Unclonable encryption from BB84 states: a simultaneous Goldreich-Levin reduction

Andrea Coladangelo* Qipeng Liu† Ziyi Xie‡

Abstract

Goldreich-Levin reductions are ubiquitous in cryptography: they convert an algorithm capable of guessing $\langle r, m\rangle \pmod 2$ for a hidden string m and a random challenge r, to one that is capable of extracting the entirety of m.

Here, we describe a "simultaneous" Goldreich-Levin reduction for two entangled parties who are capable of guessing $\langle r, m\rangle$ given uniformly random *identical* challenges r. This allows to upgrade *any* unclonable encryption scheme satisfying "search" security to one satisfying the gold standard of unclonable "indistinguishability". As a corollary, we show that the simplest candidate unclonable encryption scheme from BB84 states satisfies unclonable indistinguishability.

This result was discovered by GPT-5.6 Ultra after a few interactions. Our prompts included recent results on unclonable encryption by Ananth and Sahai [AS26] and Ragavan [Rag26].

AI usage statement. The result of this work was discovered by GPT-5.6 Ultra. In our prompts, we included the two recent results by Ananth and Sahai, and Ragavan, and we asked GPT to prove security of a scheme based solely on BB84 states (rather than Pauli states). GPT returned a correct proof. We then asked it to find a different proof based on a simultaneous Goldreich–Levin reduction, which it also found.

We find the new reduction to be quite magical. We have spent considerable time trying to understand why it works and to arrive at a good conceptual explanation of it. As Terry Tao recently put it, echoing Bill Thurston, mathematics is not just about producing theorems and proofs, but about developing "human understanding". This experience has reinforced our belief that at least part of our role as researchers in the future will indeed be centered around developing human understanding of AI-discovered proofs. We hope that our community will embrace this role alongside the other roles that we hope researchers will continue to play. In this work, we try to do our part by developing what we hope is a satisfying conceptual understanding of the security of unclonable encryption.

*University of Washington. Email: coladan@cs.washington.edu
†University of California San Diego. Email: qipengliu0@gmail.com
‡Tsinghua University. Email: xie-zy21@mails.tsinghua.edu.cn

Contents

1 Introduction **3**
1.1 Our result 5

2 Technical overview **7**

3 The model and main result **15**

4 A warmup: the independent-mask Goldreich–Levin reduction **17**
4.1 Setup 17
4.2 Independent-mask GL reduction 18

5 The common-mask GL reduction **20**
5.1 The reduction algorithm 20

6 Analysis of the reduction, part 1: the "filtered" overlap **22**

7 Analysis of the reduction, part 2: lower bounding the "filtered overlap" by the original winning probability **24**
7.1 Some setup and useful facts 24
7.2 A key idea: the geometric filter as an inverse 25

8 Uniformity and efficiency analysis **27**

A A common-mask strategy on which plain GL fails completely **32**

1 Introduction

Unclonable encryption provides a security property that no classical encryption scheme can achieve: an adversary may arbitrarily process a ciphertext into two parts *before* the secret key is revealed, yet, two non-communicating recipients should not be able to simultaneously recover the plaintext, *even after learning the secret key* [Got03; BL20].

There are two natural formulations of security for such an encryption scheme. In the *search* formulation, modeled in Figure 1 below by a game G, the challenger samples a key-message pair (k, m) from a distribution π, where $m \in \{0,1\}^n$ (for some security parameter n); sends a ciphertext or "token" state $\rho_{k,m}$ to $\mathcal{A}$, who processes the state into two parts; the two recipients, upon learning k, must both guess the *entire* message m. Security requires their joint winning probability in G to be negligible in the security parameter. This guarantee, however, does not rule out substantial shared leakage about the plaintext. For example, both recipients might learn its first bit while remaining unable to reconstruct the remaining bits. Yet, that single bit could already reveal the answer to a sensitive yes/no question encoded in the encrypted message.

The *decision* formulation, or unclonable *indistinguishability*, provides a stronger notion of security. Here, a hidden challenge bit selects one of two equal-length messages, and both recipients must identify which of the two messages was encrypted. A coordinated random guess trivially succeeds with probability $1/2$, and security requires that every attack achieves at most a negligible advantage over this baseline. In this sense, unclonable indistinguishability captures the idea that, after the ciphertext is processed or "split", at least one of the two recipients learns essentially nothing about the encrypted message. Establishing this stronger guarantee from simultaneous-search security is a central search-to-decision problem in unclonable cryptography.

The BB84 monogamy-of-entanglement game of Tomamichel, Fehr, Kaniewski, and Wehner [TFKW13] provides a basic source of search security for unclonable encryption. In this game, a referee measures n qubits in independently chosen computational or Hadamard bases, obtaining an outcome string $m \in \{0,1\}^n$. Two separated players subsequently learn the string of basis choices $\theta \in \{0,1\}^n$ and must both recover m. Tomamichel et al. [TFKW13] showed that the optimal joint success probability in this game is exactly $\cos^{2n}(\pi/8)$. Broadbent and Lord later adapted this game to construct unclonable encryption [BL20]: the ciphertext is the BB84 state $H^\theta|m\rangle = \bigotimes_i H^{\theta_i}|m_i\rangle$, and the secret key is the string of basis choices θ.

A natural route to go from search security to decision security is suggested by the classical Goldreich–Levin theorem. Goldreich–Levin turns parity prediction into full recovery: an algorithm that predicts the (mod 2) inner product $\langle r, m\rangle$ for a uniformly random mask $r \in \{0,1\}^n$ with noticeable advantage over $1/2$ can be transformed into one that recovers the hidden string m with noticeable probability [GL89]. Thus, a random inner product, which can also be viewed as the parity of a random subset of the bits of m, is a "hardcore bit" of m: any noticeable ability to predict it can be leveraged to recover the entirety of m. It turns out that the Goldreich-Levin theorem is still valid, and is even more elegant, in the quantum world [AC02]: a single superposition query to the algorithm predicting the inner product suffices to recover the entirety of m!

At first sight, this is precisely the type of reduction needed to upgrade a search-secure unclonable encryption scheme to one satisfying the stronger *indistinguishability* notion: given a search-secure scheme modeled by a game G, one can encrypt a bit b by masking it with the inner product $\langle r, m\rangle$, i.e. using the latter as a one-time pad below). We denote by $\mathsf{GL}(G)$ the corresponding *common-mask decision game* (see Figure 1 below): here, the challenger additionally samples a uniformly random mask r, and reveals (k, r) to both parties; their goal is to both guess the bit $\langle r, m\rangle$. An attack on the one-bit encryption is equivalent to a strategy for $\mathsf{GL}(G)$. Thus, if one could apply Goldreich–Levin in this *simultaneous* setting, such a strategy could then be converted into a strategy for the search

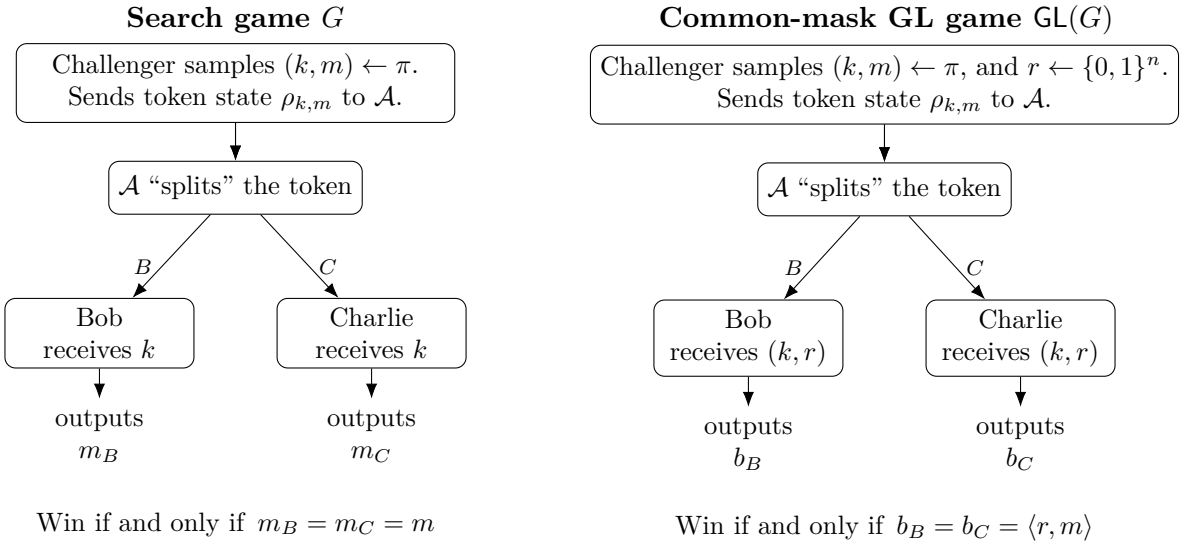

Figure 1: The search game G (left) and its common-mask Goldreich–Levin variant $\mathsf{GL}(G)$ (right).

game G in which both recipients extract m, contradicting search security.

The subtlety is that the simultaneous setting requires a form of Goldreich–Levin that is not supplied by the standard reduction. In $\mathsf{GL}(G)$, after the split, the parties receive the original key k as well as a *common* random mask: the *same* r is revealed to both parties. The natural reduction in which Bob and Charlie each locally apply a standard quantum Goldreich-Levin reduction only works for *independent* masks, i.e. when Bob and Charlie respectively receive some independently sampled r and s, and are required to guess their respective inner products $\langle r, m \rangle$ and $\langle s, m \rangle$. This distinction is fundamental: earlier work on the feasibility of unclonable encryption already identified barriers showing that standard seeded-extractor and Goldreich–Levin compilers do not directly survive in the unclonable setting [AKLL+22].

Simultaneous Goldreich–Levin and the common-mask barrier. The single-party quantum Goldreich–Levin reduction was developed by Adcock and Cleve [AC02]. Later, Coladangelo, Liu, Liu, and Zhandry established a version allowing quantum auxiliary input, which is useful in other settings of unclonable cryptography [CLLZ21]. Kundu and Tan and, independently, Ananth, Kaleoglu, and Liu extended this approach to the setting of two non-communicating parties [KT25; AKL23]. Crucially, however, these simultaneous reductions use *independent* random masks r and s for the two parties.

Extending these results to the *common* mask setting has proved substantially more elusive. Kundu and Tan explicitly observe that their techniques do not extend to the case in which both recipients receive the same mask [KT25]. Ananth and Behera formulate related identical-challenge simultaneous inner-product statements as conjectures [AB24], while Chevalier, Hermouet, and Vu propose closely related conjectures in the context of simultaneous compute-and-compare obfuscation [CHV23]. Most directly, Ananth, Kaleoglu, and Yuen identify a common-mask simultaneous Goldreich–Levin theorem as precisely the missing ingredient needed for their classical-key Goldreich–Levin compiler [AKY25]. The importance of this distinction becomes concrete for the BB84 monogamy game. Coladangelo, Liu, and Xie [CLX26] studied two decision variants of this game and showed that merely asking both players to predict a fixed parity does not suffice: a fixed parity can be guessed with constant bias for every n, so ordinary XOR repetition fails to amplify security. By contrast, when a uniformly

random mask r is revealed identically to both players, they proved exponentially small advantage for semi-classical strategies and conjectured the same behavior for arbitrary quantum strategies. Related "no-telegraphing" results establish security in similar models [CKNY25; ÇG24], but also do not yield the fully quantum common-mask search-to-decision reduction needed here.

Thus, in this work, we focus on the following question:

Does there exist a simultaneous Goldreich-Levin reduction
for two entangled quantum parties in the common-mask setting?

Recent momentous progress on unclonable encryption. In the meantime, there has been rapid progress via other methods recently culminating in a proof of the unconditional existence of unclonable encryption satisfying the gold standard of indistinguishable-security for single-bit messages. First, Bhattacharyya and Culf obtained inverse-polynomial security [BC26]. Then, in a breakthrough work, Bhattacharyya, Broadbent, and Culf obtained negligible security, albeit with inefficient encryption and decryption [BBC26a]. Two independent and concurrent AI-assisted breakthroughs of Ananth and Sahai [AS26], and Ragavan [Rag26] proved the existence of efficient constructions from Pauli eigenstates with negligible security. Finally, Bhattacharyya, Broadbent, and Culf [BBC26b] extended these results to unclonable encryption of arbitrary messages.

However, all of these results rely on direct, quite technical, analysis tailored to each of their proposed schemes, leaving open the elusive question of a general reduction from $\mathsf{GL}(G)$ to G. There are several appealing reasons for wanting such a reduction:

- The security of the simplest conjectured unclonable encryption scheme based on BB84 states [BL20; CLX26] is still open. A simultaneous Goldreich-Levin reduction would immediately prove its security.
- Such a reduction would establish that, for unclonable encryption, search security can be generically upgraded to the stronger indistinguishability guarantee. This would be a clean, modular bridge from search to decision, avoiding the need for a separate analysis for each unclonable encryption construction. In particular, *any* scheme with a search security game G would automatically inherit security for the Goldreich–Levin compiled game $\mathsf{GL}(G)$ from a bound on the optimal value in G, which is often considerably simpler to establish directly.
- Finally, given the remarkable versatility of the classical Goldreich–Levin theorem, we expect that a common-mask simultaneous Goldreich–Levin reduction would find applications beyond unclonable encryption, as a useful tool for constructing and analyzing other unclonable cryptographic primitives. For example, such a reduction (if efficient) would also allow one to lift search security to decision security in the *computational* setting, where a direct information-theoretic analysis, as in the recent sequence of papers [BBC26a; AS26; Rag26], may be hard to carry out.

1.1 Our result

We prove that a common-mask Goldreich–Levin reduction exists. Our theorem is information-theoretic and is not tied to BB84 or any other specific construction. Let G be any search unclonable-encryption game, and let $\mathsf{GL}(G)$ be its corresponding common-mask GL game (as in Figure 1). We show the following.

Theorem 1.1 (Informal)**.** *Let p_{Search} and p_{Pred} be the optimal success probabilities in G and $\mathsf{GL}(G)$, respectively, and define*

$$\Delta := 2p_{\mathrm{Pred}} - 1.$$

Then,

$$p_{\mathrm{Search}} \geq \Delta^4 .$$

More concretely, there is an explicit reduction that transforms any strategy $\mathcal{P}$ for $\mathsf{GL}(G)$, winning with probability $\frac{1}{2} + \frac{\Delta}{2}$ into a strategy $\mathcal{P}'$ for G winning with probability at least Δ^4. Thus, as long as Δ is non-negligible, the reduction yields a non-negligible success probability in G.

While the complexity of the reduction is not important when adversaries are unbounded, we note that our reduction is efficient in the following sense: if a lower bound on $0 < \gamma \leq \Delta$ is known, the reduction is uniform and has running time polynomial in the running time of $\mathcal{P}$, and $1/\gamma$. When a lower bound on Δ is not known, we show that there is still a uniform reduction with *expected* running time linear in the running time of $\mathcal{P}$, but at the cost of a polynomial loss in the search success guarantee, which becomes $p_{\mathrm{Search}} \geq \Omega(\Delta^6)$.

We note that our reduction is rooted in some ideas already present in the security analysis of Ananth and Sahai [AS26]. In particular, one key idea of considering a certain *filter* to reduce Bob and Charlie's "disagreement" appears in the security analysis of Ananth and Sahai, though it does not have any operational meaning there. The main new insight is that such a filter can be actually implemented to achieve a Goldreich-Levin reduction that applies to any unclonable encryption game.

While in this paper we focus on unclonable-encryption games as described above, we note that our reduction from Theorem 1.1 applies also to *monogamy games*, essentially unchanged. We discuss this briefly in Remark 5.2.

Our Theorem 1.1 immediately yields two new constructions of unclonable encryption based on BB84 states [TFKW13] and coset states [CLLZ21; CV22].

Corollary 1.2 (Unclonable encryption from BB84 states)**.** *Consider the following one-bit encryption scheme. To encrypt a bit b, sample $\theta, r, m \leftarrow \{0,1\}^n$, use (θ, r) as the secret key, and output the ciphertext*

$$\left(H^\theta |m\rangle, b \oplus \langle r, m\rangle\right).$$

Decryption measures the quantum register in the basis specified by θ to recover m and then removes the $\langle r, m\rangle$ one-time pad. This scheme satisfies unclonable indistinguishability. More precisely, the optimal winning probability in $\mathsf{GL}(G)$ (where G is the corresponding search game) is at most

$$\frac{1}{2} + \frac{1}{2}\cos^{n/2}\left(\frac{\pi}{8}\right).$$

The latter follows immediately from Theorem 1.1 along with security of the corresponding search game from [TFKW13], which says that the optimal winning probability in the search game is $\cos^{2n}\left(\frac{\pi}{8}\right)$. This also resolves a conjecture posed in [CLX26] (there, the conjecture is posed for monogamy games, which our reduction also applies to).

Corollary 1.3 (Unclonable encryption from subspace coset states)**.** *Let n be even. Sample a uniformly random $n/2$-dimensional subspace $A \subseteq \mathbb{F}_2^n$, a uniform $r \in \{0,1\}^n$, and uniform $s, s' \in \mathbb{F}_2^n$. For each A, fix an efficiently computable bijection*

$$m_A : (\mathbb{F}_2^n / A) \times (\mathbb{F}_2^n / A^\perp) \longrightarrow \{0,1\}^n,$$

where $\mathbb{F}_2^n/A$ and $\mathbb{F}_2^n/A^\perp$ are the spaces of cosets of A and $A^\perp$ respectively, and abbreviate $m_A(s,s') := m_A(s+A, s'+A^\perp)$. Encrypt a bit b using secret key (A,r) as

$$\left(|A_{s,s'}\rangle, b \oplus \langle r, m_A(s,s')\rangle\right), \qquad |A_{s,s'}\rangle := X^s Z^{s'} |A\rangle, \qquad |A\rangle := \frac{1}{\sqrt{|A|}} \sum_{a \in A} |a\rangle.$$

The scheme satisfies unclonable indistinguishability. The optimal winning probability in $\mathsf{GL}(G)$ (where G is the corresponding search game) is at most

$$\frac{1}{2} + \frac{e^{1/8}}{2} \cos^{n/4}\left(\frac{\pi}{8}\right).$$

Moreover, assuming post-quantum indistinguishability obfuscation and one-way functions, the ciphertext may additionally contain obfuscated membership programs for $A+s$ and $A^\perp + s'$ while retaining computational unclonable indistinguishability against quantum polynomial-time adversaries.

The latter follows from [CV22], which shows that the optimal winning probability in the corresponding search game is $\sqrt{e}\cos^n(\frac{\pi}{8})$.

We remark that the upper bounds in Corollaries 1.2 and 1.3 are most likely not tight, as they are obtained by combining our general search-to-decision reduction, which pays a quartic loss, with the corresponding search bounds. In particular, numerical evidence strongly suggests that the tight upper bound for unclonable encryption from BB84 states is actually $\frac{1}{2} + \frac{1}{2}\cos^{2n}(\frac{\pi}{8})$.

2 Technical overview

The goal and notation. Fix a search game G as in Figure 1 (a more detailed description is also given in Section 3). We use the same notation: k is the key, $m \in \{0,1\}^n$ is the hidden string, $\rho^A_{k,m}$ is the token state, which is split by $\mathcal{A}$ into $|\psi_{k,m}\rangle^{BC}$ (here the superscripts indicate registers, where B and C can be of arbitrary dimension). We take the state after the split to be pure without loss of generality since $\mathcal{A}$ can always supply any purifying register to one of the two parties (and this can only increase their success probability). Recall that $\mathcal{A}$ does not have the key k, which is only revealed to Bob and Charlie, who cannot communicate. Their goal in G is to guess m. In $\mathsf{GL}(G)$, Bob and Charlie additionally receive the same mask $r \in \{0,1\}^n$, and their goal is to guess the bit $\langle r, m\rangle := r_1 m_1 \oplus \cdots \oplus r_n m_n$.

Fix a strategy $\mathcal{P}$ for $\mathsf{GL}(G)$, and let p_{Pred} be its success probability. Set

$$\Delta := 2p_{\mathrm{Pred}} - 1.$$

Our reduction keeps the splitting procedure $\mathcal{A}$ unchanged and constructs a strategy for G where the probability of jointly recovering m is at least Δ^4.

Warmup: The single-party quantum Goldreich–Levin. We start by recalling the single-party quantum Goldreich–Levin reduction from Adcock and Cleve [AC02]. For a key k and a mask r, let Bob's local strategy in $\mathsf{GL}(G)$ be described by projectors $\Pi^B_{k,r,0}$ and $\Pi^B_{k,r,1}$. We will refer to this as Bob's *decoder* for simplicity. Equivalently, Bob's decoder in $\mathsf{GL}(G)$ consists of measuring the binary observable

$$V^B_{k,r} := \Pi^B_{k,r,0} - \Pi^B_{k,r,1}.$$

Now, notice that $V_{k,r}^B$ is nothing other than a *reflection* (and hence a unitary). So, we can use it to define the following controlled-unitary. Let R be an n-qubit auxiliary register. Define the "mask-controlled" reflection

$$\mathsf{CV}_k^B := \sum_{r\in\{0,1\}^n} \left(|r\rangle\langle r|\right)^R \otimes V_{k,r}^B\,,$$

which acts coherently as $\sum_r \alpha_r |r\rangle^R|\phi\rangle^B \;\longmapsto\; \sum_r \alpha_r |r\rangle^R V_{k,r}^B|\phi\rangle^B$.

The GL reduction circuit of [AC02] places R in a uniform superposition, applies this controlled reflection, and "recombines" the R branches with a second Hadamard transform:

$$\mathsf{GL}_k^B := (H_R^{\otimes n} \otimes I_B)\,\mathsf{CV}_k^B\,(H_R^{\otimes n} \otimes I_B).$$

For an arbitrary state $|\phi\rangle^B$, expanding these three operations gives

$$\mathsf{GL}_k^B(|0^n\rangle^R|\phi\rangle^B) = \frac{1}{2^n}\sum_{y,r}(-1)^{\langle r,y\rangle}|y\rangle^R V_{k,r}^B|\phi\rangle^B. \tag{2.1}$$

We define $V_{k,r}^C$ and GL_k^C analogously for Charlie, introducing a separate auxiliary register S (though we will not need Charlie just yet in the single-party setting). We denote the operator-valued Fourier coefficients of the two reflection families by

$$\widehat{V}_k^B(y) := \frac{1}{2^n}\sum_r (-1)^{\langle r,y\rangle} V_{k,r}^B, \qquad \widehat{V}_k^C(y) := \frac{1}{2^n}\sum_s (-1)^{\langle s,y\rangle} V_{k,s}^C.$$

Then, Equation (2.1) can be written more compactly as

$$\mathsf{GL}_k^B(|0^n\rangle^R|\phi\rangle^B) = \sum_y |y\rangle^R \widehat{V}_k^B(y)|\phi\rangle^B,$$

and similarly for Charlie. Note then that the probability that Bob's GL "extractor" outputs m is simply $\|\widehat{V}_k^B(m)|\phi\rangle^B\|^2$.

Why does the reduction work in the single-party setting? The controlled reflection $\sum_{r\in\{0,1\}^n}\left(|r\rangle\langle r|\right)^R \otimes V_{k,r}^B$ coherently converts Bob's guess for $\langle r, m\rangle$ into a phase on the branch indexed by r. If the decoder is usually correct, this phase is usually equal to $(-1)^{\langle r,m\rangle}$. The final Hadamard transform then cancels this phase precisely on the outcome m, causing the corresponding branches to interfere constructively. Equivalently, the correlation with the hidden inner product translates into weight on the branch $\widehat{V}_k^B(m)|\phi\rangle^B$ corresponding to the Fourier coefficient at m. This is the main idea behind the single-party quantum Goldreich–Levin reduction, and formalizing it is not difficult (we refer the reader, for example, to Appendix B.3 in [CLLZ21] for the details).

Independent-mask simultaneous GL. We now consider the *independent*-mask analogue of $\mathsf{GL}(G)$, in which Bob and Charlie receive independent uniformly random masks r and s respectively. In this setting, the simultaneous GL reduction is the most natural quite (and fairly straightforward to analyze): Bob and Charlie each locally apply their respective GL extractor, as described above.

Recall that we denote by $|\psi_{k,m}\rangle^{BC}$ the state of Bob and Charlie after the split. In the reduction, they apply GL_k^B and GL_k^C to their respective registers RB and SC. Then, Equation (2.1), applied jointly to registers RB and SC implies that the unnormalized leftover state on BC, conditioned on both parties outputting m (i.e. both registers R and S containing m) is

$$\begin{aligned}|\eta_{k,m}^{\mathrm{ind}}\rangle &:= (\widehat{V}_k^B(m)\otimes\widehat{V}_k^C(m))|\psi_{k,m}\rangle \\ &= \mathop{\mathbb{E}}_{r,s\leftarrow\{0,1\}^n}\left[(-1)^{\langle r,m\rangle}(-1)^{\langle s,m\rangle}\left(V_{k,r}^B\otimes V_{k,s}^C\right)\right]|\psi_{k,m}\rangle.\end{aligned} \tag{2.2}$$

Thus, the probability of simultaneous correct recovery of m is precisely $\|\eta^{\mathrm{ind}}_{k,m}\|^2$.

It is useful now to separate the two individual correct-output branches. Define

$$|b_{k,m}\rangle := \widehat{V}^B_k(m) \otimes I_C|\psi_{k,m}\rangle, \qquad |c_{k,m}\rangle := I_B \otimes \widehat{V}^C_k(m)|\psi_{k,m}\rangle.$$

The vector $|b_{k,m}\rangle$ is the unnormalized leftover state when Bob alone runs his GL circuit and obtains m; $|c_{k,m}\rangle$ has the analogous meaning when Charlie alone runs his circuit. Thus, $\||b_{k,m}\rangle\|^2$ and $\||c_{k,m}\rangle\|^2$ describe the two individual correct-output probabilities. Simultaneous recovery, instead, is captured by the branch $|\eta^{\mathrm{ind}}_{k,m}\rangle = (\widehat{V}^B_k(m) \otimes \widehat{V}^C_k(m))|\psi_{k,m}\rangle$. Now, applying Cauchy–Schwarz gives

$$\left\|\eta^{\mathrm{ind}}_{k,m}\right\|^2 \geq \left|\langle \psi_{k,m} \mid \eta^{\mathrm{ind}}_{k,m}\rangle\right|^2 = \left|\langle b_{k,m} \mid c_{k,m}\rangle\right|^2. \tag{2.3}$$

Thus, we can lower bound the joint success probability by the *overlap between the individual correct-output branches* of Bob and Charlie. Crucially, large norms of these two vectors alone do not guarantee a large joint success probability: the two vectors must also *align* well. In the independent-mask setting, however, it turns out that successful joint prediction, on average over the independent challenges r and s, guarantees precisely the required alignment. Indeed, expanding the overlap and the Fourier expressions gives

$$\langle b_{k,m} \mid c_{k,m}\rangle = \mathop{\mathbb{E}}_{r,s\leftarrow\{0,1\}^n}\left[(-1)^{\langle r,m\rangle+\langle s,m\rangle}\langle\psi_{k,m}|(V^B_{k,r}\otimes V^C_{k,s})|\psi_{k,m}\rangle\right].$$

Crucially, the independent-mask experiment samples exactly the same pairs (r,s) that occur in the latter overlap! This is what allows us to bound the RHS in terms of the probability that Bob and Charlie are simultaneously correct. To see this, note that, for fixed (k,m,r,s), the expression inside the expectation is the "correlation" between Bob's and Charlie's prediction outcomes after each outcome is multiplied by the sign corresponding to the correct parity. This correlation is $+1$ when Bob and Charlie are either both correct or both incorrect, and -1 when exactly one of them is correct. Letting p_{kmrs} denote the probability that both players are correct, the latter correlation is therefore at least $p_{kmrs} - (1 - p_{kmrs}) = 2p_{kmrs} - 1$.

Thus, letting $p^{\mathrm{ind}}_{\mathrm{Pred}} \geq 1/2$ denote the probability that both players are correct, averaged over k, m, r, s, gives

$$\mathop{\mathbb{E}}_{(k,m)\leftarrow\pi}\langle b_{k,m} \mid c_{k,m}\rangle \geq 2p^{\mathrm{ind}}_{\mathrm{Pred}} - 1. \tag{2.4}$$

Combining Equation (2.4) with (2.3) and Jensen's inequality gives the desired lower bound on the probability of successful extraction, p_{Search}, and thus on the optimal success probability in G:

$$p_{\mathrm{Search}} \geq \left(2p^{\mathrm{ind}}_{\mathrm{Pred}} - 1\right)^2 = \Delta^2.$$

The full calculation appears in the proof of Proposition 4.3.

The common-mask barrier. We are now ready to move to $\mathsf{GL}(G)$, the game that is relevant to unclonable encryption. Here, the crucial difference from earlier is that the *same* mask r is given to both parties. Using the same notation as before, the winning probability in $\mathsf{GL}(G)$, i.e. the probability that both Bob and Charlie predict correctly, is

$$p_{\mathrm{Pred}} = \mathop{\mathbb{E}}_{\substack{(k,m)\leftarrow\pi\\ r\leftarrow\{0,1\}^n}}\left[\langle\psi_{k,m}|\left(\frac{I_B + (-1)^{\langle r,m\rangle}V^B_{k,r}}{2}\otimes\frac{I_C + (-1)^{\langle r,m\rangle}V^C_{k,r}}{2}\right)|\psi_{k,m}\rangle\right]. \tag{2.5}$$

Crucially, every term in this average involves a matching pair (r, r). By contrast, if the parties simply run their two local GL circuits, the desired branch where both parties are correct corresponds instead to Equation (2.2), which averages over every pair (r, s). Thus, p_{Pred} gives no direct lower bound on the averaged overlap $\mathbb{E}_{(k,m)\leftarrow\pi}\langle b_{k,m} \mid c_{k,m}\rangle$ appearing on the LHS of (2.4). Indeed, Appendix A gives an example in which the averaged overlap is zero even though $p_{\mathrm{Pred}} > 1/2$.

A perhaps natural attempt to fix the reduction is to put the two coherent mask registers in the entangled state

$$\frac{1}{2^{n/2}}\sum_r |r\rangle^R|r\rangle^S.$$

This forces Bob's and Charlie's GL extractor calls to use the same internal mask. The problem is that now the final Hadamard transforms can only "see" the *difference* between the two output strings: indeed, for an output pair (y, z), the two Fourier phases multiply to

$$(-1)^{\langle r,y\rangle}(-1)^{\langle r,z\rangle} = (-1)^{\langle r,y\oplus z\rangle}.$$

Thus, for the desired pair $(y, z) = (m, m)$, the phase is 1 for every r, but the same is true for every equal pair (y, y), so no Fourier phase remains that distinguishes (m, m) from the other equal pairs.

The reduction from $\mathsf{GL}(G)$ to G. We now present the new reduction, taking a strategy for $\mathsf{GL}(G)$ to a strategy for G. We will then spend the rest of the technical overview explaining why the reduction works. Fix a known parameter $0 < \gamma \le \Delta$ (i.e. some known lower bound on Δ), and set $\lambda_\gamma := \frac{1}{1+2\gamma}$. The reduction is as follows.

Reduction 1 (Common-mask GL reduction)**.**

1. Keep the splitting procedure $\mathcal{A}$ unchanged.
2. Bob samples a string L according to:

   $$\Pr[L = \ell] = (1-\lambda_\gamma)\lambda_\gamma^\ell, \qquad \ell = 0, 1, 2, \ldots .$$

   Then, he samples a list of L independent and uniformly random masks $T = (t_1, \ldots, t_L)$. After learning k, Bob applies

   $$W^B_{k,T} := \begin{cases} I_B, & L = 0, \\ V^B_{k,t_L} \cdots V^B_{k,t_1}, & L \ge 1. \end{cases}$$

   He then applies the single-party Goldreich-Levin extractor GL^B_k using a fresh register R initialized to $|0^n\rangle$, measures R, and outputs the resulting string m_B.
3. Charlie just applies GL^C_k to a fresh register S initialized to $|0^n\rangle$, measures S, and outputs the resulting string m_C. He does not use T or apply any other operation.

The main idea: a polynomial "filter" to align Bob and Charlie's individual correct-output branches. *What is this mysterious reduction achieving?* The only difference from the reduction in the independent-mask setting is that Bob now applies the random sequence of unitaries

$$W^B_{k,T} = V^B_{k,t_L} \cdots V^B_{k,t_1}$$

before running his Goldreich-Levin extractor. If we follow an analogous analysis as in the independent-mask setting for the new reduction, we see that the probability of successful extraction becomes

$$p_{\text{search}} = \mathbb{E}_T \left\| (\widehat{V}_k^B(m) W_{k,T}^B) \otimes \widehat{V}_k^C(m) |\psi_{k,m}\rangle \right\|^2 .$$

(the only change is the insertion of $W_{k,T}^B$ and the averaging over T). Now, there is one degree of freedom that turns out to be crucial in our analysis. Observe that we can rewrite the latter as

$$\mathbb{E}_T \left\| (\widehat{V}_k^B(m) W_{k,T}^B) \otimes (U_T^C \widehat{V}_k^C(m)) |\psi_{k,m}\rangle \right\|^2 ,$$

for any arbitrary family of unitaries U_T^C acting only on register C (since these do not affect the norm). From here, as in the previous analysis, we can lower-bound the expression, again via Cauchy-Schwarz, by

$$\mathbb{E}_T \left| \langle b_{k,m} \mid W_{k,T}^B \otimes U_T^C \mid c_{k,m} \rangle \right|^2 .$$

where $|b_{k,m}\rangle$ and $|c_{k,m}\rangle$ are as before. Comparing this expression to the plain overlap $|\langle b_{k,m} \mid c_{k,m}\rangle|^2$, the new overlap is *filtered* by $W_{k,T}^B \otimes U_T^C$. The nice thing is that we now get to pick any U_T^C of our choice to help us relate the expression to the success probability in the common-mask game. It turns out, as we will attempt to justify shortly, that one good choice of U_T^C is nothing other than $U_T^C = W_{k,T}^C$!

We can now further lower bound the last expression by Jensen's inequality, and reduce the problem to analyzing

$$\langle b_{k,m} | \mathbb{E}_T (W_{k,T}^B \otimes W_{k,T}^C) | c_{k,m} \rangle =: \langle b_{k,m} | F_{\gamma,k} | c_{k,m} \rangle$$

(and relating the latter to the winning probability in the original decision game).

Now, for fixed k, γ, the "filter" operator $F_{\gamma,k}$ can actually be understood as a *polynomial*. Define

$$\mathsf{J}_{k,t} := V_{k,t}^B \otimes V_{k,t}^C, \qquad \mathsf{J}_k := \mathop{\mathbb{E}}_t [\mathsf{J}_{k,t}].$$

Then, notice that, since the masks in the sequence T are sampled independently of each other,

$$F_{\gamma,k} := (1 - \lambda_\gamma) \sum_{\ell \geq 0} \lambda_\gamma^\ell \mathsf{J}_k^\ell .$$

where this power series is well-defined since $\|J_k\|_{op} \leq 1$.

Crucially, the "shape" of the polynomial is determined by the distribution over L (the length of the list T). Then, the name of the game is to try pick the right polynomial so that the resulting filter $F_{\gamma,k}$ does a good job at aligning the vectors $|b_{k,m}\rangle$ and $F_{\gamma,k}|c_{k,m}\rangle$. The filter will do so by applying the appropriate function to the eigenvectors of J_k. In our case, the power series is precisely the Taylor series of an *inverse* function:

$$(1 - \lambda_\gamma) \sum_{\ell \geq 0} \lambda_\gamma^\ell \mathsf{J}_k^\ell = (1 - \lambda_\gamma)(I - \lambda_\gamma \mathsf{J}_k)^{-1} .$$

Now, when one is not worried about efficiency of the reduction, there is no issue with having an infinite series (which simply corresponds to arbitrarily long sequences of unitaries applied by Bob). When efficiency of the reduction is a concern, one has to necessarily truncate the polynomial appropriately (since efficiency is determined by the degree of the polynomial), and the hope is to find a good enough low-degree approximation of the desired filter. This idea is reminiscent of the design of quantum algorithms via the QSVT framework [GSLW19], or, even more so, of constructions of "approximate ground state projectors" (AGSP) as polynomials in the Hamiltonian terms, e.g. in [ALV12; AAG22]. Below, we do our best to explain why the filter resulting from the particular choice of a geometric distribution succeeds.

Relating the "filtered" overlap to the advantage Δ. Note that a good choice of filter should not only make the overlap $\langle b_{k,m}|F_{\gamma,k}|c_{k,m}\rangle$ as large as possible, but it should also allow us to relate this overlap to the advantage Δ in $\mathsf{GL}(G)$. Eventually, we show that our filter $F_{\gamma,k}$ satisfies:

$$|\mathbb{E}_{k,m}\langle b_{k,m}|F_{\gamma,k}|c_{k,m}\rangle| \geq \gamma\Delta\,,$$

which, in turn, implies $p_{\text{search}} \geq (\gamma\Delta)^2$, based on the previous argument. Since we can run the reduction with any $0 < \gamma \leq \Delta$, this ultimately shows $p_{\text{search}} \geq \Delta^4$. So, our first goal will be to rewrite the overlap in a form that is more amenable to being related to Δ.

Fix k, m for now. For readability, we drop the k, m subscripts for now (we will reintroduce them later!). Denote $|b\rangle = |b_{k,m}\rangle$ and $|c\rangle = |c_{k,m}\rangle$ and $F = F_{\gamma,k}$. Define

$$|g\rangle := \frac{|b\rangle + |c\rangle}{2}, \qquad |h\rangle := \frac{|b\rangle - |c\rangle}{2}.$$

Then $|b\rangle = |g\rangle + |h\rangle$ and $|c\rangle = |g\rangle - |h\rangle$. A first convenient observation is the following. We can rewrite

$$\langle b|F|c\rangle = \langle g|F|g\rangle - \langle h|F|h\rangle + \langle h|F|g\rangle - \langle g|F|h\rangle\,,$$

and since the last two terms are purely imaginary, we have

$$\mathrm{Re}\langle b|F|c\rangle = \langle g|F|g\rangle - \langle h|F|h\rangle.$$

Since $|\langle b|F|c\rangle| \geq \mathrm{Re}\langle b|F|c\rangle$, it will suffice for us to bound the latter expression. Eventually, the right choice of F will have to accomplish the following: increase the relative contribution of $\langle g|F|g\rangle$ compared to $\langle h|F|h\rangle$. Now, do these overlaps have any *operational* meaning? Indeed, they do.

Recall that

$$\begin{aligned}|b\rangle = |b_{k,m}\rangle &= \widehat{V}_k^B(m) \otimes I_C|\psi_{k,m}\rangle =: \mathsf{B}|\psi\rangle \ \text{, and}\\ |c\rangle = |c_{k,m}\rangle &= I_B \otimes \widehat{V}_k^C(m)|\psi_{k,m}\rangle =: \mathsf{C}|\psi\rangle\end{aligned}$$

Let us introduce the operators

$$S_{\text{sum}} := \frac{1}{2}(\mathsf{B} + \mathsf{C}), \qquad R_{\text{diff}} := \frac{1}{2}(\mathsf{B} - \mathsf{C}).$$

Then, we have $|g\rangle = S_{\text{sum}}|\psi\rangle$ and $|h\rangle = R_{\text{diff}}|\psi\rangle$, and thus

$$\mathrm{Re}\langle b|F|c\rangle = \langle\psi|S_{\text{sum}}FS_{\text{sum}}|\psi\rangle - \langle\psi|R_{\text{diff}}FR_{\text{diff}}|\psi\rangle. \tag{2.6}$$

Observe that the R_{diff} operator captures the difference between the post-measurement states conditioned on Bob (resp. Charlie) correctly guessing m and Charlie (resp. Bob) doing nothing. Concretely, $\|R_{\text{diff}}|\psi\rangle\|$ is precisely this distance. It turns out that R_{diff} is a good proxy for the *disagreement* between Bob and Charlie, in a formal sense. Recall that we defined

$$\mathsf{J}_{k,t} := V_{k,t}^B \otimes V_{k,t}^C, \qquad \mathsf{J}_k := \mathop{\mathbb{E}}_t[\mathsf{J}_{k,t}].$$

Note that $\mathsf{J}_{k,t}$ is an observable that precisely measures the "agreement" between Bob and Charlie in $\mathsf{GL}(G)$ (on key k and mask t). This is because the observable has eigenvalue $+1$ on the eigenspace where Bob and Charlie return the same answer, and -1 where they return different answers. We can then define the operator $D_k = \frac{1}{2}(I - \mathsf{J}_k)$, which captures *disagreement*. It turns out that one can show:

$$-D_k \preceq R_{\text{diff}} \preceq D_k\,,$$

i.e. R_{diff} is dominated by the disagreement operator D_k. On the other hand, S_{sum} does not quite capture "agreement", but it simply captures the sum of Bob and Charlie's individual correct-output probabilities. So, at a high level, the filter F should accomplish the following: suppress the negative contribution to the overlap coming from the disagreement, while maintaining enough positive contribution from Bob and Charlie's individual success probabilities.

We now reintroduce the k, m dependence for clarity and write $S^{k,m}_{\text{sum}}$ and $R^{k,m}_{\text{diff}}$. We also introduce the following compact notation:

$$\langle X \rangle_\pi := \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \langle \psi_{k,m} | X | \psi_{k,m} \rangle .$$

For example,

$$\langle S_{\text{sum}} \rangle_\pi := \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \langle \psi_{k,m} | S^{k,m}_{\text{sum}} | \psi_{k,m} \rangle .$$

Now, a crucial observation is that, together, S_{sum} and D_k are sufficient to determine the advantage Δ. This observation is what lets us connect $\text{Re}\langle b|F|c\rangle$ to the advantage Δ. Rewriting Equation (2.5) as

$$\begin{aligned}
p_{\text{Pred}} &= \frac{1}{4} \mathop{\mathbb{E}}_{\substack{(k,m)\leftarrow\pi \\ r\leftarrow\{0,1\}^n}} \left[\langle \psi_{k,m} | \left(\frac{I_B + (-1)^{\langle r,m \rangle} V^B_{k,r}}{2} \otimes \frac{I_C + (-1)^{\langle r,m \rangle} V^C_{k,r}}{2} \right) | \psi_{k,m} \rangle \right] \\
&= \frac{1}{4} \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \left[\langle \psi_{k,m} | \left(I_{BC} + \hat{V}^B_k(m) \otimes I_C + I_B \otimes \hat{V}^C_k(m) + \mathop{\mathbb{E}}_{r\leftarrow\{0,1\}^n} V^B_{k,r} \otimes V^C_{k,r} \right) | \psi_{k,m} \rangle \right] \\
&= \frac{1}{4} \langle I_{BC} + \mathsf{B}_{k,m} + \mathsf{C}_{k,m} + \mathsf{J}_k \rangle_\pi
\end{aligned}$$

and using $S^{k,m}_{\text{sum}} = \frac{1}{2}(\mathsf{B}_{k,m} + \mathsf{C}_{k,m})$, $D_k = \frac{1}{2}(I - \mathsf{J}_k)$ and $\Delta = 2p_{\text{Pred}} - 1$, we have

$$\Delta = \langle S_{\text{sum}} \rangle_\pi - \langle D \rangle_\pi . \tag{2.7}$$

While the above is a straightforward calculation, it establishes a somewhat surprising fact that is worth highlighting: it relates Bob and Charlie's advantage Δ in the decision game $\mathsf{GL}(G)$ to: their "disagreement" in $\mathsf{GL}(G)$, captured by D, and the sum of their individual *correct-extraction* probabilities when running a local Goldreich-Levin extractor, captured by S_{sum}.

Notice now that Equation (2.7) is quite similar to (2.6), when we rewrite the latter using the new notation we introduced (and average over k, m):

$$\text{Re}\langle \mathsf{B}F\mathsf{C} \rangle_\pi = \langle S_{\text{sum}} F S_{\text{sum}} \rangle_\pi - \langle R_{\text{diff}} F R_{\text{diff}} \rangle_\pi . \tag{2.8}$$

This suggests a proof strategy: choose a filter F such that

1. $\langle R_{\text{diff}} F R_{\text{diff}} \rangle_\pi \leq \alpha \langle D \rangle_\pi$, and
2. $\langle S_{\text{sum}} F S_{\text{sum}} \rangle_\pi \geq \alpha \langle S_{\text{sum}} \rangle_\pi$

for some factor $\alpha > 0$ that is not too small. We will show that one can obtain $\alpha = \gamma$, which implies the desired bound $\text{Re}\langle \mathsf{B}F\mathsf{C} \rangle_\pi \geq \gamma\Delta$.

Why the chosen filter works. Recall that we defined $\lambda_\gamma := \frac{1}{1+2\gamma}$, and that our reduction uses the particular filter

$$\begin{aligned} F_{\gamma,k} &= (1-\lambda_\gamma)\sum_{\ell\geq 0}(\lambda_\gamma \mathsf{J}_k)^\ell \\ &= (1-\lambda_\gamma)(I-\lambda_\gamma \mathsf{J}_k)^{-1} \\ &= f(D_k)\,, \end{aligned} \tag{2.9}$$

where a straightforward calculation gives $f(d) = \frac{\gamma}{d+\gamma}$. In other words, $F_{\gamma,k}$ is diagonal in an eigenbasis of D_k, and its diagonal elements are obtained by applying f to the eigenvalues of D_k. Note what such a filter F achieves: it is close to the identity on components with little disagreement ($d < \gamma$) but suppresses, inversely in d, components with large disagreement ($d >> \gamma$). In this sense, F is *approximately a projection onto the low-disagreement* component of the state.

The $R_{\text{diff}} F_{\gamma,k} R_{\text{diff}}$ bound. We aim to show

$$\langle R_{\text{diff}} F R_{\text{diff}} \rangle_\pi \leq \gamma \langle D \rangle_\pi\,.$$

Recall that R_{diff} is dominated by the disagreement operator D_k:

$$-D_k \preceq R_{\text{diff}} \preceq D_k,$$

Assume for simplicity that R_{diff} and D_k commute, and thus share an eigenbasis. Then, we have

$$R_{\text{diff}} F_{\gamma,k} R_{\text{diff}} \preceq D_k F_{\gamma,k} D_k\,.$$

(since F is positive and also diagonal in any eigenbasis of D). From here, one can immediately see that

$$D_k F_{\gamma,k} D_k \preceq \gamma D_k\,,$$

since any component with eigenvalue d gets transformed to $\frac{\gamma d^2}{d+\gamma}$ by the LHS, which is less than γd. Thus, this gets us the desired bound. A short block-matrix argument (which we defer to the main body) makes this intuition valid even when R_{diff} does not commute with D_k.

The $S_{\text{sum}} F S_{\text{sum}}$ bound. We are now left with showing

$$\langle S_{\text{sum}} F S_{\text{sum}} \rangle_\pi \geq \gamma \langle S_{\text{sum}} \rangle_\pi\,.$$

Suppressing k, m for a moment, let $|x\rangle = S_{\text{sum}}|\psi\rangle$. The difficulty is that the LHS overlap $\langle x|F|x\rangle$ may be very small due to the filter, even though the RHS overlap $\langle \psi|x\rangle = \langle \psi|S_{\text{sum}}|\psi\rangle$ is large. This can happen if $|x\rangle$ has a lot of its weight on eigenspaces of F corresponding to small eigenvalues. More concretely, if F has a tiny eigenvalue μ on some component, then that component can contribute to $\langle x|F|x\rangle$ at a "discount rate" of μ. More generally, among all vectors $|x\rangle$ having a fixed overlap with $|\psi\rangle$, the vector with smallest contribution to the filtered overlap points in the direction $F_{\gamma,k}^{-1}|\psi\rangle$. Formally, this follows from a Cauchy-Schwarz inequality:

$$\left\langle S_{\text{sum}} F_{\gamma,k} S_{\text{sum}} \right\rangle_\pi \geq \frac{\left|\langle S_{\text{sum}} \rangle_\pi\right|^2}{\langle F_{\gamma,k}^{-1} \rangle_\pi}.$$

Now, by Equation (2.7), the numerator on the RHS is $(\Delta + d)^2$, where $d = \langle D_k \rangle_\pi$. On the other hand, by Equation (2.9), the filter is designed so that the cost of applying the inverse filter is exactly linear in the disagreement d (up to a rescaling by γ):

$$\langle F_{\gamma,k}^{-1} \rangle_\pi = \frac{\gamma + \langle D_k \rangle_\pi}{\gamma} = \frac{\gamma + d}{\gamma}.$$

Using $\gamma \le \Delta$, we get precisely

$$\langle S_{\mathrm{sum}} F_{\gamma,k} S_{\mathrm{sum}} \rangle_\pi \ge \gamma(\Delta + d) = \gamma \langle S_{\mathrm{sum}} \rangle_\pi .$$

as desired.

Finally, putting our $R_{\mathrm{diff}} F_{\gamma,k} R_{\mathrm{diff}}$ and $S_{\mathrm{sum}} F S_{\mathrm{sum}}$ bounds together with (2.7) and (2.8) yields the desired bound $\mathrm{Re}\langle \mathsf{B} F \mathsf{C} \rangle_\pi \ge \gamma\Delta$. This, in turn gives $p_{\mathrm{search}} \ge \Delta^4$ (choosing $\gamma = \Delta$ in the reduction).

As a final remark, notice that the filter achieves a very delicate balance: it suppresses high-disagreement components enough to control the negative contribution from $R_{\mathrm{diff}} F_{\gamma,k} R_{\mathrm{diff}}$, while its inverse grows only linearly with disagreement, which is just what is needed to have enough of the positive $S_{\mathrm{sum}} F S_{\mathrm{sum}}$ contribution survive.

Relation to the filter used by Ananth and Sahai [AS26]. The filter used in our reduction is analogous to a filter used by Ananth and Sahai [AS26]. Both filters have the geometric series form

$$J = (1 - \lambda_\gamma) \sum_{\ell=0}^{\infty} \lambda_\gamma^\ell \mathsf{J}^\ell = (1 - \lambda_\gamma)(I - \lambda_\gamma \mathsf{J})^{-1} ,$$

but they are used in different ways in their respective analyses. Ananth and Sahai prove that, for every integer $\ell \ge 0$,

$$\left\| \mathsf{E}_B \mathsf{J}^\ell \mathsf{E}_C \right\|_{\mathrm{op}} \le \left\| \mathsf{E}_B \mathsf{E}_C \right\|_{\mathrm{op}} ,$$

where E_B and E_B are operators capturing Bob and Charlie's individual success probabilities. They then use this inequality to bound certain "agreement moments" that arise in their analysis.

For us, the filter has a clear operational meaning. It arises from Bob's actual sequence of unitaries $W_{k,T}^B = V_{k,t_L} \cdots V_{k,t_1}$ in our Goldreich-Levin reduction (Reduction 1) (after inserting matching operations on Charlie's side for the purpose of the analysis).

3 The model and main result

We start by formulating "search" security for an unclonable encryption scheme. This can be modeled generically as a "cloning game" as in [AKL23].

Definition 3.1 ("Search" unclonable encryption game)**.** *Fix an integer $n \ge 1$ (which can be understood as the security parameter), a finite key space $\mathcal{K}$, and a distribution π on $\mathcal{K} \times \{0,1\}^n$.*

- *(sampling) The challenger samples $(k, m) \leftarrow \pi$.*
- *(token state generation) For every (k, m), a challenger prepares a token state $\rho_{k,m}^A$ on a finite-dimensional register A.*
- *(splitting) The adversary applies a fixed cloning/splitting procedure $\mathcal{A}$ producing*

$$\sigma_{k,m}^{BC} := \mathcal{A}(\rho_{k,m}^A).$$

The state $\sigma_{k,m}^{BC}$ may in general be mixed and arbitrarily entangled across B and C. We think of registers B and C as going to two parties Bob and Charlie respectively. In the rest of the paper, we will take the state to be pure, and denote it as $|\psi_{k,m}\rangle^{BC}$. This is without loss of generality, since $\mathcal{A}$ can always provide any purifying register to one of the two parties (and this can only increase their success probability).

- *(Key is revealed and Bob and Charlie guess) The challenger reveals the key k to Bob and Charlie, who respectively apply a local measurement on B and C to obtain a guess for m (this measurement can be taken to be projective without loss of generality since $\mathcal{A}$ can provide any required auxiliary register).*

The game is won if both *Bob and Charlie's guesses equal m.*

Remark 3.2. *The simplest example of such a game is based on BB84 states. There, k encodes a basis $\{0,1\}^n$ and $\rho^A_{k,m}$ is the state encoding m under the basis k: $\bigotimes_i H^{k_i}|m_i\rangle$. The resulting game, considered by Broadbent and Lord in [BL20] in the setting of unclonable encryption, is "secure" (in the sense of having an exponentially small optimal winning probability). The latter follows from a monogamy-of-entanglement property of BB84 states established in [TFKW13].*

We now define a Goldreich-Levin mapping from a search game G as above to a decision game $\mathsf{GL}(G)$. The latter corresponds to the following mapping of an unclonable encryption scheme that satisfies "search" security into one that satisfies "indistinguishability" security. To encrypt the bit b: use the search secure scheme to encrypt a uniformly random string $m \in \{0,1\}^n$ via the ciphertext $|\mathsf{Enc}(k,m)\rangle$; sample another uniformly random "mask" $r \in \{0,1\}^n$; the new ciphertext is $(|\mathsf{Enc}(k,m)\rangle, r, \langle r,m\rangle \oplus b)$ (i.e. use $\langle r,m\rangle$ as a one-time pad). Although this transformation is typical in classical (and quantum) cryptography, its validity for unclonable encryption games has remained a conjecture since [AKLL+22]. As discussed, the primary source of difficulty and barrier is that, in the corresponding decision game, the "mask" r is identical for Bob and Charlie (since it is part of the original ciphertext), rather than Bob and Charlie having their own independently sampled masks.

Definition 3.3 ("Common-mask" Goldreich-Levin unclonable encryption game)**.** *Let G be a search unclonable-encryption game as in Definition 3.1. We define the game $\mathsf{GL}(G)$ as follows:*

- *The sampling, token state generation, and splitting phases are identical as in G.*
- *The only difference is in the final phase: the challenger additionally samples a uniformly random $r \in \{0,1\}^n$, and reveals k and r to Bob and Charlie. Then, Bob and Charlie each return a guess for $\langle r,m\rangle$. The game is won if both guesses are correct.*

We refer to r as the "mask".

The main result of this work is that the "Common-mask" Goldreich–Levin search-to-decision reduction works for any unclonable encryption game.

Theorem 3.4. *The following holds for any search unclonable encryption game G (as in Definition 3.1). Let p_{Search} and p_{Pred} be the optimal success probabilities in G and $\mathsf{GL}(G)$ respectively (where $\mathsf{GL}(G)$ is as in Definition 3.3). Then,*

$$p_{\text{Search}} \geq (2p_{\text{Pred}} - 1)^4.$$

Note that $2p_{\text{Pred}} - 1$ is (twice) the optimal "advantage" over $\frac{1}{2}$ (i.e. random guessing) in $\mathsf{GL}(G)$.

Moreover, there is an explicit reduction that takes any strategy for $\mathsf{GL}(G)$ that wins with probability $\frac{1}{2} + \frac{1}{2}\Delta$ (for $\Delta \geq 0$) to a strategy for G that wins with probability $\gamma^2\Delta^2$, where γ is any known lower bound on Δ. The reduction is uniform; and takes expected polynomial time in the runtime of the strategy for $\mathsf{GL}(G)$ and $1/\gamma$.

The runtime of the reduction can be made *worst-case* polynomial-time at the cost of a factor of $\frac{1}{2}$ loss in the p_{Search} lower bound (see Theorem 8.2). Moreover, there is a reduction that does not need to know a lower bound γ at all, and achieves $p_{\text{Search}} \geq \Omega(\Delta^6)$ (see Theorem 8.3).

Going forward, for convenience, we will often refer to Δ as the advantage of a strategy, even though it is technically twice the advantage.

The rest of the paper is devoted to the proof of Theorem 3.4.

4 A warmup: the independent-mask Goldreich–Levin reduction

In this section, we consider a simpler version of the Goldreich-Levin unclonable encryption game from Definition 3.3. The only difference will be that here Bob and Charlie receive *independently* sampled masks r and r' (rather than identical) and will have to respectively guess $\langle r, m\rangle$ and $\langle r', m\rangle$.

This "independent-mask" version of the game has been analyzed before. The analysis follows straightfowardly from the standard quantum Goldreich-Levin reduction of Adcock and Cleve [AC02], as observed in, e.g. [CLLZ21; AKL23]. We recall this analysis here in broad strokes as it will help us introduce some relevant notation, as well as pinpoint the crux in analyzing to the common-mask version of the game.

4.1 Setup

All inner products are over $\mathbb{F}_2$, and we write

$$\chi_m(r) := (-1)^{\langle r,m\rangle}, \qquad r \in \{0,1\}^n.$$

For the rest of this section, we fix the strategy for the "splitting" phase. For k and $m \in \{0,1\}^n$ sampled by the challenger, let $|\psi_{k,m}\rangle$ denote the joint state of Bob and Charlie after the splitting phase (which we are assuming is pure without loss of generality). The reduction we will describe does not change the splitting phase at all, only Bob and Charlie's operations.

For key k and mask $r \in \{0,1\}^n$, let $V^B_{k,r}$ and $V^C_{k,r}$ be the binary observables measured by Bob and Charlie respectively to computer their answer bit. Then, their corresponding projective measurements are $\{\Pi^B_{k,r,b}\}_{b\in\{0,1\}}$ and $\{\Pi^C_{k,r,b}\}_{b\in\{0,1\}}$ where

$$\Pi^B_{k,r,b} := \frac{I + (-1)^b V^B_{k,r}}{2}, \qquad \Pi^C_{k,r,b} := \frac{I + (-1)^b V^C_{k,r}}{2}.$$

The "correct answer" projectors are

$$\Pi^B_{k,r,\langle r,m\rangle} = \frac{I + \chi_m(r) V^B_{k,r}}{2}, \qquad \Pi^C_{k,r,\langle r,m\rangle} = \frac{I + \chi_m(r) V^C_{k,r}}{2}.$$

The winning probability for this strategy is the following.

Definition 4.1 (Independent-mask unclonable-encryption game winning probability)**.**

$$p^{\text{ind}}_{\text{Pred}} := \mathop{\mathbb{E}}_{\substack{(k,m)\leftarrow\pi \\ r,s\leftarrow\{0,1\}^n}} \left[\langle\psi_{k,m}| \left(\frac{I + \chi_m(r) V^B_{k,r}}{2} \otimes \frac{I + \chi_m(s) V^C_{k,s}}{2} \right) |\psi_{k,m}\rangle \right]$$

(where the expectation is over $r, s \in \{0,1\}^n$).

Note that a winning probability of 1/2 is trivially achievable: give the ciphertext to Bob and let Charlie do a random guess. We write

$$\Delta^{\text{ind}} := 2p^{\text{ind}}_{\text{Pred}} - 1$$

for the advantage over this baseline. We henceforth consider only the case $\Delta^{\text{ind}} \geq 0$.

4.2 Independent-mask GL reduction

We now describe a reduction that takes a strategy achieving advantage Δ^{ind} in the independent-mask unclonable-encryption game, and obtains a strategy that wins with probability at least $(\Delta^{\text{ind}})^2$ in the "search" unclonable-encryption game of Definition 3.1, i.e. Bob and Charlie simultaneously extract m with probability at least $(\Delta^{\text{ind}})^2$. The reduction is simple: Bob and Charlie each locally run the [AC02] Goldreich-Levin extractor, and return the output. We recall how this extractor works.

The [AC02] Goldreich-Levin extractor. Let R be an n-qubit register. For a key k, define Bob's controlled reflection

$$\mathsf{CV}_k^B := \sum_{r\in\{0,1\}^n} |r\rangle\langle r|^R \otimes V_{k,r}^B \,.$$

Then, the [AC02] Goldreich–Levin extractor applies the following unitary

$$\mathsf{GL}_k^B := (H_R^{\otimes n} \otimes I_B)\mathsf{CV}_k^B(H_R^{\otimes n} \otimes I_B)\,, \tag{4.1}$$

followed by a measurement of the R register. The guarantee proven in [AC02], which follows from a straightforward calculation, is that, for any state $|\phi\rangle^B$ which provides advantage ϵ at guessing $\langle r, m\rangle$, the extractor above will output m with probability at least ϵ^2 when run on $|0\rangle^R|\phi\rangle^B$. In the independent-mask GL reduction, Bob and Charlie both locally run the [AC02] extractor in parallel. So, let us define CV_k^C and GL_k^C analogously.

Now, for $y \in \{0,1\}^n$, define

$$\widehat{V}_k^B(y) := 2^{-n}\sum_r (-1)^{\langle r,y\rangle} V_{k,r}^B, \qquad \widehat{V}_k^C(y) := 2^{-n}\sum_r (-1)^{\langle r,y\rangle} V_{k,r}^C. \tag{4.2}$$

The following lemma will be important later.

Lemma 4.2. *For every state $|\phi\rangle^B$,*

$$\mathsf{GL}_k^B(|0^n\rangle^R|\phi\rangle^B) = \sum_{y\in\{0,1\}^n} |y\rangle^R \widehat{V}_k^B(y)|\phi\rangle^B.$$

Equivalently, the following operator identity holds for all y:

$$\widehat{V}_k^B(y). = (\langle y|^R \otimes I_B)\mathsf{GL}_k^B(|0^n\rangle^R \otimes I_B)\,.$$

The analogous identities hold for Charlie.

Proof. Fix an arbitrary input state $|\phi\rangle^B$. Then,

$$\begin{aligned}
\mathsf{GL}_k^B(|0^n\rangle^R|\phi\rangle^B) &= (H_R^{\otimes n} \otimes I_B)\mathsf{CV}_k^B(H_R^{\otimes n} \otimes I_B) \\
&= 2^{-n} \sum_{r,y\in\{0,1\}^n} (-1)^{\langle r,y\rangle} |y\rangle^R V_{k,r}^B |\phi\rangle^B \\
&= \sum_{y\in\{0,1\}^n} |y\rangle^R \widehat{V}_k^B(y)|\phi\rangle^B
\end{aligned}$$

The sums are finite, so we may group together all terms with the same value of y:

$$\mathsf{GL}_k^B(|0^n\rangle^R|\phi\rangle^B) = \sum_{y\in\{0,1\}^n} |y\rangle^R \left(2^{-n} \sum_{r\in\{0,1\}^n} (-1)^{\langle r,y\rangle} V_{k,r}^B\right) |\phi\rangle^B$$

$$= \sum_{y\in\{0,1\}^n} |y\rangle^R \widehat{V}_k^B(y)|\phi\rangle^B .$$

Finally, projecting the R register onto $\langle y|$, and noting that the resulting identity holds for all $|\phi\rangle$, yields the operator identity

$$(\langle y|^R \otimes I_B)\mathsf{GL}_k^B(|0^n\rangle^R \otimes I_B) = \widehat{V}_k^B(y) .$$

The argument is analogous on Charlie's side. □

The independent-mask GL reduction. As mentioned earlier, the independent-mask Goldreich-Levin reduction is simple: Bob and Charlie each locally run the [AC02] Goldreich-Levin extractor.

Proposition 4.3 (Independent-mask GL reduction)**.** *Let $p_{\text{Pred}}^{\text{ind}}$ be the probability that Bob and Charlie win the independent-mask game (as in Definition 4.1). Assume $p_{\text{Pred}}^{\text{ind}} \geq 1/2$ and set $\Delta^{\text{ind}} := 2p_{\text{Pred}}^{\text{ind}} - 1$. If Bob and Charlie each run the Goldreich-Levin extractor locally, they both output the correct m with probability at least*

$$p_{\text{Search}} \geq (\Delta^{\text{ind}})^2 = \left(2p_{\text{Pred}}^{\text{ind}} - 1\right)^2.$$

Proof. A direct consequence of Lemma 4.2 is that, upon running their respective Goldreich-Levin extractors and obtaining the same correct outcome m, the unnormalized leftover state on registers B and C is

$$\widehat{V}_k^B(m) \otimes \widehat{V}_k^C(m)|\psi_{k,m}\rangle .$$

Consequently, the probability that Bob and Charlie output the same correct outcome m given key k in the first place is precisely

$$\begin{aligned} p_{k,m} &:= \left\| \widehat{V}_k^B(m) \otimes \widehat{V}_k^C(m)|\psi_{k,m}\rangle \right\|^2 \\ &\geq \left| \langle\psi_{k,m}| \widehat{V}_k^B(m) \otimes \widehat{V}_k^C(m)|\psi_{k,m}\rangle \right|^2 \\ &= |\langle b_{k,m} \mid c_{k,m}\rangle|^2 . \end{aligned} \tag{4.3}$$

where the first inequality is by Cauchy–Schwarz, and in the last equality we defined

$$|b_{k,m}\rangle := \widehat{V}_k^B(m) \otimes I|\psi_{k,m}\rangle, \qquad |c_{k,m}\rangle := I \otimes \widehat{V}_k^C(m)|\psi_{k,m}\rangle.$$

Let X and Y denote the ± 1-valued random variables for Bob and Charlie respectively guessing their own inner product correctly (where $+1$ denotes a correct guess). Then, notice that for fixed k, m, r, s, we have

$$\mathbb{E}[XY \mid k, m, r, s] = \chi_m(r)\chi_m(s)\langle\psi_{k,m}|V_{k,r}^B \otimes V_{k,s}^C|\psi_{k,m}\rangle.$$

Using the definitions of $\widehat{V}_k^B(m)$ and $\widehat{V}_k^C(m)$ from (4.2), we have

$$\begin{aligned} \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \langle b_{k,m} \mid c_{k,m}\rangle &= \mathop{\mathbb{E}}_{\substack{(k,m)\leftarrow\pi \\ r,s}} \left[\chi_m(r)\chi_m(s)\langle\psi_{k,m}|V_{k,r}^B \otimes V_{k,s}^C|\psi_{k,m}\rangle\right] \\ &= \mathbb{E}[XY] \end{aligned} \tag{4.4}$$

Now, let $P_{u,v} = \Pr[X = u, Y = v]$. By definition, $P_{+1,+1} = p_{\text{Pred}}^{\text{ind}}$, and therefore

$$\mathbb{E}[XY] = P_{+1,+1} + P_{-1,-1} - P_{+1,-1} - P_{-1,+1}$$

$$= 2(P_{+1,+1} + P_{-1,-1}) - 1$$
$$\geq 2p^{\text{ind}}_{\text{Pred}} - 1 = \Delta^{\text{ind}}. \tag{4.5}$$

Finally, averaging Equation (4.3) over $(k, m) \leftarrow \pi$, applying Jensen's inequality and Equations (4.4) and (4.5), we have

$$p_{\text{Search}} \geq \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} |\langle b_{k,m} \mid c_{k,m}\rangle|^2$$
$$\geq \left| \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \langle b_{k,m} \mid c_{k,m}\rangle \right|^2 \geq (\Delta^{\text{ind}})^2 ,$$

as desired. □

What fails in the common-mask setting. We draw the reader's attention to the key assumption that made the above argument go through: the challenges r and s are drawn independently (and $(\Delta^{\text{ind}})^2$ is the probability that they succeed simultaneously on this independent-challenge distribution). This is what allows us to relate the overlap $\langle b_{k,m} \mid c_{k,m}\rangle$ to $\mathbb{E}[XY|k, m]$ in Equation (4.4). In the "common-mask" game, which we will discuss next, the hypothesis is that Bob and Charlie have advantage Δ when receiving *identical* challenges. Thus, the success hypothesis only controls the diagonal terms $r = s$, and therefore need not bound the same overlap. In this case, what can happen is that $|b_{k,m}\rangle$ and $|c_{k,m}\rangle$ can each have substantial norm, but their overlap is tiny or even zero. In Appendix A, we give an explicit two-dimensional example in which the common-mask simultaneous success probability is 25/48, yet $|b_{k,m}\rangle$ and $|c_{k,m}\rangle$ are orthogonal. To overcome this, we will have to devise a new reduction that works in the common-mask setting.

5 The common-mask GL reduction

For a search unclonable encryption game G, recall the definition of the common-mask unclonable encryption game $\mathsf{GL}(G)$ from Definition 3.3: the only difference from the independent-mask game discussed in Section 4 is that Bob and Charlie now receive identical challenges r. In this section, we will describe how to modify the reduction from the previous section so that it works in the common-mask setting. This will allows us to prove our main Theorem 3.4, relating the success probability in G to the success probability in $\mathsf{GL}(G)$.

5.1 The reduction algorithm

Fix a strategy for Bob and Charlie. We use the same notation as in Section 4 to describe it: for key k and mask $r \in \{0, 1\}^n$, we let $V^B_{k,r}$ and $V^C_{k,r}$ be the binary observables (i.e. reflections) that they measure to compute their answer bit. For key k and $m \in \{0, 1\}^n$, we denote by $|\psi_{k,m}\rangle$ the state on Bob and Charlie's registers after the splitting phase. As before, we denote by p_{Pred} the winning probability of this strategy in $\mathsf{GL}(G)$, and let $\Delta = 2p_{\text{Pred}} - 1$.

The reduction algorithm is "slightly" non-black-box. It needs to know a lower bound $0 < \gamma \leq \Delta$ on the advantage. The success probability p_{Search} of the resulting search strategy will then depend on γ: we will eventually show $p_{\text{Search}} \geq \gamma^2\Delta^2$.

Moreover, the reduction algorithm is actually (fundamentally) randomized. As part of the reduction, Bob samples a non-negative integer L according to the following geometric distribution. Let

$$\lambda_\gamma := \frac{1}{1 + 2\gamma}.$$

Then,

$$\Pr[L = \ell] = (1 - \lambda_\gamma)\lambda_\gamma^\ell, \qquad \ell = 0, 1, 2, \ldots. \tag{5.1}$$

Here is the reduction algorithm.

Reduction 2 (Common-mask GL reduction)**.**

Input: A strategy for $\mathsf{GL}(G)$ (for which we use the notation introduced above), and a parameter $0 < \gamma \leq \Delta$.

Output: The following strategy for G.

The splitting phase is identical as in $\mathsf{GL}(G)$, and results in Bob and Charlie obtaining a state $|\psi_{k,m}\rangle$ on registers B, C, for some key k and $m \in \{0,1\}^n$. Then:

1. Bob samples L from the distribution in (5.1). He then samples $T = (t_1, \ldots, t_\ell)$ where the $t_i \leftarrow \{0,1\}^n$ are independent and uniformly random masks.

2. After receiving k, Bob applies unitary

$$W^B_{k,T} = V^B_{k,t_\ell} \cdots V^B_{k,t_1},$$

to register B. Then Bob initializes a register R to $|0^n\rangle$, and apply the [AC02] Golreich-Levin unitary GL^B_k (as in Equation 4.1) on registers RB. Finally, he measures R in the standards basis, and returns the output.

3. After receiving k, Charlie initializes a register S to $|0^n\rangle$. He applies GL^C_k on registers SC, measures S in the computational basis, and returns the output. Note that, unlike Bob, Charlie does not apply any other unitary before GL^C_k.

We will spend most of the remainder of this paper proving the following theorem.

Theorem 5.1. *Let p_{Pred} be the winning probability of a strategy S for game $\mathsf{GL}(G)$. Let $\Delta = 2p_{\text{Pred}} - 1$, and assume $\Delta > 0$. Let $0 < \gamma \leq \Delta$, and let S' be the strategy for game G obtained by running Reduction 2 on input S and γ. Let p_{Search} be the winning probability of S' in G. Then,*

$$p_{\text{Search}} \geq \gamma^2 \Delta^2.$$

For a high level (but detailed) explanation of what this reduction is accomplishing, see the Technical Overview first (Section 2).

Remark 5.2 (GL reduction for monogamy-of-entanglement games)**.** *We remark that this reduction, and its analysis, are essentially unchanged for* monogamy-of-entanglement *games. In a monogamy game, it is $\mathcal{A}$ who first prepares a state on registers ABC, sending A to the challenger, and B and C to Bob and Charlie respectively. Then, the challenger measures A in a basis determined by the key k, obtaining an outcome m. Everything else is identical. One can then analogously define a decision monogamy game $\mathsf{GL}(G)$ from a search monogamy game G. The GL reduction for monogamy games is analogous to Reduction 2: keep $\mathcal{A}$ unchanged, and have Bob and Charlie act exactly as in Reduction 2. The analysis of this reduction is identical, and achieves the same guarantee, since the analysis only needs to consider the state on BC and the respective operations performed by Bob and Charlie on it.*

6 Analysis of the reduction, part 1: the "filtered" overlap

In this section, we begin the analysis of Reduction 2. We will formalize the high-level argument from the Technical Overview, lower-bounding the probability of successful joint extraction in terms of an appropriately "filtered" overlap. We fix a strategy for $\mathsf{GL}(G)$, and use the notation introduced at the start of Section 5 to describe it.

For a key k and a mask $t \in \{0,1\}^n$, define

$$\mathsf{J}_{k,t} = V^B_{k,t} \otimes V^C_{k,t}, \qquad \mathsf{J}_k = \mathop{\mathbb{E}}_t[\mathsf{J}_{k,t}]. \tag{6.1}$$

The operator J_k is Hermitian and satisfies $\|\mathsf{J}_k\|_{\mathrm{op}} \le 1$ since $V^B_{k,t}$ and $V^C_{k,t}$ are reflections. For a finite sequence $T = (t_1, \ldots, t_\ell)$, define the unitaries

$$W^B_{k,T} = V^B_{k,t_\ell} \cdots V^B_{k,t_1}, \qquad W^C_{k,T} = V^C_{k,t_\ell} \cdots V^C_{k,t_1}. \tag{6.2}$$

We let both unitaries be the identity when T is empty.

As before, we denote by p_{PRed} the winning probability of the strategy in $\mathsf{GL}(G)$, and let $\Delta = 2p_{\mathrm{PRed}} - 1$. For a reduction parameter $0 < \gamma \le \Delta$, recall that we defined $\lambda_\gamma := \frac{1}{1+2\gamma}$. Here, we also introduce the following operator, which will be crucial in our analysis:

$$F_{\gamma,k} := (1-\lambda_\gamma) \sum_{\ell=0}^{\infty} \lambda_\gamma^\ell \mathsf{J}_k^\ell,$$

where $\mathsf{J}_k^0 = I_{BC}$. Since $0 < \lambda_\gamma < 1$ and $\|\mathsf{J}_k\| \le 1$, the series converges in operator norm.

We have the following straightforward identities.

Lemma 6.1. *For $\ell \in \mathbb{N}$,*

$$\mathop{\mathbb{E}}_{T:\,|T|=\ell} \left[W^B_{k,T} \otimes W^C_{k,T} \right] = \mathsf{J}_k^\ell.$$

Consequently,

$$\mathop{\mathbb{E}}_T \left[W^B_{k,T} \otimes W^C_{k,T} \right] = F_{\gamma,k}.$$

Proof. Let $X_t := V^B_{k,t} \otimes V^C_{k,t}$. The definitions of the synchronized-sequence unitaries in Equation (6.2), including their matching order, give

$$W^B_{k,T} \otimes W^C_{k,T} = X_{t_\ell} \cdots X_{t_1}.$$

Successively averaging the independent variables $t_\ell, \ldots, t_1$ yields

$$\mathop{\mathbb{E}}_{T:|T|=\ell}[X_{t_\ell} \cdots X_{t_1}] = \left(\mathop{\mathbb{E}}_t X_t\right)^\ell = \mathsf{J}_k^\ell.$$

Averaging this conditional identity over the geometric length law $\Pr[L = \ell] = (1-\lambda_\gamma)\lambda_\gamma^\ell$ yields $F_{\gamma,k}$, as desired. □

Now, for fixed (k, m, T), we can apply Lemma 4.2 to deduce that the unnormalized leftover state on BC conditioned on both Bob and Charlie's outputs being m is

$$|\eta_{k,m,T}\rangle = \left(\widehat{V}^B_k(m) W^B_{k,T} \otimes \widehat{V}^C_k(m) \right) |\psi_{k,m}\rangle.$$

Thus, the correct-output probability is $\|\eta_{k,m,T}\|^2$.

Now, as informally outlined at the end of Section 5, a key step before applying Cauchy-Schwarz is to leverage a unitary degree of freedom on Charlie's side. We let

$$|\eta'_{k,m,T}\rangle = \left(\widehat{V}_k^B(m) W_{k,T}^B \otimes W_{k,T}^C \widehat{V}_k^C(m)\right) |\psi_{k,m}\rangle\,.$$

Then, we have

$$\begin{aligned} \|\eta_{k,m,T}\|^2 &= \|\eta'_{k,m,T}\|^2 \\ &\geq |\langle \psi_{k,m} \mid \eta'_{k,m,T}\rangle|^2. \end{aligned} \tag{6.3}$$

Lemma 6.2. *Let p_{Search} denote the success probability of the strategy obtained from the reduction of Reduction 2. Then,*

$$p_{\text{Search}} \geq \left| \mathop{\mathbb{E}}_{(k,m)\leftarrow\pi} \langle \psi_{k,m}| (\widehat{V}_k^B(m) \otimes I_C) F_{\gamma,k} (I_B \otimes \widehat{V}_k^C(m)) |\psi_{k,m}\rangle \right|^2,$$

Proof. Averaging Equation (6.3) and applying Jensen's inequality gives

$$\begin{aligned} p_{\text{Search}} &\geq \mathop{\mathbb{E}}_{k,m,T} |\langle \psi_{k,m} \mid \eta'_{k,m,T}\rangle|^2 \\ &\geq \left| \mathop{\mathbb{E}}_{k,m,T} \langle \psi_{k,m} \mid \eta'_{k,m,T}\rangle \right|^2. \end{aligned} \tag{6.4}$$

By Lemma 6.1, the average inside the absolute value is

$$\begin{aligned} \mathop{\mathbb{E}}_{k,m,T} \langle \psi_{k,m} \mid \eta'_{k,m,T}\rangle &= \mathop{\mathbb{E}}_{k,m,T} \langle \psi_{k,m}| \left(\widehat{V}_k^B(m) W_{k,T}^B \otimes W_{k,T}^C \widehat{V}_k^C(m)\right) |\psi_{k,m}\rangle \\ &= \mathop{\mathbb{E}}_{k,m} \langle \psi_{k,m}| (\widehat{V}_k^B(m) \otimes I_C) F_{\gamma,k} (I_B \otimes \widehat{V}_k^C(m)) |\psi_{k,m}\rangle\,, \end{aligned}$$

as desired. □

It will be convenient for us to rewrite the lower bound in Equation (6.4) more compactly as a filtered "trace overlap" (squared) in the form $|\mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})|^2$, for an appropriately chosen mixed state ρ and operators $\mathsf{B}, F_\gamma, \mathsf{C}$. Define the following block-diagonal Hermitian operators:

$$\begin{aligned} \varrho &:= \bigoplus_{k,m} \pi(k,m) \cdot |\psi_{k,m}\rangle\langle\psi_{k,m}|^{BC}, \\ \mathsf{B} &:= \bigoplus_{k,m} \left(\widehat{V}_k^B(m) \otimes I_C\right), \\ \mathsf{C} &:= \bigoplus_{k,m} \left(I_B \otimes \widehat{V}_k^C(m)\right), \\ \mathsf{J} &:= \bigoplus_{k,m} \mathsf{J}_k\,, \end{aligned} \tag{6.5}$$

where J_k as defined in Equation (6.1). Note that all of these operators have norm at most 1. Note also that the (k,m) block of J depends on k but not on m. Finally, define also

$$F_\gamma = (1-\lambda_\gamma) \sum_{\ell=0}^{\infty} \lambda_\gamma^\ell \mathsf{J}^\ell = \bigoplus_{k,m} F_{\gamma,k}\,. \tag{6.6}$$

Then, we have the following.

Corollary 6.3 (Lower bounding search success by a "filtered" overlap)**.** *Let p_{Search} denote the success probability of the strategy obtained from the reduction of Reduction 2. Then,*

$$p_{\text{Search}} \geq |\text{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})|^2 \ .$$

Proof. This follows immediately from Lemma 6.2, by rewriting the inner product as a trace and a direct computation. □

7 Analysis of the reduction, part 2: lower bounding the "filtered overlap" by the original winning probability

Having lower-bounded the winning probability p_{Search} of the search strategy by (the square of) the filtered overlap $\text{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})$ in Corollary 6.3, we are left with lower-bounding $\text{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})$ itself. We will show the following.

Lemma 7.1. *Let $\rho, \mathsf{B}, \mathsf{C}, F_\gamma$ be as defined in* (6.5) *and* (6.6)*. Let p_{Pred} be the winning probability of the strategy for $\mathsf{GL}(G)$. Let $\Delta = 2p_{\text{Pred}} - 1$, and assume $\Delta > 0$. Let $0 < \gamma \leq \Delta$. Then,*

$$|\text{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})| \geq \gamma \Delta \, .$$

Together, Lemma 7.1 and Corollary 6.3 prove Theorem 5.1 (which completes the proof of Theorem 3.4). So, the rest of this section is devoted to the proof of Lemma 7.1.

7.1 Some setup and useful facts

Recall the definition of $\rho, \mathsf{B}, \mathsf{C}, \mathsf{J}, F_\gamma$ from Equations (6.5) and (6.6). For $u, v \in \{+1, -1\}$, define

$$\Gamma_{u,v} := \frac{1}{4} \left(I + u\mathsf{B} + v\mathsf{C} + uv\mathsf{J} \right).$$

Every $\Gamma_{u,v}$ is positive semidefinite. Indeed, its (k, m) block is

$$\mathop{\mathbb{E}}_t \left[\frac{I + u \cdot \chi_m(t) \cdot V^B_{k,t}}{2} \otimes \frac{I + v \cdot \chi_m(t) \cdot V^C_{k,t}}{2} \right],$$

and each tensor factor is a projection since the $V^B_{k,t}$ and $V^C_{k,t}$ are reflections. The signs u and v record whether Bob and Charlie are correct or incorrect. In particular, note that the winning probability of the strategy in $\mathsf{GL}(G)$ is exactly

$$p_{\text{Pred}} = \text{Tr}(\varrho \Gamma_{+,+}).$$

Set

$$S := \frac{1}{2}(\mathsf{B} + \mathsf{C}), \qquad D := \frac{1}{2}(I - \mathsf{J}), \qquad R := \frac{1}{2}(\mathsf{B} - \mathsf{C}).$$

Letting $\Delta = 2p_{\text{Pred}} - 1$, and using the expression for p_{Pred} above, and $\text{Tr}(\varrho) = 1$, we can write the advantage (over $\frac{1}{2}$) directly in terms of S and D:

$$\begin{aligned} \Delta &= 2p_{\text{Pred}} - 1 \\ &= \frac{1}{2} \text{Tr}[\varrho(I + \mathsf{B} + \mathsf{C} + \mathsf{J})] - 1 \\ &= \frac{1}{2} \text{Tr}[\varrho(\mathsf{B} + \mathsf{C} + \mathsf{J} - I)] \end{aligned}$$

$$= \mathrm{Tr}(\varrho(S - D)) \,. \tag{7.1}$$

This identity nicely separates the sum S of the two local averages B and C from the "disagreement" operator D.

Since $\mathsf{J} \preceq I$, we have $D \succeq 0$. Importantly, we also have

$$D + R = 2\Gamma_{+,-} \succeq 0, \qquad D - R = 2\Gamma_{-,+} \succeq 0. \tag{7.2}$$

7.2 A key idea: the geometric filter as an inverse

Recall that

$$F_\gamma = (1 - \lambda_\gamma) \sum_{\ell=0}^{\infty} \lambda_\gamma^\ell \mathsf{J}^\ell, \qquad \lambda_\gamma = \frac{1}{1 + 2\gamma}, \qquad D = \frac{1}{2}(I - \mathsf{J}).$$

Since $\|\mathsf{J}\| \leq 1$ and $0 < \lambda_\gamma < 1$, we have $\|\lambda_\gamma \mathsf{J}\| < 1$, which implies that the geometric series converges in operator norm. In particular, the classic geometric sum identity, extended to sums of commuting operators, gives

$$\sum_{\ell=0}^{\infty} \lambda_\gamma^\ell \mathsf{J}^\ell = (I - \lambda_\gamma \mathsf{J})^{-1}.$$

Next, substituting $\lambda_\gamma = (1 + 2\gamma)^{-1}$ gives

$$1 - \lambda_\gamma = \frac{2\gamma}{1 + 2\gamma}, \qquad I - \lambda_\gamma \mathsf{J} = \frac{1}{1 + 2\gamma}\big((1 + 2\gamma)I - \mathsf{J}\big).$$

Combining these identities and then using $(1 + 2\gamma)I - \mathsf{J} = (I - \mathsf{J}) + 2\gamma I = D + 2\gamma I$, we obtain

$$\begin{aligned} F_\gamma &= (1 - \lambda_\gamma)(I - \lambda_\gamma \mathsf{J})^{-1} \\ &= 2\gamma\big((1 + 2\gamma)I - \mathsf{J}\big)^{-1} \\ &= \gamma(D + \gamma I)^{-1}. \end{aligned} \tag{7.3}$$

Because $D \succeq 0$, we have $D + \gamma I \succ 0$. Moreover, $D + \gamma I \succeq \gamma I$, so the last expression also shows that $0 \prec F_\gamma \preceq I$. We are now ready to tackle the proof that $|\mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})| \geq \gamma\Delta$. In fact, we will show the stronger statement that $\mathrm{Re}\,\mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C})$ already satisfies the bound.

Lemma 7.2. $\mathrm{Re}\,\mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C}) \geq \gamma\Delta$.

Proof. First, since $\mathsf{B} = S + R$ and $\mathsf{C} = S - R$, we have

$$\begin{aligned} \mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C}) = \; & \mathrm{Tr}(\varrho S F_\gamma S) - \mathrm{Tr}(\varrho S F_\gamma R) \\ & + \mathrm{Tr}(\varrho R F_\gamma S) - \mathrm{Tr}(\varrho R F_\gamma R). \end{aligned}$$

Since S, R, and F_γ are Hermitian, the first and last terms are real, while

$$\mathrm{Tr}(\varrho R F_\gamma S) = \overline{\mathrm{Tr}(\varrho S F_\gamma R)}.$$

Thus the mixed terms are purely imaginary, and taking real parts gives

$$\mathrm{Re}\,\mathrm{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C}) = \mathrm{Tr}[\varrho S F_\gamma S] - \mathrm{Tr}[\varrho R F_\gamma R] \,.$$

We will spend the rest of the proof bounding $\mathrm{Tr}[\varrho S F_\gamma S]$ and $\mathrm{Tr}[\varrho R F_\gamma R]$ separately.

We first lower-bound $\mathrm{Tr}[\varrho S F_\gamma S]$. Since $F_\gamma \succ 0$ (by (7.3)), Cauchy–Schwarz for the Hilbert-Schmidt inner product and the cyclic property of the trace give

$$\begin{aligned}\left|\mathrm{Tr}(\varrho S)\right|^2 &= \left|\mathrm{Tr}\left((F_\gamma^{-1/2}\rho^{1/2})\cdot(F_\gamma^{1/2}\,S\,\rho^{1/2})^\dagger\right)\right|^2 \\ &= \left|\left\langle F_\gamma^{-1/2}\varrho^{1/2}, F_\gamma^{1/2} S \varrho^{1/2}\right\rangle_{\mathrm{HS}}\right|^2 \\ &\le \mathrm{Tr}(\varrho F_\gamma^{-1})\,\mathrm{Tr}(\varrho S F_\gamma S).\end{aligned}$$

Rearranging gives

$$\mathrm{Tr}(\varrho S F_\gamma S) \ge \frac{\left|\mathrm{Tr}(\varrho S)\right|^2}{\mathrm{Tr}(\varrho F_\gamma^{-1})} \tag{7.4}$$

We will simplify both numerator and denominator on the RHS. Let $p_{\neq}$ denote the probability that Bob and Charlie output different bits. Let X and Y denote the ±1-valued random variables for Bob and Charlie respectively guessing their own inner product correctly (where $+1$ denotes a correct guess). So, for fixed (k,m,t), we have

$$\begin{aligned}\mathbb{E}[XY \mid k,m,t] &= \langle\psi_{k,m}|\left(\chi_m(t)V_{k,t}^B \otimes \chi_m(t)V_{k,t}^C\right)|\psi_{k,m}\rangle \\ &= \langle\psi_{k,m}|\left(V_{k,t}^B \otimes V_{k,t}^C\right)|\psi_{k,m}\rangle,\end{aligned}$$

Averaging over (k,m,t) and using the definitions of ϱ and J (which can be thought of as an "agreement" operator), gives

$$\mathbb{E}[XY] = \mathrm{Tr}(\varrho\mathsf{J}).$$

Moreover, $XY = +1$ with probability $1-p_{\neq}$ and $XY=-1$ with probability $p_{\neq}$, so $\mathbb{E}[XY] = 1-2p_{\neq}$. Recalling that $D = \frac{1}{2}(I-\mathsf{J})$, we therefore obtain

$$\mathrm{Tr}(\varrho D) = \frac{1}{2}(1-\mathrm{Tr}(\varrho\mathsf{J})) = \frac{1}{2}(1-\mathbb{E}[XY]) = p_{\neq}. \tag{7.5}$$

Combining Equation (7.1) with the above gives

$$\mathrm{Tr}(\varrho S) = \Delta + p_{\neq}\,. \tag{7.6}$$

By Equation (7.3) and the invertibility of $D+\gamma I$,

$$F_\gamma^{-1} = \frac{1}{\gamma}(D+\gamma I) = I + \frac{D}{\gamma}.$$

Hence, since ϱ is a state and $\mathrm{Tr}(\varrho D) = p_{\neq}$,

$$\mathrm{Tr}(\varrho F_\gamma^{-1}) = 1 + \frac{p_{\neq}}{\gamma}. \tag{7.7}$$

Substituting (7.6) and (7.7) back into (7.4) yields

$$\mathrm{Tr}(\varrho S F_\gamma S) \ge \frac{4\gamma(\Delta+p_{\neq})^2}{\gamma+p_{\neq}}. \tag{7.8}$$

We now turn to upper bounding $\mathrm{Tr}[\varrho R F_\gamma R]$. First, note that

$$\begin{pmatrix} D & R \\ R & D\end{pmatrix} \succeq 0\,,$$

since the operator is unitarily equivalent to $(D+R)\oplus(D-R)$, and both $D+R$ and $D-R$ are positive semidefinite by (7.2). Since $\begin{pmatrix} 0 & 0 \\ 0 & \gamma I \end{pmatrix} \succ 0$, we also have

$$\begin{pmatrix} D & R \\ R & D+\gamma I \end{pmatrix} \succeq 0.$$

Because $D \succeq 0$ and $\gamma > 0$, the lower-right block $D+\gamma I$ is strictly positive (and hence invertible). Thus, the Schur-complement criterion applied to

$$\begin{pmatrix} D & R \\ R & D+\gamma I \end{pmatrix} \succeq 0$$

gives

$$D - R(D+\gamma I)^{-1}R \succeq 0,$$

or equivalently $R(D+\gamma I)^{-1}R \preceq D$. Multiplying by γ and using $F_\gamma = \gamma(D+\gamma I)^{-1}$ yields

$$RF_\gamma R \preceq \gamma D\,.$$

Since $\varrho \succeq 0$, the latter implies

$$\operatorname{Tr}(\varrho R F_\gamma R) \leq \gamma \operatorname{Tr}(\varrho D) = \gamma p_{\neq}\,, \tag{7.9}$$

where the last equality is by Equation (7.5).

Finally, combining the S-term lower bound of (7.8) with the R-term upper bound of (7.9) yields

$$\begin{aligned} \operatorname{Re}\operatorname{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C}) &= \operatorname{Tr}[\varrho S F_\gamma S] - \operatorname{Tr}[\varrho R F_\gamma R] \\ &\geq \frac{\gamma(\Delta + p_{\neq})^2}{\gamma + p_{\neq}} - \gamma p_{\neq}\,. \end{aligned}$$

Finally, since $0 < \gamma \leq \Delta$ and $p_{\neq} \geq 0$, we have

$$\frac{\gamma(\Delta + p_{\neq})^2}{\gamma + p_{\neq}} - \gamma p_{\neq} \geq \Delta(\gamma + p_{\neq}) - \gamma p_{\neq} \geq \gamma\Delta,$$

which yields

$$\operatorname{Re}\operatorname{Tr}(\varrho \mathsf{B} F_\gamma \mathsf{C}) \geq \gamma\Delta.$$

□

Lemma 7.1 implies Lemma 7.2, which, together with Corollary 6.3, completes the proof of Theorem 5.1 (and Theorem 3.4).

8 Uniformity and efficiency analysis

Theorem 5.1 establishes the success guarantee for Reduction 2. In this section, we study its runtime. Theorem 8.1 below says that, given a lower bound on $0 < \gamma \leq \Delta$, this reduction is black-box, uniform, and has an *expected* runtime polynomial in $\frac{1}{\gamma}$ and in the runtime of the original strategy. Here, "black-box" means that, apart from the numerical parameter γ, the reduction uses the strategy for $\mathsf{GL}(G)$ only via an application of the splitting channel $\mathcal{A}_n$, and coherent controlled applications

of the eflections $V^B_{k,r}$ and $V^C_{k,r}$; it does not use any additional structure of the strategy. Theorem 8.2 says that an appropriate truncation of Reduction 2 has the desired *worst-case* runtime.

Finally, we describe a uniform reduction that does *not* require knowing a lower bound on Δ. Theorem 8.3 says that this reduction can be made to have expected runtime runtime *linear* in that of the original strategy, at the cost of a polynomial loss in the resulting search success probability.

Theorem 8.1. *Suppose the strategy for* $\mathsf{GL}(G)$ *has a uniform circuit implementation, and that an efficiently computable parameter* $0 < \gamma \leq \Delta$ *is given. Then Reduction 2 is uniform and has expected running time polynomial in* $1/\gamma$ *and in the running time of the strategy for* $\mathsf{GL}(G)$*.*

Proof. We examine Reduction 2 step by step.

1. The splitting procedure is exactly the same as in the common-mask GL unclonable-encryption game. Thus, the reduction uses the original uniform circuit for $\mathcal{A}_n$ without modification.

2. Bob samples L and T. Conditioned on $L = \ell$, he samples ℓ independent uniform n-bit masks and applies
$$W^B_{k,T} = V^B_{k,t_\ell} \cdots V^B_{k,t_1}$$
to register B. He then initializes a register R to $|0^n\rangle$, applies the Goldreich–Levin unitary GL^B_k to registers RB, measures R in the computational basis, and returns the outcome. Conditioned on $L = \ell$, this requires ℓ applications of Bob's eflections for $W^B_{k,T}$, one further coherent controlled application inside GL^B_k, and time $O(n\ell)$ to sample the masks. Hence, its running time is polynomial in n, ℓ, and the running time of Bob's decoder.

3. Charlie applies GL^C_k to his registers and measures the output register in the computational basis. This has a uniform circuit implementation with running time polynomial in n and the running time of Charlie's decoder.

The geometric random variable L can be sampled by repeated independent Bernoulli trials with success probability $1 - \lambda_\gamma$. Since L has the distribution in Equation (5.1) and $\lambda_\gamma = (1 + 2\gamma)^{-1}$,
$$\mathbb{E}[L] = \frac{\lambda_\gamma}{1 - \lambda_\gamma} = \frac{1}{2\gamma}.$$
Thus, the expected number of sampled masks and decoder-reflection applications is $O(1/\gamma)$. The claimed uniform expected-running-time bound follows. In particular, if γ is inverse polynomial in n and the original strategy is polynomial time, then the resulting strategy is expected polynomial time in n. □

We can obtain a worst-case polynomial-time variant by just modifying Bob's sampling of L (by setting a maximum value). This will achieve the desired worst-case runtime at the cost of a small loss in the search success probability. Define
$$L_{\max} := \left\lceil \frac{\log(2/\gamma^4)}{\log(1 + 2\gamma)} \right\rceil.$$
Bob generates at most $L_{\max}$ independent Bernoulli bits, each equal to 0 with probability λ_γ and to 1 with probability $1 - \lambda_\gamma$. If the first 1 occurs after exactly $\ell < L_{\max}$ preceding zeros, he sets $L = \ell$ and continues; if all $L_{\max}$ bits are zero, he aborts. Let $p^{(<L_{\max})}_{\text{Search}}$ denote the joint winning probability of this truncated procedure.

Theorem 8.2 (Truncated variant)**.** *Suppose the strategy for* $\mathsf{GL}(G)$ *has a uniform circuit implementation, and that an efficiently computable parameter* $0 < \gamma \le \Delta$ *is given. Then, the "truncated" procedure described above is uniform, has worst-case running time polynomial in* $1/\gamma$ *and in the running time of the strategy for* $\mathsf{GL}(G)$*, and satisfies*

$$p_{\text{Search}}^{(<L_{\max})} \ge \frac{1}{2}\gamma^2\Delta^2 \ge \frac{1}{2}\gamma^4.$$

Proof. The probability that the original geometric sampler produces $L \ge L_{\max}$ is

$$(1-\lambda_\gamma)\sum_{\ell=L_{\max}}^{\infty}\lambda_\gamma^\ell = \lambda_\gamma^{L_{\max}}.$$

Aborting on this event can decrease the winning probability by at most its probability. Therefore, Theorem 5.1 gives

$$\begin{aligned} p_{\text{Search}}^{(<L_{\max})} &\ge p_{\text{Search}} - \lambda_\gamma^{L_{\max}} \\ &\ge \gamma^2\Delta^2 - \lambda_\gamma^{L_{\max}}. \end{aligned}$$

By the definition of $L_{\max}$,

$$\lambda_\gamma^{L_{\max}} = (1+2\gamma)^{-L_{\max}} \le \frac{\gamma^4}{2} \le \frac{\gamma^2\Delta^2}{2},$$

where the last inequality uses $\gamma \le \Delta$. Consequently,

$$p_{\text{Search}}^{(<L_{\max})} \ge \frac{1}{2}\gamma^2\Delta^2 \ge \frac{1}{2}\gamma^4.$$

Finally, since $0 < \gamma \le \Delta \le 1$ and $\log(1+2\gamma) \ge 2\gamma/(1+2\gamma) \ge 2\gamma/3$,

$$L_{\max} = O\Big(\gamma^{-1}\log(2/\gamma)\Big) = O(\gamma^{-2}).$$

Thus, the truncated procedure is uniform, and uses at most polynomially many masks and decoder-reflection applications. In particular, it has uniform worst-case polynomial in n running time whenever $1/\gamma$ and the running time of the original strategy are polynomial in n. □

Finally, we show that it is not necessary to know any lower bound on Δ, thereby making the expected-time reduction completely uniform.

Theorem 8.3 (Variant when a lower bound on Δ is not known)**.** *Let* $\Delta > 0$*. There exists a uniform reduction that does not take* Δ*, or any lower bound on* Δ*, as input. Its expected running time is linear in the running time of the strategy for* $\mathsf{GL}(G)$ *(and independent of* Δ*), and the resulting strategy for* G *succeeds with probability at least*

$$\frac{3}{16}\Delta^6.$$

Proof. The reduction first samples an integer $I \ge 1$ according to

$$\Pr[I = i] = 3 \cdot 4^{-i}, \qquad i = 1, 2, \ldots, \tag{8.1}$$

sets $\gamma = 2^{-I}$, and then runs the procedure in Reduction 2 with this value of γ. The procedure is well defined for every $\gamma > 0$; the condition $\gamma \le \Delta$ is needed only for its success guarantee.

The distribution in (8.1) can be sampled uniformly as follows: repeatedly sample two bits, and stop at the first pair that is not 00. If this occurs on the ith trial, set $I = i$.

Let $j \geq 1$ be such that

$$\frac{\Delta}{2} \leq 2^{-j} \leq \Delta.$$

Such a j exists for every $0 < \Delta \leq 1$ and is used only in the analysis; the reduction does not need to know it. Conditioned on the event $I = j$, Theorem 5.1 applies and gives success probability at least $(2^{-j}\Delta)^2$. Ignoring all other branches, which contribute nonnegative success probability, we obtain

$$\begin{aligned} p_{\mathrm{Search}} &\geq \Pr[I = j]\,(2^{-j}\Delta)^2 \\ &= 3 \cdot 4^{-j}\,(2^{-j}\Delta)^2 \\ &= 3(2^{-j})^4\Delta^2 \\ &\geq 3\left(\frac{\Delta}{2}\right)^4 \Delta^2 = \frac{3}{16}\Delta^6. \end{aligned}$$

It remains to bound the expected running time. Conditioned on $I = i$, we have $\gamma = 2^{-i}$, and hence the expected sequence length in Reduction 2 is

$$\mathbb{E}[L \mid I = i] = \frac{1}{2\gamma} = 2^{i-1}.$$

Therefore,

$$\mathbb{E}[L] = \sum_{i=1}^{\infty} 3 \cdot 4^{-i}\, 2^{i-1} = \frac{3}{2}\sum_{i=1}^{\infty} 2^{-i} = \frac{3}{2}.$$

Thus, the expected number of sampled masks and decoder-reflection calls is constant, and the overall expected running time is linear in the running time of the original strategy for $\mathsf{GL}(G)$. □

References

[GL89] O. Goldreich and L. A. Levin. "A Hard-Core Predicate for All One-Way Functions". In: *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing (STOC '89)*. ACM Press, 1989, pp. 25–32. DOI: 10.1145/73007.73010.

[AC02] M. Adcock and R. Cleve. "A Quantum Goldreich–Levin Theorem with Cryptographic Applications". In: *STACS 2002*. Ed. by H. Alt and A. Ferreira. Vol. 2285. Lecture Notes in Computer Science. Springer, 2002, pp. 323–334. DOI: 10.1007/3-540-45841-7_26.

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

[CLLZ21] A. Coladangelo, J. Liu, Q. Liu, and M. Zhandry. "Hidden Cosets and Applications to Unclonable Cryptography". In: *Advances in Cryptology—CRYPTO 2021, Part I*. Ed. by T. Malkin and C. Peikert. Vol. 12825. Lecture Notes in Computer Science. Springer, 2021, pp. 556–584. DOI: 10.1007/978-3-030-84242-0_20.

[CV22] Eric Culf and Thomas Vidick. "A Monogamy-of-Entanglement Game for Subspace Coset States". In: *Quantum* 6 (2022), p. 791. DOI: 10.22331/q-2022-09-01-791.

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

[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)*. Ed. by Steven T. Flammia. 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.

[TFKW13] 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. DOI: 10.1088/1367-2630/15/10/103002.

[AKLL+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*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 212–241. DOI: 10.1007/978-3-031-15979-4_8.

[KT25] Srijita Kundu and Ernest Y.-Z. Tan. "Device-Independent Uncloneable Encryption". In: *Quantum* 9 (2025), p. 1582. DOI: 10.22331/q-2025-01-08-1582.

[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)*. Ed. by Raghu Meka. 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.

[AB24] Prabhanjan Ananth and Amit Behera. "A Modular Approach to Unclonable Cryptography". In: *Advances in Cryptology—CRYPTO 2024, Part VII*. Ed. by Leonid Reyzin and Douglas Stebila. Vol. 14926. Lecture Notes in Computer Science. Springer, 2024, pp. 3–37. DOI: 10.1007/978-3-031-68394-7_1.

[CHV23] Céline Chevalier, Paul Hermouet, and Quoc-Huy Vu. *Towards Unclonable Cryptography in the Plain Model*. Cryptology ePrint Archive, Report 2023/1825. 2023. arXiv: 2311.16663. URL: https://eprint.iacr.org/2023/1825.

[CLX26] 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)*. Ed. by Shubhangi Saraf. Vol. 362. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026, 41:1–41:20. DOI: 10.4230/LIPIcs.ITCS.2026.41.

[CKNY25] Jeffrey Champion, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. "Untelegraphable Encryption and Its Applications". In: *Theory of Cryptography—TCC 2025, Part III*. Vol. 16270. Lecture Notes in Computer Science. Springer, 2025, pp. 3–35. DOI: 10.1007/978-3-032-12296-4_1.

[ÇG24] Alper Çakan and Vipul Goyal. *Unbounded Leakage-Resilient Encryption and Signatures*. Cryptology ePrint Archive, Report 2024/1876. 2024. URL: https://eprint.iacr.org/2024/1876.

[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.

[BBC26a] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. *The Uncloneable Bit Exists*. 2026. arXiv: 2603.08916. URL: https://arxiv.org/abs/2603.08916.

[AS26] Prabhanjan Ananth and Amit Sahai. *Unconditional Unclonable Encryption*. Cryptology ePrint Archive, Report 2026/1511. 2026. arXiv: 2607.21551. URL: https://eprint.iacr.org/2026/1511.

[Rag26] Seyoon Ragavan. *Efficient Unclonable Encryption from Pauli Eigenstates*. Cryptology ePrint Archive, Report 2026/1509. 2026. arXiv: 2607.21811. URL: https://eprint.iacr.org/2026/1509.

[AAG22] Anurag Anshu, Itai Arad, and David Gosset. "An area law for 2d frustration-free spin systems". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 12–18.

[ALV12] Itai Arad, Zeph Landau, and Umesh Vazirani. "Improved one-dimensional area law for frustration-free systems". In: *Physical Review B—Condensed Matter and Materials Physics* 85.19 (2012), p. 195145.

[BBC26b] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. "Statistically secure uncloneable encryption of arbitrary messages". In: *arXiv preprint arXiv:2607.28561* (2026).

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. "Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics". In: *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*. 2019, pp. 193–204.

A A common-mask strategy on which plain GL fails completely

We describe a simple two-bit cloning game and a strategy for its common-GL mask game whose prediction probability is strictly larger than 1/2, while the corresponding plain Goldreich–Levin simultaneous extraction probability is zero (and hence the overlap between vectors $|b_{k,m}\rangle, |c_{k,m}\rangle$ is also zero).

We consider the following distribution over (k, m): uniform over all pairs satisfying $k = m$. Although k completely reveals the message m in this setting, it is still interesting to observe that plain GL fails.

The challenger prepares the following fixed quantum state (regardless of k, m) and Alice does nothing:

$$|\psi\rangle^{A_1A_2} := \frac{1}{\sqrt{6}}\left(|00\rangle + |01\rangle + 2|11\rangle\right).$$

Consider a splitting procedure that sends A_1 to Bob and A_2 to Charlie. Thus, the post-splitting state is exactly $|\psi\rangle^{BC}$. For masks $r \in \{00, 01, 10, 11\}$, define the local reflections

r	00	01	10	11
B_r	Z	X	I	$-X$
C_r	Z	$-Z$	I	$-X$

and let $\chi_x(r) := (-1)^{\langle r,x\rangle}$. Then, define Bob's and Charlie's decoder reflections to be

$$V^B_{k,r} := \chi_k(r)B_r, \qquad V^C_{k,r} := \chi_k(r)C_r.$$

Since $k = m$, note that the reflections about the respective correct-output subspaces are exactly

$$\chi_m(r)V^B_{k,r} = B_r, \qquad \chi_m(r)V^C_{k,r} = C_r,$$

which act as $+1$ on the correct-output subspace and -1 on incorrect. Then, the four projectors corresponding to simultaneous correct guesses, for $r = 00, 01, 10, 11$ respectively, are

$$|0\rangle\langle 0| \otimes |0\rangle\langle 0|, \qquad |+\rangle\langle +| \otimes |1\rangle\langle 1|, \qquad I, \qquad |-\rangle\langle -| \otimes |-\rangle\langle -|.$$

Their expectations with respect to $|\psi\rangle$ can be computed to be

$$\frac{1}{6}, \qquad \frac{3}{4}, \qquad 1, \qquad \frac{1}{6}.$$

So, the common-mask prediction probability is

$$p_{\text{Pred}} = \frac{1}{4}\left(\frac{1}{6} + \frac{3}{4} + 1 + \frac{1}{6}\right) = \frac{25}{48} > \frac{1}{2}.$$

Now, consider what happens when Bob and Charlie each apply their plain GL extractors associated with the above strategy. Since $k = m$, the relevant Fourier coefficients are

$$\widehat{V}^B_k(m) = \frac{1}{4}\sum_r \chi_m(r)V^B_{k,r} = \frac{1}{4}\sum_r B_r = \frac{I+Z}{4} = \frac{1}{2}|0\rangle\langle 0|,$$
$$\widehat{V}^C_k(m) = \frac{1}{4}\sum_r \chi_m(r)V^C_{k,r} = \frac{1}{4}\sum_r C_r = \frac{I-X}{4} = \frac{1}{2}|-\rangle\langle -|.$$

However, $(\langle 0|\langle -|)|\psi\rangle = \frac{1}{\sqrt{2}}\left(\langle 00| - \langle 01|\right)|\psi\rangle = 0$, which implies that the entire simultaneous correct-output branch vanishes:

$$\left(\widehat{V}^B_k(m) \otimes \widehat{V}^C_k(m)\right)|\psi\rangle = 0.$$

Therefore, the plain GL extraction probability is exactly

$$p_{\text{plainGL}} = \left\|\left(\widehat{V}^B_k(m) \otimes \widehat{V}^C_k(m)\right)|\psi\rangle\right\|^2 = 0.$$

One can also check that the individual correct-output vectors $|b_{k,m}\rangle = \left(\widehat{V}^B_k(m) \otimes I\right)|\psi\rangle$ and $|c_{k,m}\rangle = \left(I \otimes \widehat{V}^C_k(m)\right)|\psi\rangle$, satisfy

$$\langle b_{k,m} \mid c_{k,m}\rangle = \langle\psi|\left(\widehat{V}^B_k(m) \otimes \widehat{V}^C_k(m)\right)|\psi\rangle = 0.$$

Thus, $p_{\text{Pred}} = \frac{25}{48} > \frac{1}{2}$, but $p_{\text{plainGL}} = \langle b_{k,m} \mid c_{k,m}\rangle = 0$.