{E}\to\mathrm{BC}}(M_{\mathrm{E}})\right) = \mathrm{Tr}\left[(M_{\mathrm{A}}^T \otimes N_{\mathrm{BC}})(I_{\mathrm{A}} \otimes \Phi_{\mathrm{E}\to\mathrm{BC}})(|\Omega\rangle\langle\Omega|_{\mathrm{AE}})\right].$$

Proof. Suppose the $(j,k)^{th}$ entry of M is $\alpha_{j,k}$. Thus, the $(k,j)^{th}$ entry of M^T is $\alpha_{j,k}$.

$$\begin{aligned}
&\mathrm{Tr}\Big((M_{\mathrm{A}}^T \otimes N_{\mathrm{BC}})(I_{\mathrm{A}} \otimes \Phi_{\mathrm{E}\to\mathrm{BC}})(|\Omega\rangle\langle\Omega|_{\mathrm{AE}})\Big)\\
&= \frac{1}{\dim(\mathrm{A})} \sum_{j,k\in[\dim(\mathrm{A})]} \mathrm{Tr}\Big((M_{\mathrm{A}}^T \otimes N_{\mathrm{BC}})(I_{\mathrm{A}} \otimes \Phi_{\mathrm{E}\to\mathrm{BC}})(|j\rangle\langle k|_{\mathrm{A}} \otimes |j\rangle\langle k|_{\mathrm{E}})\Big)\\
&= \frac{1}{\dim(\mathrm{A})} \sum_{j,k\in[\dim(\mathrm{A})]} \mathrm{Tr}\Big(M_{\mathrm{A}}^T\,|j\rangle\langle k|_{\mathrm{A}}\Big) \cdot \mathrm{Tr}(N_{\mathrm{BC}}\Phi_{\mathrm{E}\to\mathrm{BC}}(|j\rangle\langle k|_{\mathrm{E}}))\\
&= \frac{1}{\dim(\mathrm{A})} \sum_{j,k\in[\dim(\mathrm{A})]} \alpha_{j,k} \cdot \mathrm{Tr}(N_{\mathrm{BC}}\Phi_{\mathrm{E}\to\mathrm{BC}}(|j\rangle\langle k|_{\mathrm{E}}))
\end{aligned}$$

$$= \frac{1}{\dim(\mathbf{A})} \cdot \mathrm{Tr}\left(N_{\mathbf{BC}} \Phi_{\mathbf{E}\to\mathbf{BC}} \left(\sum_{j,k\in[\dim(\mathbf{A})]} \alpha_{j,k} |j\rangle\langle k|_{\mathbf{E}} \right) \right)$$

$$= \frac{1}{\dim(\mathbf{A})} \cdot \mathrm{Tr}(N_{\mathbf{BC}} \Phi_{\mathbf{E}\to\mathbf{BC}} (M_{\mathbf{E}}))$$

□

2.1 Paulis and Cliffords

We describe the single-qubit Pauli matrices below:

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The phase gate is

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix},$$

and the Hadamard gate is

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

We state a basic fact below.

Fact 2.2. *Let P be a single-qubit Pauli and let $|\psi\rangle$ be a single-qubit state such that $P|\psi\rangle = \theta|\psi\rangle$, where $\theta \in \{+1, -1\}$. Then measuring $|\psi\rangle$ according to the observable P – i.e., measuring $|\psi\rangle$ according to the two-outcome projective measurement $\left\{ \underbrace{\frac{I+P}{2}}_{\text{outcome } +1}, \underbrace{\frac{I-P}{2}}_{\text{outcome } -1} \right\}$ – yields the outcome θ with probability 1.*

Notation. We use the following notation. For strings $x \in \{0,1\}^n, z \in \{0,1\}^n$,

$$P_{x,z} = \bigotimes_{j\in[n]} P_{x_j,z_j},$$

where $P_{x_j,z_j} = i^{x_j z_j} X^{x_j} Z^{z_j}$. In other words,

$$P_{x_j,z_j} = \begin{cases} I & x_j = z_j = 0, \\ X & x_j = 1, z_j = 0, \\ Z & x_j = 0, z_j = 1, \\ Y & x_j = 1, z_j = 1 \end{cases}$$

We write $d := 2^n$, $P_k := P_{x,z}$ for $k = (x, z)$, and $X^x := \bigotimes_j X^{x_j}$, $Z^z := \bigotimes_j Z^{z_j}$.

Identities. The following are well-known identities:

1. $HXH = Z$
2. $(SH)^\dagger Y(SH) = Z$
3. $X^2 = Y^2 = Z^2 = I$
4. $\mathrm{Tr}(X) = \mathrm{Tr}(Y) = \mathrm{Tr}(Z) = 0$
5. $\mathrm{Tr}(XY) = \mathrm{Tr}(YX) = \mathrm{Tr}(XZ) = \mathrm{Tr}(ZX) = \mathrm{Tr}(YZ) = \mathrm{Tr}(ZY) = 0$

Hilbert–Schmidt orthogonality.

Proposition 2.3 (Orthogonality of Pauli operators)**.** *For every $k = (x, z) \in \mathbb{F}_2^n \times \mathbb{F}_2^n$, $P_{x,z}$ is Hermitian and $P_{x,z}^2 = I$. If $k \neq (0^n, 0^n)$, then $P_{x,z}$ also has trace zero. Moreover, for $k = (x, z)$ and $k' = (x', z')$,*

$$\mathrm{Tr}(P_{x,z} P_{x',z'}) = d \cdot \mathbb{1}[k = k']. \tag{13}$$

The same statements hold for the transposed Paulis $\{P_{x,z}^T : (x, z) \in \mathbb{F}_2^n \times \mathbb{F}_2^n\}$.

Proof. For one qubit, the four matrices $i^{xz} X^x Z^z$ are I, Z, X, Y for $(x, z) = (0, 0), (0, 1), (1, 0), (1, 1)$, respectively. Each is Hermitian and squares to the identity. Tensor products preserve these properties, so $P_{x,z} = P_{x,z}^\dagger$ and $P_{x,z}^2 = I$.

If $k \neq (0^n, 0^n)$, at least one local factor is X, Y, or Z, and therefore has trace zero. Trace is multiplicative over tensor products, so $\mathrm{Tr}(P_{x,z}) = 0$ for every non-identity Pauli. The one-qubit Paulis are Hilbert–Schmidt orthogonal: $\mathrm{Tr}(PP') = 2\mathbb{1}[P = P']$. Hence

$$\begin{aligned} \mathrm{Tr}(P_{x,z} P_{x',z'}) &= \prod_{j=1}^{n} \mathrm{Tr}\Big(P_{x_j,z_j} P_{x'_j,z'_j}\Big) \\ &= 2^n \mathbb{1}[(x, z) = (x', z')] = d\mathbb{1}[k = k']. \end{aligned}$$

Every tensor Pauli satisfies $P_{x,z}^T = \pm P_{x,z}$, so $P_{x,z}^T$ is also Hermitian, and traceless when $k \neq (0^n, 0^n)$. Finally,

$$\mathrm{Tr}\Big(P_{x,z}^T P_{x',z'}^T\Big) = \mathrm{Tr}\Big((P_{x',z'} P_{x,z})^T\Big) = \mathrm{Tr}(P_{x',z'} P_{x,z}),$$

which proves the orthogonality statement for the transposed Paulis. □

Commutation-anticommutation structure.

Definition 2.4 (Pauli symplectic form)**.** *For $k = (x, z), k' = (x', z') \in \mathbb{F}_2^n \times \mathbb{F}_2^n$, we define the symplectic form to be:*

$$[k, k'] = \langle x, z' \rangle + \langle z, x' \rangle \in \mathbb{F}_2.$$

Lemma 2.5 (Pauli commutation relations)**.** *For any $k, k' \in \mathbb{F}_2^n \times \mathbb{F}_2^n$, we have:*

$$P_k P_{k'} = (-1)^{[k,k']} P_{k'} P_k. \tag{14}$$

The same relation holds for the transposed Paulis.

Proof. Write $k = (x, z)$ and $k' = (x', z')$. Additionally, let $|x \wedge z| \in \mathbb{N}$ denote the number of entries where x and z are both 1. Then we have:

$$P_k P_{k'} = i^{|x\wedge z|+|x'\wedge z'|} X^x Z^z X^{x'} Z^{z'} = i^{|x\wedge z|+|x'\wedge z'|} \cdot (-1)^{\langle x', z\rangle} X^{x+x'} Z^{z+z'}.$$

By symmetry, we also have:

$$P_{k'} P_k = i^{|x\wedge z|+|x'\wedge z'|} \cdot (-1)^{\langle x, z'\rangle} X^{x+x'} Z^{z+z'}.$$

Comparing the above two equalities implies (14). Transposing (14) reverses both products and preserves the sign, which gives the same relation for P_k^T. □

Definition 2.6 (Symplectic character matrix)**.** *Define $\Xi^{\text{Pauli}} \in \{-1, 1\}^{4^n \times 4^n}$ to be the matrix with rows and columns indexed by elements of $\mathbb{F}_2^n \times \mathbb{F}_2^n$. The entry in row k and column k' is $\Xi^{\text{Pauli}}_{k,k'} = (-1)^{[k,k']}$.*

Lemma 2.7 (Pauli symplectic structure is spectrally balanced)**.** *We have $\left\|\Xi^{\text{Pauli}}\right\|_\infty = 2^n$.*

Proof. For $k = (x, z)$ and $k' = (x', z')$, the corresponding entry is

$$\Xi^{\text{Pauli}}_{(x,z),(x',z')} = (-1)^{\langle x,z'\rangle+\langle z,x'\rangle}.$$

Permute the columns by exchanging their two coordinates. The resulting matrix Ξ' has entries

$$\Xi'_{(x,z),(x',z')} = (-1)^{\langle x,x'\rangle+\langle z,z'\rangle}.$$

Thus $2^{-n}\Xi'$ is the normalized $2^{2n} \times 2^{2n}$ Hadamard matrix, which is unitary. Since column permutations preserve operator norm,

$$\|\Xi^{\text{Pauli}}\|_\infty = \|\Xi'\|_\infty = 2^n.$$

□

2.2 Conditions for PSD 2×2 Block Matrices

The following is a well-known property of Schur complements; regardless, we include its proof for completeness.

Lemma 2.8 (PSD 2×2 block matrices)**.** *Let A, B, C be square Hermitian operators of the same dimension. Then:*

1. *If C is positive-definite, then $\begin{bmatrix} A & B \\ B & C \end{bmatrix}$ is PSD if and only if $A - BC^{-1}B$ is PSD.*
2. *The block matrix $\begin{bmatrix} A & B \\ B & A \end{bmatrix}$ is PSD if and only if $A - B$ and $A + B$ are both PSD.*

Proof. To prove Item 1, note that if we let $T = \begin{bmatrix} I & 0 \\ C^{-1}B & I \end{bmatrix}$, we have:

$$T^\dagger \begin{bmatrix} A - BC^{-1}B & 0 \\ 0 & C \end{bmatrix} T = \begin{bmatrix} I & BC^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} A - BC^{-1}B & 0 \\ 0 & C \end{bmatrix} \begin{bmatrix} I & 0 \\ C^{-1}B & I \end{bmatrix}$$

$$= \begin{bmatrix} I & BC^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} A - BC^{-1}B & 0 \\ B & C \end{bmatrix}$$
$$= \begin{bmatrix} A & B \\ B & C \end{bmatrix}.$$

Since T is invertible, it follows that $\begin{bmatrix} A & B \\ B & C \end{bmatrix}$ being PSD is equivalent to $\begin{bmatrix} A - BC^{-1}B & 0 \\ 0 & C \end{bmatrix}$ being PSD, which is in turn equivalent to $A - BC^{-1}B$ and C both being PSD. We are already assuming that C is positive-definite, so this is in turn equivalent to $A - BC^{-1}B$ being PSD.

Now we turn to Item 2. Note that $U = \frac{1}{\sqrt{2}} \begin{bmatrix} I & I \\ I & -I \end{bmatrix}$ is invertible. Conjugating by U, we have:

$$U \begin{bmatrix} A & B \\ B & A \end{bmatrix} U^\dagger = \frac{1}{2} \begin{bmatrix} I & I \\ I & -I \end{bmatrix} \begin{bmatrix} A & B \\ B & A \end{bmatrix} \begin{bmatrix} I & I \\ I & -I \end{bmatrix}$$
$$= \frac{1}{2} \begin{bmatrix} I & I \\ I & -I \end{bmatrix} \begin{bmatrix} A+B & A-B \\ A+B & B-A \end{bmatrix}$$
$$= \begin{bmatrix} A+B & 0 \\ 0 & A-B \end{bmatrix}.$$

So $\begin{bmatrix} A & B \\ B & A \end{bmatrix}$ is PSD if and only if $\begin{bmatrix} A+B & 0 \\ 0 & A-B \end{bmatrix}$ is PSD, which is in turn true if and only if $A+B$ and $A-B$ are both PSD. □

2.3 Unclonable Encryption and Unclonable Indistinguishability

We recall the definition of unclonable encryption below; our formulation follows the unclonable IND-CPA definition of Ananth, Kaleoglu, Li, Liu, and Zhandry [Ana+22], building on Broadbent–Lord and Ananth–Kaleoglu [BL20; AK21].

Definition 2.9 (One-time private-key unclonable encryption)**.** *A one-time private-key unclonable encryption scheme is a triple of quantum algorithms $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\mathcal{M}_n$:*

- $\mathsf{Gen}(1^n)$ *outputs a classical secret key k;*
- $\mathsf{Enc}(k, m)$*, for $m \in \mathcal{M}_n$, outputs a quantum ciphertext state ρ_{ct};*
- $\mathsf{Dec}(k, \rho_{\mathsf{ct}})$ *outputs a classical message.*

The scheme is correct if there is a negligible function μ such that, for every $m \in \mathcal{M}_n$,

$$\Pr\left[\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m : k \xleftarrow{\$} \mathsf{Gen}(1^n)\right] \geq 1 - \mu(n).$$

Perfect correctness means that the probability is 1 for every supported key and every message.

Definition 2.10 (One-time unclonable-indistinguishability game)**.** *Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a private-key quantum encryption scheme. An adversary is a tuple $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$, where $\mathcal{G}$ chooses two equal-length challenge messages, $\mathcal{A}$ is the pre-key splitting algorithm, and $\mathcal{B}, \mathcal{C}$ are the non-communicating post-split algorithms. The experiment proceeds as follows:*

1. $(m_0, m_1, \rho_S) \leftarrow \mathcal{G}(1^n)$, *with* $m_0, m_1 \in \mathcal{M}_n$;
2. *Sample* $k \leftarrow \mathsf{Gen}(1^n)$ *and* $b \xleftarrow{\$} \{0,1\}$;
3. *Prepare* $\rho_{\mathsf{ct}} \leftarrow \mathsf{Enc}(k, m_b)$ *and give* $\rho_{\mathsf{ct}} \otimes \rho_S$ *to* $\mathcal{A}$;
4. $\mathcal{A}$ *applies a channel and outputs a bipartite state* ρ_{BC};
5. *Reveal the same key* k *to* $\mathcal{B}$ *and* $\mathcal{C}$*; they output bits* $b_{\mathcal{B}}$ *and* $b_{\mathcal{C}}$ *without communicating;*
6. $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$ *wins exactly when* $b_{\mathcal{B}} = b_{\mathcal{C}} = b$.

We denote the winning probability by $\Pr\left[\mathrm{Win}^{\mathsf{UI}}_{\Pi,\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(n)\right]$. *For the one-bit message space* $\{0,1\}$, *we fix* $(m_0, m_1) = (0,1)$, *absorb any auxiliary-state preparation into* $\mathcal{A}$, *and write* $\Pr\left[\mathrm{Win}^{\mathsf{UI}}_{\Pi,\mathcal{A},\mathcal{B},\mathcal{C}}(n)\right]$ *for the winning probability in this specialized game.*

Definition 2.11 (Unclonable indistinguishability)**.** *The scheme* Π *is information-theoretically* $\varepsilon(n)$*-unclonable-indistinguishable if, for every finite-dimensional adversary* $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$,

$$\Pr\left[\mathrm{Win}^{\mathsf{UI}}_{\Pi,\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(n)\right] \leq \frac{1}{2} + \varepsilon(n).$$

It is computationally secure if the same inequality holds for every QPT *adversary with* $\varepsilon = \mathrm{negl}$. *The baseline* $1/2$ *is achieved by sending one shared random guess to both recipients. We denote the advantage of* $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$ *to be* $\Pr\left[\mathrm{Win}^{\mathsf{UI}}_{\Pi,\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(n)\right] - \frac{1}{2}$.

Remark 2.12. *For the message space* $\{0,1\}$, *the fixed-message game above suffices to establish security: equal challenge messages give winning probability at most* $1/2$, *while the pair* $(1,0)$ *reduces to* $(0,1)$ *by complementing both receivers' output bits. Conditioning on the chosen message pair also covers randomized message selection and its associated auxiliary state. Hence, throughout the construction and security proof, we consider adversaries of the form* $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ *in this specialized game.*

3 Construction

We describe a construction of an unclonable encryption scheme for 1-bit messages as follows. We first describe the ciphertext as a mixed state, then give explicit implementations in Section 3.1.

- $\underline{\mathsf{Gen}(1^n)}$: On input a statistical security parameter $n \geq 1$, it samples and outputs

$$k = (x, z) \xleftarrow{\$} (\{0,1\}^n \times \{0,1\}^n) \setminus \{(0^n, 0^n)\}.$$

- $\underline{\mathsf{Enc}(k, m \in \{0,1\})}$: On input a key k parsed as (x, z) and message $m \in \{0,1\}$, prepare the mixed state:

$$\rho^{(k,m)}_{\mathsf{ct}} := \frac{1}{2^n}(I + (-1)^m P_{x,z}).$$

  Note that the above is a well-defined state since $(x, z) \neq (0^n, 0^n)$.

- $\mathsf{Dec}(k, \sigma)$: On input the secret key k parsed as (x, z) and an n-qubit ciphertext state σ, do the following: measure σ using the $P_{x,z}$ observable. That is, measure σ using the two-outcome projective measurement $\mathcal{M} = \left\{ \underbrace{\frac{I + P_{x,z}}{2}}_{\text{outcome } +1}, \underbrace{\frac{I - P_{x,z}}{2}}_{\text{outcome } -1} \right\}$ to obtain the outcome θ. If θ is $+1$ then output 0, else output 1.

Lemma 3.1. *The above scheme satisfies correctness.*

Proof. For $m' \in \{0,1\}$, we have:

$$\begin{aligned}
\Pr\left[\mathsf{Dec}(k, \rho_{\mathsf{ct}}^{(k,m)}) = m'\right] &= \mathrm{Tr}\left[\left(\frac{I + (-1)^{m'} P_{x,z}}{2}\right) \rho_{\mathsf{ct}}^{(k,m)}\right] \\
&= \mathrm{Tr}\left[\left(\frac{I + (-1)^{m'} P_{x,z}}{2}\right)\left(\frac{I + (-1)^{m} P_{x,z}}{2^n}\right)\right] \\
&= \frac{1}{2^{n+1}}\left(\underbrace{\mathrm{Tr}\left[I\right]}_{=2^n} + \left((-1)^m + (-1)^{m'}\right)\underbrace{\mathrm{Tr}\left[P_{x,z}\right]}_{=0} + (-1)^{m+m'}\underbrace{\mathrm{Tr}\left[I\right]}_{=2^n}\right) \\
&= \frac{1}{2}(1 + (-1)^{m+m'}),
\end{aligned}$$

which is 1 if $m' = m$ and 0 otherwise. □

3.1 Explicit Instantiation

Encryption. To prepare the mixed state $\rho_{\mathsf{ct}}^{(k,m)}$, do the following:

- For every $j \in [n]$, set U_{x_j,z_j} as:

$$U_{x_j,z_j} = \begin{cases} I & x_j = 0 \\ H & x_j = 1, z_j = 0 \\ SH & x_j = 1, z_j = 1 \end{cases}$$

  Set $U_k = U_{x_1,z_1} \otimes \cdots \otimes U_{x_n,z_n}$.

- By the definition of Gen, there exists a minimal $j^* \in [n]$ for which $P_{x_{j^*},z_{j^*}} \neq I$ (equivalently, $(x_{j^*}, z_{j^*}) \neq (0,0)$). For each $j \neq j^*$, sample $r_j \xleftarrow{\$} \{0,1\}$. Then set r_{j^*} as follows:

$$r_{j^*} = m \bigoplus \left(\bigoplus_{\substack{j \in [n] \\ j \neq j^*}} q_j \cdot r_j \right),$$

  where q_j is defined as follows:

$$q_j = \begin{cases} 0 & P_{x_j,z_j} = I, \\ 1 & P_{x_j,z_j} \neq I. \end{cases}$$

  (Equivalently, we are sampling $(r_1, \ldots, r_n)$ uniformly subject to the constraint that $\bigoplus_{j \in [n]} q_j \cdot r_j = m$.)

- Output the n-qubit pure state $|\psi_{k,m;r}\rangle = U_k|r_1 \cdots r_n\rangle$.

Let us check that this indeed prepares the mixed state $\rho_{\mathsf{ct}}^{(k,m)}$:

Proposition 3.2. *We have:*

$$\rho_{\mathsf{ct}}^{(k,m)} = \frac{1}{2^n}(I + (-1)^m P_{x,z}) = \frac{1}{2^{n-1}} \sum_{\substack{r\in\{0,1\}^n \\ \text{s.t. } m=\bigoplus_{j=1}^n q_j r_j}} U_k \, |r_1 \cdots r_n\rangle\langle r_1 \cdots r_n| \, U_k^\dagger. \tag{15}$$

Proof. The vectors $|r_1 \cdots r_n\rangle$ with $\bigoplus_j q_j r_j = m$ form an orthonormal basis of the $(-1)^m$ eigenspace of $Z^{q_1} \otimes \cdots \otimes Z^{q_n}$. Consequently,

$$\sum_{\substack{r\in\{0,1\}^n \\ \bigoplus_j q_j r_j = m}} |r_1 \cdots r_n\rangle\langle r_1 \cdots r_n| = \frac{1}{2}\left(I + (-1)^m \bigotimes_{j=1}^n Z^{q_j}\right). \tag{16}$$

At least one q_j is 1, so this eigenspace has dimension 2^{n-1}. The preparation procedure samples its basis vectors uniformly. For each qubit, $U_{x_j,z_j} Z^{q_j} U_{x_j,z_j}^\dagger = P_{x_j,z_j}$, and therefore

$$U_k \left(\bigotimes_{j=1}^n Z^{q_j}\right) U_k^\dagger = P_{x,z}.$$

Conjugating (16) by U_k and dividing by 2^{n-1} gives

$$\frac{1}{2^{n-1}} \sum_{\substack{r\in\{0,1\}^n \\ \bigoplus_j q_j r_j = m}} U_k \, |r_1 \cdots r_n\rangle\langle r_1 \cdots r_n| \, U_k^\dagger = \frac{1}{2^n}(I + (-1)^m P_{x,z}) = \rho_{\mathsf{ct}}^{(k,m)}.$$

This is the claimed ciphertext state. □

Decryption. To measure the observable $P_{x,z}$, we use the fact that this observable is in tensor product across the n qubits and accordingly do the following:

- For every $j \in [n]$, measure the j^{th} qubit using the P_{x_j,z_j} observable. That is, measure the j^{th} qubit using the two-outcome projective measurement $\mathcal{M}_j = \left\{ \underbrace{\frac{I + P_{x_j,z_j}}{2}}_{\text{outcome } +1}, \underbrace{\frac{I - P_{x_j,z_j}}{2}}_{\text{outcome } -1} \right\}$ to obtain the outcome θ_j.
- Compute $\theta = \prod_{j=1}^n \theta_j \in \{+1, -1\}$. If $\theta = +1$ then output 0, else output 1. (When $P_{x_j,z_j} = I$, the outcome is $\theta_j = +1$ with certainty, so qubit j need not be measured.)

4 Security

Theorem 4.1. *Let* $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$*, where* Gen, Enc *and* Dec *are as defined above. For every* $n \geq 1$ *and every adversary* $(\mathcal{A}, \mathcal{B}, \mathcal{C})$*, we have:*

$$\Pr\left[\mathrm{Win}^{\mathsf{UI}}_{\Pi,\mathcal{A},\mathcal{B},\mathcal{C}}(n)\right] \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^n}{4^n - 1}}, \tag{17}$$

where we use the one-bit specialization of Definition 2.10, as justified in Remark 2.12.

The proof has four parts. First, we express the winning probability in terms of the largest eigenvalue of an operator. Next, we use the Pauli structure to bound the product of Bob's and Charlie's correlation operators. We then show that the bound still holds after inserting the operator that accounts for disagreement. Finally, we use this estimate to bound the largest eigenvalue and hence the winning probability.

4.1 Part 1: Ricochet Reduction

Fix an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$. By Remark 2.12, we may take the challenge messages to be 0 and 1. Let $\Phi_{\mathrm{E}\to\mathrm{BC}} : \mathcal{L}(\mathrm{E}) \to \mathcal{L}(\mathrm{B} \otimes \mathrm{C})$ be the cloner's channel, where E is the ciphertext register and B, C are sent to Bob and Charlie. Any auxiliary state prepared before receiving the ciphertext can be included in the definition of Φ. For each key k, write the receivers' binary POVMs as $\{\Lambda^{(k)}_{\mathcal{B},0}, \Lambda^{(k)}_{\mathcal{B},1}\}$ and $\{\Lambda^{(k)}_{\mathcal{C},0}, \Lambda^{(k)}_{\mathcal{C},1}\}$, and set

$$B_k = \Lambda^{(k)}_{\mathcal{B},0} - \Lambda^{(k)}_{\mathcal{B},1}, \qquad C_k = \Lambda^{(k)}_{\mathcal{C},0} - \Lambda^{(k)}_{\mathcal{C},1}.$$

Thus their measurement operators for a guess m are $(I_{\mathrm{B}} + (-1)^m B_k)/2$ and $(I_{\mathrm{C}} + (-1)^m C_k)/2$. The observables B_k, C_k are Hermitian contractions, and the winning probability is

$$\begin{aligned} p &= \Pr[(\mathcal{A}, \mathcal{B}, \mathcal{C}) \text{ wins}] \\ &= \mathbb{E}_{k=(x,z),m} \operatorname{Tr}\left[\left(\frac{I_{\mathrm{B}} + (-1)^m B_k}{2} \otimes \frac{I_{\mathrm{C}} + (-1)^m C_k}{2}\right) \Phi_{\mathrm{E}\to\mathrm{BC}}(\rho^{(k,m)}_{\mathsf{ct}})\right]. \end{aligned} \tag{18}$$

The averages are over the sampled key and a uniform bit m.

Let A be a copy of E, and let

$$\tau_{\mathrm{ABC}} = (I_{\mathrm{A}} \otimes \Phi_{\mathrm{E}\to\mathrm{BC}})(|\Omega\rangle\langle\Omega|_{\mathrm{AE}})$$

be the normalized Choi state. Applying Theorem 2.1 and the ciphertext formula (15) to (18) gives

$$p = \mathbb{E}_{k=(x,z),m} \operatorname{Tr}\left[\left((I + (-1)^m P^T_{x,z})_{\mathrm{A}} \otimes \frac{I_{\mathrm{B}} + (-1)^m B_k}{2} \otimes \frac{I_{\mathrm{C}} + (-1)^m C_k}{2}\right) \tau_{\mathrm{ABC}}\right]. \tag{19}$$

The terms with an odd power of $(-1)^m$ vanish on averaging over m. The four remaining terms give

$$\begin{aligned} p = \frac{1}{4}\mathbb{E}_{k=(x,z)} \operatorname{Tr}\Big[&(I_{\mathrm{ABC}} + (P^T_{x,z})_{\mathrm{A}} \otimes B_k \otimes I_{\mathrm{C}} \\ &+ (P^T_{x,z})_{\mathrm{A}} \otimes I_{\mathrm{B}} \otimes C_k + I_{\mathrm{A}} \otimes B_k \otimes C_k)\,\tau_{\mathrm{ABC}}\Big]. \end{aligned} \tag{20}$$

Define

$$\mathbf{G} = \frac{1}{2}\mathbb{E}_{k=(x,z)}\Big[(P_{x,z}^T)_{\mathbf{A}} \otimes B_k \otimes I_{\mathbf{C}} + (P_{x,z}^T)_{\mathbf{A}} \otimes I_{\mathbf{B}} \otimes C_k + I_{\mathbf{A}} \otimes B_k \otimes C_k - I_{\mathbf{ABC}}\Big]. \tag{21}$$

Since τ has trace one, we obtain

$$p = \frac{1}{2}\left(1 + \mathrm{Tr}(\mathbf{G}\tau)\right) \leq \frac{1}{2} + \frac{1}{2}\lambda_{\max}(\mathbf{G}), \tag{22}$$

where the inequality is (12). It remains to bound the largest eigenvalue of $\mathbf{G}$.

Definition 4.2 (Keys and decoding observables)**.** *For the rest of the proof, let*

$$\mathcal{K}_n := (\{0,1\}^n \times \{0,1\}^n) \setminus \{(0^n, 0^n)\},$$
$$d := 2^n, \qquad K := |\mathcal{K}_n| = 4^n - 1 = d^2 - 1.$$

For $k = (x,z) \in \mathcal{K}_n$, write Θ_k for the transposed Pauli on the reference register. Recall the decoding observables from the reduction:

$$\Theta_k := (P_{x,z}^T)_{\mathbf{A}}, \qquad B_k \in \mathcal{L}(\mathbf{B}), \qquad C_k \in \mathcal{L}(\mathbf{C}). \tag{23}$$

Since B_k, C_k arise from binary POVMs, they are self-adjoint contractions:

$$B_k = B_k^\dagger, \quad C_k = C_k^\dagger, \quad -I_{\mathbf{B}} \preceq B_k \preceq I_{\mathbf{B}}, \quad -I_{\mathbf{C}} \preceq C_k \preceq I_{\mathbf{C}}. \tag{24}$$

Definition 4.3 (Averaged operators)**.** *Define the following operators on $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$:*

$$\mathsf{E}_B := \frac{1}{K}\sum_{k\in\mathcal{K}_n} \Theta_k \otimes B_k \otimes I_{\mathbf{C}}, \tag{25}$$

$$\mathsf{E}_C := \frac{1}{K}\sum_{k\in\mathcal{K}_n} \Theta_k \otimes I_{\mathbf{B}} \otimes C_k, \tag{26}$$

$$\mathsf{E}_{BC} := \frac{1}{K}\sum_{k\in\mathcal{K}_n} I_{\mathbf{A}} \otimes B_k \otimes C_k. \tag{27}$$

$$\mathsf{D} := I_{\mathbf{ABC}} - \mathsf{E}_{BC}. \tag{28}$$

The operator $\mathbf{G}$ defined in Equation (21) can equivalently be written as

$$\mathbf{G} = \frac{1}{2}\left(\mathsf{E}_B + \mathsf{E}_C + \mathsf{E}_{BC} - I_{\mathbf{ABC}}\right). \tag{29}$$

Each of $\mathsf{E}_B, \mathsf{E}_C, \mathsf{E}_{BC}$ is an average of Hermitian contractions, so it is itself a Hermitian contraction. In particular, $0 \preceq \mathsf{D} \preceq 2I$. For the Choi state τ of the fixed attack,

$$\mathrm{Tr}(\mathsf{E}_B\tau) = 2\mathsf{Pr}[\textit{Bob guesses correctly}] - 1,$$
$$\mathrm{Tr}(\mathsf{E}_C\tau) = 2\mathsf{Pr}[\textit{Charlie guesses correctly}] - 1,$$
$$\mathrm{Tr}(\mathsf{E}_{BC}\tau) = \mathsf{Pr}[\textit{Bob and Charlie agree}] - \mathsf{Pr}[\textit{Bob and Charlie disagree}],$$
$$\mathrm{Tr}(\mathsf{D}\tau) = 2\mathsf{Pr}[\textit{Bob and Charlie disagree}].$$

These identities follow by assigning signs $+1, -1$ to the binary measurement outcomes, as in the reduction above.

4.2 Part 2: Bounding the Product of the Correlation Operators

We now use the Pauli structure to show that the product of the two correlation operators has small norm.

Lemma 4.4 (Product of the correlation operators)**.** *Let $\gamma := d/K$. Then*

$$\|\mathsf{E}_B\mathsf{E}_C\|_\infty \le \gamma. \tag{30}$$

We present two proofs; the first is adapted from [Rag26] and the other is from [AS26].

4.2.1 Proof A: Using Pauli Commutation-Anticommutation Structure

Proof A of Lemma 4.4. We will make use of the matrix Ξ^{Pauli} defined in Definition 2.6. Recalling the definitions from Definition 4.3, we compute:

$$\begin{aligned}
\mathsf{E}_B\mathsf{E}_C &= \frac{1}{K^2}\sum_{k,k'\in\mathcal{K}_n} (\Theta_k \otimes B_k \otimes I_{\mathbf{C}})\,(\Theta_{k'} \otimes I_{\mathbf{B}} \otimes C_{k'}) \\
&= \frac{1}{K^2}\sum_{k,k'\in\mathcal{K}_n} \Theta_k\Theta_{k'} \otimes B_k \otimes C_{k'} \\
&= \frac{1}{K^2}\sum_{k,k'\in\mathcal{K}_n} \Xi^{\text{Pauli}}_{k',k}\Theta_{k'}\Theta_k \otimes B_k \otimes C_{k'} \\
&= \frac{1}{K^2}\sum_{k,k'\in\mathcal{K}_n} \Xi^{\text{Pauli}}_{k',k}\,(\Theta_{k'} \otimes I_{\mathbf{B}} \otimes C_{k'})\,(\Theta_k \otimes B_k \otimes I_{\mathbf{C}})\,.
\end{aligned}$$

To help rewrite the above expression, let $\tilde{d}$ be the dimension of the Hilbert space of $\mathbf{ABC}$, and let $M_C \in \mathbb{C}^{\tilde{d}\times\tilde{d}|\mathcal{K}_n|}$ be a block matrix with $|\mathcal{K}_n|$ block columns indexed by $k' \in \mathcal{K}_n$. The block column corresponding to k' will be $\Theta_{k'} \otimes I_{\mathbf{B}} \otimes C_{k'}$. Similarly define $M_B \in \mathbb{C}^{\tilde{d}|\mathcal{K}_n|\times\tilde{d}}$ to have row blocks indexed by $k \in \mathcal{K}_n$ and containing $\Theta_k \otimes B_k \otimes I_{\mathbf{C}}$ in the block corresponding to k. Finally, let $M_\Xi \in \mathbb{C}^{\tilde{d}|\mathcal{K}_n|\times\tilde{d}|\mathcal{K}_n|}$ be defined to be:

$$M_\Xi := \Xi^{\text{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n} \otimes I_{\tilde{d}}.$$

In words, M_Ξ has block rows and columns indexed by $k', k \in \mathcal{K}_n$ respectively. The corresponding $\tilde{d}\times\tilde{d}$ block is $\Xi^{\text{Pauli}}_{k',k}$ times the identity matrix. The restriction $\Xi^{\text{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n}$ deletes the row and column indexed by the identity Pauli. In these terms, the outcome of the above computation can be rewritten as:

$$\begin{aligned}
\mathsf{E}_B\mathsf{E}_C &= \frac{1}{K^2} M_C M_\Xi M_B \\
&= \frac{1}{K^2} M_C(\Xi^{\text{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n} \otimes I_{\tilde{d}})M_B \\
\Rightarrow \|\mathsf{E}_B\mathsf{E}_C\|_\infty &= \frac{1}{K^2}\left\|M_C(\Xi^{\text{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n} \otimes I_{\tilde{d}})M_B\right\|_\infty \\
&\le \frac{1}{K^2}\|M_C\|_\infty \cdot \left\|\Xi^{\text{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n}\right\|_\infty \cdot \|M_B\|_\infty.
\end{aligned}$$

We bound each term of the RHS above as follows:

- For any unit $|\psi\rangle$, the vector $M_B|\psi\rangle$ will be the concatenation of $(\Theta_k \otimes B_k \otimes I_{\mathrm{C}})\,|\psi\rangle$ for every $k \in \mathcal{K}_n$. Since Θ_k is unitary and B_k a contraction, each of these chunks has norm at most 1, so the total ℓ_2 norm will be at most $\sqrt{|\mathcal{K}_n|} = \sqrt{K}$. Consequently, we have $\|M_B\|_\infty \le \sqrt{K}$.
- Applying the same argument to $M_C^\dagger$, whose row blocks are $\Theta_{k'} \otimes I_{\mathrm{B}} \otimes C_{k'}$, gives $\|M_C\|_\infty = \|M_C^\dagger\|_\infty \le \sqrt{K}$.
- Finally, we have $\left\|\Xi^{\mathrm{Pauli}}_{\mathcal{K}_n,\mathcal{K}_n}\right\|_\infty \le \left\|\Xi^{\mathrm{Pauli}}\right\|_\infty$ since taking a submatrix cannot increase operator norm. The latter quantity is exactly $2^n = d$ by Lemma 2.7.

Putting these together yields:

$$\|\mathsf{E}_B\mathsf{E}_C\|_\infty \le \frac{1}{K^2} \cdot K \cdot d = \frac{d}{K},$$

as desired. □

4.2.2 Proof B: Using Conditional Overlap and Hilbert–Schmidt Orthogonality

The second proof uses the following inequality for operators that share one register. It bounds their product using partial traces over that register.

Lemma 4.5 (Conditional overlap)**.** *Let $U_{\mathbf{AB}} \in \mathcal{L}(\mathbf{A} \otimes \mathbf{B})$ and $V_{\mathbf{AC}} \in \mathcal{L}(\mathbf{A} \otimes \mathbf{C})$. Extend each operator to $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$ by the identity on its missing register, and denote these extensions by $\widetilde{U}_{\mathbf{ABC}}$ and $\widetilde{V}_{\mathbf{ABC}}$. Then*

$$\|\widetilde{U}_{\mathbf{ABC}}\widetilde{V}_{\mathbf{ABC}}\|_\infty \le \sqrt{\|\operatorname{Tr}_{\mathbf{A}}(U_{\mathbf{AB}}^\dagger U_{\mathbf{AB}})\|_\infty}\sqrt{\|\operatorname{Tr}_{\mathbf{A}}(V_{\mathbf{AC}}V_{\mathbf{AC}}^\dagger)\|_\infty}.$$

Proof. Write $d_{\mathbf{A}} = \dim(\mathbf{A})$ and decompose the operators into blocks:

$$U_{\mathbf{AB}} = \sum_{i,j\in[d_{\mathbf{A}}]} |i\rangle\langle j|_{\mathbf{A}} \otimes (U_{i,j})_{\mathbf{B}}, \qquad V_{\mathbf{AC}} = \sum_{j,k\in[d_{\mathbf{A}}]} |j\rangle\langle k|_{\mathbf{A}} \otimes (V_{j,k})_{\mathbf{C}}.$$

The extensions are therefore

$$\widetilde{U}_{\mathbf{ABC}} = \sum_{i,j} |i\rangle\langle j|_{\mathbf{A}} \otimes (U_{i,j})_{\mathbf{B}} \otimes I_{\mathbf{C}},$$
$$\widetilde{V}_{\mathbf{ABC}} = \sum_{j,k} |j\rangle\langle k|_{\mathbf{A}} \otimes I_{\mathbf{B}} \otimes (V_{j,k})_{\mathbf{C}}.$$

For each j, k, set

$$T_{j,k}^{\mathbf{AB}} = \sum_i |i\rangle\langle k|_{\mathbf{A}} \otimes (U_{i,j})_{\mathbf{B}}.$$

Multiplying the blocks gives

$$\widetilde{U}_{\mathbf{ABC}}\widetilde{V}_{\mathbf{ABC}} = \sum_{j,k} T_{j,k}^{\mathbf{AB}} \otimes (V_{j,k})_{\mathbf{C}}.$$

Fix unit vectors $|\xi\rangle, |\eta\rangle$ on $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$. Since $T_{j,k}^{\mathbf{AB}}$ and $(V_{j,k})_{\mathbf{C}}$ act on disjoint registers, Cauchy–Schwarz gives

$$\left|\langle\eta|\widetilde{U}_{\mathbf{ABC}}\widetilde{V}_{\mathbf{ABC}}|\xi\rangle\right|^2$$

$$= \left| \sum_{j,k} \langle \eta | (I_{\mathbf{AB}} \otimes (V_{j,k})_{\mathbf{C}}) (T_{j,k}^{\mathbf{AB}} \otimes I_{\mathbf{C}}) | \xi \rangle \right|^2$$
$$\leq \left(\sum_{j,k} \| (I_{\mathbf{AB}} \otimes (V_{j,k})_{\mathbf{C}}^\dagger) | \eta \rangle \|^2 \right) \left(\sum_{j,k} \| (T_{j,k}^{\mathbf{AB}} \otimes I_{\mathbf{C}}) | \xi \rangle \|^2 \right).$$

The first sum is bounded by the second partial trace in the statement:

$$\begin{aligned} \sum_{j,k} \| (I_{\mathbf{AB}} \otimes (V_{j,k})_{\mathbf{C}}^\dagger) | \eta \rangle \|^2 &= \langle \eta | \left(I_{\mathbf{AB}} \otimes \sum_{j,k} (V_{j,k})_{\mathbf{C}} (V_{j,k})_{\mathbf{C}}^\dagger \right) | \eta \rangle \\ &= \langle \eta | \left(I_{\mathbf{AB}} \otimes \mathrm{Tr}_{\mathbf{A}} (V_{\mathbf{AC}} V_{\mathbf{AC}}^\dagger) \right) | \eta \rangle \\ &\leq \| \mathrm{Tr}_{\mathbf{A}} (V_{\mathbf{AC}} V_{\mathbf{AC}}^\dagger) \|_\infty. \end{aligned}$$

For the other sum, direct multiplication gives

$$\begin{aligned} \sum_{j,k} (T_{j,k}^{\mathbf{AB}})^\dagger T_{j,k}^{\mathbf{AB}} &= \sum_{i,j,k} |k\rangle\langle k|_{\mathbf{A}} \otimes (U_{i,j})_{\mathbf{B}}^\dagger (U_{i,j})_{\mathbf{B}} \\ &= I_{\mathbf{A}} \otimes \mathrm{Tr}_{\mathbf{A}} (U_{\mathbf{AB}}^\dagger U_{\mathbf{AB}}). \end{aligned}$$

Hence

$$\sum_{j,k} \| (T_{j,k}^{\mathbf{AB}} \otimes I_{\mathbf{C}}) | \xi \rangle \|^2 \leq \| \mathrm{Tr}_{\mathbf{A}} (U_{\mathbf{AB}}^\dagger U_{\mathbf{AB}}) \|_\infty.$$

Combining the two bounds and taking the supremum over unit $|\xi\rangle, |\eta\rangle$ proves the lemma. □

Proof B of Lemma 4.4. To apply the lemma, write the two correlation operators on their respective registers:

$$U_{\mathbf{AB}} := \frac{1}{K} \sum_{k \in \mathcal{K}_n} \Theta_k \otimes B_k, \qquad V_{\mathbf{AC}} := \frac{1}{K} \sum_{k \in \mathcal{K}_n} \Theta_k \otimes C_k.$$

Their canonical embeddings in $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$ are exactly E_B and E_C.

Using Equation (13) and preserving the order of the Bob operators,

$$\begin{aligned} \mathrm{Tr}_{\mathbf{A}} \left(U_{\mathbf{AB}}^\dagger U_{\mathbf{AB}} \right) &= \frac{1}{K^2} \sum_{k,k'} \mathrm{Tr}(\Theta_k \Theta_{k'}) B_k B_{k'} \\ &= \frac{d}{K^2} \sum_k B_k^2 \preceq \frac{d}{K} I_{\mathbf{B}}. \end{aligned}$$

Here $B_k^2 \preceq I_{\mathbf{B}}$ because B_k is a Hermitian contraction.

Similarly,

$$\mathrm{Tr}_{\mathbf{A}} \left(V_{\mathbf{AC}} V_{\mathbf{AC}}^\dagger \right) = \frac{d}{K^2} \sum_k C_k^2 \preceq \frac{d}{K} I_{\mathbf{C}}.$$

Applying the conditional-overlap lemma (Lemma 4.5) proved above gives

$$\| \mathsf{E}_B \mathsf{E}_C \|_\infty \leq \sqrt{\frac{d}{K}} \sqrt{\frac{d}{K}} = \frac{d}{K}.$$

□

4.3 Part 3: Accounting for Disagreement

The product bound alone does not control joint success: the receivers' individual advantages may come from different events, as in the random-forwarding example in the overview. We therefore insert an operator that gives more weight to directions with less disagreement. The next proposition lets us do this using a power series in E_{BC}.

Proposition 4.6 (Inserting powers of the agreement operator)**.** *For every integer $\ell \geq 0$,*

$$\|\mathsf{E}_B \mathsf{E}_{BC}^{\ell} \mathsf{E}_C\|_\infty \leq \|\mathsf{E}_B \mathsf{E}_C\|_\infty \leq \gamma. \tag{31}$$

Proof. For $\ell \geq 1$, expand the inserted power and move only operators on disjoint registers:

$$\begin{aligned} \mathsf{E}_B \mathsf{E}_{BC}^{\ell} \mathsf{E}_C = \frac{1}{K^\ell} \sum_{k_1,\ldots,k_\ell \in \mathcal{K}_n} & \left(I_{\mathbf{A}} \otimes I_{\mathbf{B}} \otimes C_{k_1} \cdots C_{k_\ell}\right) \mathsf{E}_B \mathsf{E}_C \\ & \cdot \left(I_{\mathbf{A}} \otimes B_{k_1} \cdots B_{k_\ell} \otimes I_{\mathbf{C}}\right). \end{aligned} \tag{32}$$

The order of the Bob operators, and separately of the Charlie operators, is unchanged. Each outside product has norm at most one, since every B_k, C_k is a contraction. The triangle inequality therefore gives

$$\|\mathsf{E}_B \mathsf{E}_{BC}^{\ell} \mathsf{E}_C\|_\infty \leq \frac{1}{K^\ell} \sum_{k_1,\ldots,k_\ell \in \mathcal{K}_n} \|\mathsf{E}_B \mathsf{E}_C\|_\infty = \|\mathsf{E}_B \mathsf{E}_C\|_\infty \leq \gamma.$$

The case $\ell = 0$ is exactly the product bound of Lemma 4.4. □

Definition 4.7 (The filter F)**.** *Let $t > 0$ be a real number to be chosen later. Recall that D is positive semidefinite and measures disagreement. Write its spectral decomposition as*

$$\mathsf{D} = \sum_j d_j \left|v_j\right\rangle\!\left\langle v_j\right|, \qquad d_j \geq 0.$$

Define

$$\mathsf{F}(t) := 2t(2tI + \mathsf{D})^{-1} = \sum_j \frac{2t}{2t + d_j} \left|v_j\right\rangle\!\left\langle v_j\right|. \tag{33}$$

Write $\mathsf{F} = \mathsf{F}(t)$ when t is fixed. Its eigenvalues $2t/(2t + d_j)$ lie in $(0, 1]$ and decrease as d_j increases. Thus F gives more weight to directions with less disagreement, and $0 \prec \mathsf{F} \preceq I$. The same diagonal representation gives

$$(2tI + \mathsf{D})\mathsf{F} = \mathsf{F}(2tI + \mathsf{D}) = 2tI. \tag{34}$$

In particular, F commutes with D and with $\mathsf{E}_{BC} = I - \mathsf{D}$; no commutation with E_B or E_C is asserted.

The preceding power bounds imply that inserting $\mathsf{F}(t)$ preserves the product estimate:

Lemma 4.8. *For any choice of $t > 0$, we have:*

$$\|\mathsf{E}_B \mathsf{F}(t) \mathsf{E}_C\|_\infty \leq \gamma. \tag{35}$$

Proof. We will prove this by considering a power series expansion of $\mathsf{F}(t)$. Indeed, we have:

$$\begin{aligned} 2tI + \mathsf{D} &= 2tI + (I - (I - \mathsf{D})) \\ &= (2t+1)\left(I - \frac{I - \mathsf{D}}{2t+1}\right) \\ &= (2t+1)\left(I - \frac{\mathsf{E}_{BC}}{2t+1}\right). \end{aligned}$$

Since E_{BC} is an average of Hermitian contractions, we have $\|\mathsf{E}_{BC}\|_\infty \leq 1 \Rightarrow \|\mathsf{E}_{BC}/(2t+1)\|_\infty < 1$. Therefore we can indeed express the inverse with a power series as follows:

$$\begin{aligned} (2tI + \mathsf{D})^{-1} &= (2t+1)^{-1} \sum_{\ell=0}^{\infty} \left(\frac{\mathsf{E}_{BC}}{2t+1}\right)^\ell \\ \Rightarrow \mathsf{F} &= \frac{2t}{2t+1} \sum_{\ell=0}^{\infty} \left(\frac{\mathsf{E}_{BC}}{2t+1}\right)^\ell. \end{aligned} \tag{36}$$

Hence we have:

$$\begin{aligned} \|\mathsf{E}_B \mathsf{F} \mathsf{E}_C\|_\infty &= \left\| \frac{2t}{2t+1} \sum_{\ell=0}^{\infty} \frac{1}{(2t+1)^\ell} \mathsf{E}_B \mathsf{E}_{BC}^\ell \mathsf{E}_C \right\|_\infty \\ &\leq \frac{2t}{2t+1} \sum_{\ell=0}^{\infty} \frac{1}{(2t+1)^\ell} \left\| \mathsf{E}_B \mathsf{E}_{BC}^\ell \mathsf{E}_C \right\|_\infty \\ &\leq \frac{2t}{2t+1} \sum_{\ell=0}^{\infty} \frac{\gamma}{(2t+1)^\ell} \text{ (using Proposition 4.6)} \\ &= \frac{2t}{2t+1} \cdot \gamma \cdot \frac{1}{1 - \frac{1}{2t+1}} \\ &= \gamma, \end{aligned}$$

as desired. □

4.4 Part 4: Completing the Proof

We now relate the largest eigenvalue of $\mathbf{G}$ to the quantity just bounded.

Lemma 4.9. *Let $t = \lambda_{\max}(\mathbf{G})$. If $t > 0$, then we have $t^2 \leq \|\mathsf{E}_B \mathsf{F}(t) \mathsf{E}_C\|_\infty$.*

Let us first see why this would suffice to complete the security proof:

Proof of Theorem 4.1, assuming Lemma 4.9. Let $t = \lambda_{\max}(\mathbf{G})$. We claim that $t \leq \sqrt{d/K}$. If $t \leq 0$ this is automatic. If $t > 0$, then combining Lemma 4.9 with Lemma 4.8 implies that:

$$t^2 \leq \|\mathsf{E}_B \mathsf{F}(t) \mathsf{E}_C\|_\infty \leq \frac{d}{K} \Rightarrow t \leq \sqrt{\frac{d}{K}}.$$

Substituting this into Equation (22) gives:

$$p \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{d}{K}}$$

$$= \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^n}{4^n - 1}}.$$

This proves Equation (17) and completes the proof. □

To prove Lemma 4.9, recall from Section 1.1 that the high-level idea is to use the following identity:

$$(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B + \mathsf{E}_C) = (\mathsf{E}_B - \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B - \mathsf{E}_C) + 2\mathsf{E}_B\mathsf{F}(t)\mathsf{E}_C + 2\mathsf{E}_C\mathsf{F}(t)\mathsf{E}_B. \tag{37}$$

The latter two terms of the RHS are handled by the norm $\|\mathsf{E}_B\mathsf{F}\mathsf{E}_C\|_\infty$. To handle the first term on the RHS, we have a few lemmas:

Proposition 4.10. *Let* D *be as defined in Equation* (28)*. Then we have:*

$$-\mathsf{D} \preceq \mathsf{E}_B - \mathsf{E}_C \preceq \mathsf{D}. \tag{38}$$

Proof. For every $u, v \in \{+1, -1\}$, define

$$\Gamma^{u,v} := \frac{1}{4}\left(I_{\mathbf{ABC}} + u\mathsf{E}_B + v\mathsf{E}_C + uv\mathsf{E}_{BC}\right). \tag{39}$$

We claim that $\Gamma^{u,v} \succeq 0$ for all u, v. To see this, for a fixed key k, define $W_{k,B} := \Theta_k \otimes B_k \otimes I_{\mathbf{C}}$, and $W_{k,C} := \Theta_k \otimes I_{\mathbf{B}} \otimes C_k$. These are Hermitian operators and have operator norm at most 1. They commute, because $W_{k,B}W_{k,C} = \Theta_k^2 \otimes B_k \otimes C_k = I_{\mathbf{A}} \otimes B_k \otimes C_k = W_{k,C}W_{k,B}$. Thus $I + uW_{k,B}$ and $I + vW_{k,C}$ are commuting positive semidefinite operators (which means that they can be simultaneously diagonalized), so

$$\Gamma_k^{u,v} := \frac{1}{4}(I + uW_{k,B})(I + vW_{k,C}) \succeq 0.$$

Expanding and averaging uniformly over k yields $\Gamma^{u,v} \succeq 0$ as claimed. Now the conclusion follows from noting that:

$$4\Gamma^{+,-} = \mathsf{D} + (\mathsf{E}_B - \mathsf{E}_C), \qquad 4\Gamma^{-,+} = \mathsf{D} - (\mathsf{E}_B - \mathsf{E}_C),$$

and we have established that $\Gamma^{+,-}$ and $\Gamma^{-,+}$ are both positive semidefinite. □

Proposition 4.11 (Quadratic sandwich bound)**.** *For any* $t > 0$*, we have* $(\mathsf{E}_B - \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B - \mathsf{E}_C) \preceq 2t\mathsf{D}$.

Proof. By Proposition 4.10 and Item 2 of Lemma 2.8, the following matrix:

$$\begin{bmatrix} \mathsf{D} & \mathsf{E}_B - \mathsf{E}_C \\ \mathsf{E}_B - \mathsf{E}_C & \mathsf{D} \end{bmatrix}$$

is PSD. Consequently, we also have:

$$\begin{bmatrix} \mathsf{D} & \mathsf{E}_B - \mathsf{E}_C \\ \mathsf{E}_B - \mathsf{E}_C & 2tI + \mathsf{D} \end{bmatrix} = \begin{bmatrix} \mathsf{D} & \mathsf{E}_B - \mathsf{E}_C \\ \mathsf{E}_B - \mathsf{E}_C & \mathsf{D} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 2tI \end{bmatrix} \succeq 0.$$

Now applying Item 1 of Lemma 2.8 (noting that $2tI + \mathsf{D}$ is positive definite) implies that:

$$\begin{aligned} \mathsf{D} - (\mathsf{E}_B - \mathsf{E}_C)(2tI + \mathsf{D})^{-1}(\mathsf{E}_B - \mathsf{E}_C) &\succeq 0 \\ \Leftrightarrow 2t\mathsf{D} &\succeq (\mathsf{E}_B - \mathsf{E}_C)2t(2tI + \mathsf{D})^{-1}(\mathsf{E}_B - \mathsf{E}_C) \\ &= (\mathsf{E}_B - \mathsf{E}_C)\mathsf{F}(\mathsf{E}_B - \mathsf{E}_C), \end{aligned}$$

as desired. □

Corollary 4.12. *We have* $(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(t)(\mathsf{E}_B + \mathsf{E}_C) \preceq 2t\mathsf{D} + 2\mathsf{E}_B\mathsf{F}\mathsf{E}_C + 2\mathsf{E}_C\mathsf{F}\mathsf{E}_B$.

Proof. This is immediate from (37) after plugging in Proposition 4.11. □

With this, we can complete the proof of Lemma 4.9:

Proof of Lemma 4.9. Assume $t = \lambda_{\max}(\mathbf{G}) > 0$ and choose a unit eigenvector $|\psi\rangle$ on $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$ such that

$$\mathbf{G}|\psi\rangle = t|\psi\rangle. \tag{40}$$

Using (29), this eigenvector equation becomes

$$(\mathsf{E}_B + \mathsf{E}_C)|\psi\rangle = (2tI + \mathsf{D})|\psi\rangle. \tag{41}$$

Recall that $(2tI + \mathsf{D})\mathsf{F} = 2tI$. Applying the eigenvector equation on both sides of the quadratic form gives

$$\begin{aligned}\langle\psi|(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(\mathsf{E}_B + \mathsf{E}_C)|\psi\rangle &= \langle\psi|(2tI + \mathsf{D})\mathsf{F}(2tI + \mathsf{D})|\psi\rangle \\ &= 4t^2 + 2t\langle\psi|\mathsf{D}|\psi\rangle.\end{aligned} \tag{42}$$

On the other hand, Corollary 4.12 implies that:

$$\langle\psi|(\mathsf{E}_B + \mathsf{E}_C)\mathsf{F}(\mathsf{E}_B + \mathsf{E}_C)|\psi\rangle \leq 2t\langle\psi|\mathsf{D}|\psi\rangle + 2\langle\psi|\mathsf{E}_B\mathsf{F}\mathsf{E}_C|\psi\rangle + 2\langle\psi|\mathsf{E}_C\mathsf{F}\mathsf{E}_B|\psi\rangle. \tag{43}$$

Chaining these together and using $|\langle\psi|\mathsf{E}_B\mathsf{F}\mathsf{E}_C|\psi\rangle| \leq \|\mathsf{E}_B\mathsf{F}\mathsf{E}_C\|_\infty$ implies:

$$4t^2 + 2t\langle\psi|\mathsf{D}|\psi\rangle \leq 2t\langle\psi|\mathsf{D}|\psi\rangle + 4\|\mathsf{E}_B\mathsf{F}\mathsf{E}_C\|_\infty,$$

which implies $t^2 \leq \|\mathsf{E}_B\mathsf{F}(t)\mathsf{E}_C\|_\infty$ as desired. □

5 Statement on AI Use

Activity	**AI used**
Asking the research question	No
Coming up with the approach	Yes
Proof development	Yes
Writing and exposition	Yes
Checking for bugs	Yes
Other supporting tasks	Yes

Details (elaborating on the points above). The authors are fully accountable for the contents and correctness of this paper. The construction and proof were found by (Codex) GPT-5.6 Sol Ultra in concurrent and independent extended conversations with the authors, after which we decided to merge our original versions [AS26; Rag26]. This was in part using harness ideas generated by the authors based on the UCLA Moonshot Harness [ZHC+26] and [Ope26]. Preliminary versions of this paper were then drafted by GPT-5.6 Sol Ultra, which the human authors carefully verified and significantly cleaned up. GPT-6 Astra, GPT-5.6 Sol, and Claude Fable 5 were all used for proofreading and editorial tasks. Codex GPT-5.6 Sol and GPT-6 Astra were used to produce Lean 4 code for our proofs.

6 Acknowledgements

PA is supported by the National Science Foundation under the grants FET-2329938, CAREER-2341004, and FET-2530160.

SR thanks Alexander Poremba and Vinod Vaikuntanathan for introducing him to the problem of unclonable encryption and for a wonderful collaboration on this front. SR additionally thanks Vinod Vaikuntanathan and Aparna Gupte for helpful feedback on drafts of this paper, and Alper Çakan, Eli Goldin, Yao-Ting Lin, and Justin Raizes for helpful discussions. SR's access to ChatGPT Pro was supported by the UK AISI alignment project and the Simons Institute for the Theory of Computing, and his access to Claude Max was supported by Jane Street. SR's work on this was partly while a student researcher at Google Quantum AI; the remainder was supported by NSF CNS-2534400 and in part by Jane Street.

AS sincerely thanks Fatih Kaleoglu for patiently teaching him about unclonable encryption during a visit to UCSB long ago!

References

[AB24] Prabhanjan Ananth and Amit Behera. "A Modular Approach to Unclonable Cryptography". In: *Advances in Cryptology – CRYPTO 2024*. Springer, 2024, pp. 3–37. DOI: 10.1007/978-3-031-68394-7_1. URL: https://eprint.iacr.org/2023/1797.

[AK21] Prabhanjan Ananth and Fatih Kaleoglu. "Unclonable Encryption, Revisited". In: *Theory of Cryptography*. Vol. 13042. Lecture Notes in Computer Science. Springer, 2021, pp. 299–329. DOI: 10.1007/978-3-030-90459-3_11.

[AKL23] Prabhanjan Ananth, Fatih Kaleoglu, and Qipeng Liu. "Cloning Games: A General Framework for Unclonable Primitives". In: *Advances in Cryptology—CRYPTO 2023, Part V*. Vol. 14085. Lecture Notes in Computer Science. Springer, 2023, pp. 66–98. DOI: 10.1007/978-3-031-38554-4_3.

[AKY25] Prabhanjan Ananth, Fatih Kaleoglu, and Henry Yuen. "Simultaneous Haar Indistinguishability with Applications to Unclonable Cryptography". In: *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*. Vol. 325. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025, 7:1–7:23. DOI: 10.4230/LIPIcs.ITCS.2025.7.

[Ana+22] Prabhanjan Ananth, Fatih Kaleoglu, Xingjian Li, Qipeng Liu, and Mark Zhandry. "On the Feasibility of Unclonable Encryption, and More". In: *Advances in Cryptology—CRYPTO 2022, Part II*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 212–241. DOI: 10.1007/978-3-031-15979-4_8.

[Ana26] Prabhanjan Ananth. *Non-Local Search-to-Decision Reduction over $\mathbb{F}_2$, and More*. 2026. arXiv: 2608.19091 [quant-ph]. URL: https://arxiv.org/abs/2608.19091.

[AS26] Prabhanjan Ananth and Amit Sahai. *Unconditional Unclonable Encryption*. 2026. arXiv: 2607.21551 [quant-ph]. URL: https://arxiv.org/abs/2607.21551.

[BBC26a] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. *Statistically Secure Uncloneable Encryption of Arbitrary Messages*. 2026. arXiv: 2607.28561 [quant-ph]. URL: https://arxiv.org/abs/2607.28561.

[BBC26b] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. *The Uncloneable Bit Exists*. 2026. arXiv: 2603.08916 [quant-ph].

[BC26] Archishna Bhattacharyya and Eric Culf. "Uncloneable Encryption from Decoupling". In: *Nature Physics* 22.2 (2026), pp. 315–318. DOI: 10.1038/s41567-025-03154-7.

[BCR26] Anne Broadbent, Eric Culf, and Denis Rochette. "Optimal Untelegraphable Encryption and Implications for Uncloneable Encryption". In: *IACR Communications in Cryptology* 3.2 (2026). DOI: 10.62056/akgy4fbmo.

[BDF26] Pierre Botteron, Sébastien Designolle, and Omar Fawzi. *Pauli Encodings & Unclonable Encryption*. 2026. arXiv: 2607.28602 [quant-ph]. URL: https://arxiv.org/abs/2607.28602.

[BG26a] James Bartusek and Eli Goldin. *A Note on Boosting Uncloneable Encryption in Microcrypt*. 2026. arXiv: 2605.27647 [cs.CR]. URL: https://arxiv.org/abs/2605.27647.

[BG26b] James Bartusek and Eli Goldin. “Unclonable Encryption in the Haar Random Oracle Model”. In: *Advances in Cryptology – CRYPTO 2026, 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–20, 2026, Proceedings, Part V*. Ed. by Nadia Heninger and Mike Rosulek. Lecture Notes in Computer Science. Springer Nature Switzerland, 2026, pp. 311–340. DOI: 10.1007/978-3-032-35409-9_11.

[BL20] Anne Broadbent and Sébastien Lord. “Uncloneable Quantum Encryption via Oracles”. In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. DOI: 10.4230/LIPIcs.TQC.2020.4.

[Bot+26] Pierre Botteron, Anne Broadbent, Eric Culf, Ion Nechita, Clément Pellegrini, and Denis Rochette. “Towards Unconditional Uncloneable Encryption”. In: *Quantum* 10 (2026), p. 2157. DOI: 10.22331/q-2026-07-08-2157.

[CG24] Andrea Coladangelo and Sam Gunn. “How to Use Quantum Indistinguishability Obfuscation”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. DOI: 10.1145/3618260.3649779. URL: https://eprint.iacr.org/2023/1756.

[Cho75] Man-Duen Choi. “Completely Positive Linear Maps on Complex Matrices”. In: *Linear Algebra and its Applications* 10.3 (1975), pp. 285–290. DOI: 10.1016/0024-3795(75)90075-0.

[CLX26a] Andrea Coladangelo, Qipeng Liu, and Ziyi Xie. “The Curious Case of "XOR Repetition" of Monogamy-Of-Entanglement Games”. In: *17th Innovations in Theoretical Computer Science Conference, ITCS 2026, Bocconi University, Milan, Italy, January 27-30, 2026*. Ed. by Shubhangi Saraf. Vol. 362. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026, 41:1–41:20. DOI: 10.4230/LIPICS.ITCS.2026.41. URL: https://doi.org/10.4230/LIPIcs.ITCS.2026.41.

[CLX26b] Andrea Coladangelo, Qipeng Liu, and Ziyi Xie. *Unclonable Encryption from BB84 States: A Simultaneous Goldreich–Levin Reduction*. 2026. arXiv: 2608.17629 [quant-ph]. URL: https://arxiv.org/abs/2608.17629.

[Die82] Dennis Dieks. “Communication by EPR Devices”. In: *Physics Letters A* 92.6 (1982), pp. 271–272. DOI: 10.1016/0375-9601(82)90084-6.

[Got03] Daniel Gottesman. “Uncloneable Encryption”. In: *Quantum Information and Computation* 3.6 (2003), pp. 581–602. DOI: 10.26421/QIC3.6-2.

[Jam72] Andrzej Jamiołkowski. “Linear Transformations Which Preserve Trace and Positive Semidefiniteness of Operators”. In: *Reports on Mathematical Physics* 3.4 (1972), pp. 275–278. DOI: 10.1016/0034-4877(72)90011-0.

[MH25] Fermi Ma and Hsin-Yuan Huang. “How to Construct Random Unitaries”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*. Ed. by Michal Koucký and Nikhil Bansal. ACM, 2025, pp. 806–809. DOI: 10.1145/3717823.3718254. URL: https://doi.org/10.1145/3717823.3718254.

[MST21] Christian Majenz, Christian Schaffner, and Mehrdad Tahmasbi. *Limitations on Uncloneable Encryption and Simultaneous One-Way-to-Hiding*. 2021. arXiv: 2103.14510 [quant-ph].

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. "A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations". In: *New Journal of Physics* 10.7 (2008), p. 073013.

[Ope26] OpenAI. *Prompt Used for "A Proof of the Cycle Double Cover Conjecture"*. July 9, 2026. URL: https://cdn.openai.com/pdf/04d1d1e4-bc75-476a-97cf-49055cd98d31/cdc_prompt.pdf (visited on 07/22/2026).

[Rag26] Seyoon Ragavan. *Efficient Unclonable Encryption from Pauli Eigenstates*. 2026. arXiv: 2607.21811 [quant-ph]. URL: https://arxiv.org/abs/2607.21811.

[Tom+13] Marco Tomamichel, Serge Fehr, Jędrzej Kaniewski, and Stephanie Wehner. "A monogamy-of-entanglement game with applications to device-independent quantum cryptography". In: *New Journal of Physics* 15.10 (2013), p. 103002.

[Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. DOI: 10.1017/9781316848142.

[WZ82] William K. Wootters and Wojciech H. Zurek. "A Single Quantum Cannot Be Cloned". In: *Nature* 299 (1982), pp. 802–803. DOI: 10.1038/299802a0.

[ZHC+26] Junyi Zhang*, Xinjie He*, Hyunsik Chae, Ethan Ji, Eric Jiang, Rushil Raghavan, Yiwen Kou, Alex Taylor, Kai-Wei Chang†, Raghu Meka†, Violet Peng†, Amit Sahai†, Terence Tao†, and Wei Wang†. *UCLA Moonshot Harness*. *Co-first authors with equal contribution. The remaining students are ordered by contribution to the harness. †Principal investigators, listed at the end in alphabetical order by last name. 2026.

A Glossary of Notation

The table lists the notation used throughout the construction and security proof. Auxiliary quantities used in a single proof are defined where they appear.

Symbol	Definition and Meaning
$n,\ d$	The ciphertext has n qubits and Hilbert-space dimension $d = 2^n$.
$\mathbf{A},\ \mathbf{E},\ \mathbf{B},\ \mathbf{C}$	The reference register, ciphertext register, and Bob's and Charlie's output registers, respectively. The reference register is a copy of $\mathbf{E}$.
$\mathcal{K}_n,\ K$	The non-identity Pauli key space, of size $K = 4^n - 1 = d^2 - 1$; see Definition 4.2.
$k = (x, z),\ P_{x,z}$	The secret key and its tensor Pauli $P_{x,z} = \bigotimes_{j=1}^{n} i^{x_j z_j} X^{x_j} Z^{z_j}$.
$[k, k']$	For $k = (x, z)$ and $k' = (x', z')$, the Pauli symplectic form $\langle x, z' \rangle + \langle z, x' \rangle \in \mathbb{F}_2$.
Ξ^{Pauli}	The $4^n \times 4^n$ symplectic character matrix with entries $(-1)^{[k,k']}$; see Definition 2.6.
$m,\ b$	The plaintext bit and challenge bit, respectively; they coincide in the fixed-message one-bit game.
q_j	Equals 1 when the jth Pauli factor is non-identity, and 0 otherwise. The preparation bits satisfy $\bigoplus_j q_j r_j = m$.
U_k	The product of the local basis changes used to prepare the ciphertext; see Section 3.1.
$\rho_{\mathsf{ct}}^{(k,m)}$	The ciphertext density operator $d^{-1}(I + (-1)^m P_{x,z})$.
$\lvert\Omega\rangle_{\mathbf{AE}}$	The normalized maximally entangled state between the reference and ciphertext registers; see Theorem 2.1.
$\Phi,\ \tau$	The cloner's channel from $\mathbf{E}$ to $\mathbf{B} \otimes \mathbf{C}$, and its normalized Choi state on $\mathbf{A} \otimes \mathbf{B} \otimes \mathbf{C}$.
Θ_k	The transposed Pauli $(P_{x,z}^T)_{\mathbf{A}}$ on the reference register.
$B_k,\ C_k$	Bob's and Charlie's Hermitian decoding observables for key k, with norm at most one.
$\mathsf{E}_B,\ \mathsf{E}_C$	The averaged reference–receiver correlations defined in (25) and (26). Their expectations in τ are twice the corresponding success probabilities minus one.
E_{BC}	The averaged receiver correlation in (27). Its expectation is the probability of agreement minus the probability of disagreement.
D	$I - \mathsf{E}_{BC}$. Its expectation is twice the probability of disagreement.
$\mathbf{G}$	$\frac{1}{2}(\mathsf{E}_B + \mathsf{E}_C - \mathsf{D})$, whose largest eigenvalue controls the winning probability.
p	The winning probability in the fixed-message one-bit game: $p = \frac{1}{2}(1 + \mathrm{Tr}(\mathbf{G}\tau))$; see (22).
γ	d/K, the upper bound on $\lVert\mathsf{E}_B\mathsf{E}_C\rVert_\infty$.
t	A positive parameter in $\mathsf{F}(t)$; in the last step of the proof, $t = \lambda_{\max}(\mathbf{G})$ and the case $t \le 0$ is handled separately.
$\mathsf{F}(t)$	$2t(2tI + \mathsf{D})^{-1}$, defined for $t > 0$. Write $\mathsf{F} = \mathsf{F}(t)$ when t is fixed.