Non-Local Search-to-Decision over $\mathbb{F}_2$, and More*

Prabhanjan Ananth†
UCSB

Abstract

Non-local search-to-decision asks whether the difficulty of two non-communicating (non-local) parties both predicting x given a bipartite state (that possibly depends on x) implies the difficulty of non-local parties both predicting the inner product $\langle r, x \rangle$, for a uniformly random r. This problem and its variants have been extensively studied and are motivated by applications to unclonable cryptographic primitives such as unclonable encryption and copy-protection. We study the identical-challenge setting, in which both parties receive the same uniformly random vector r, in contrast to the independently sampled challenges considered in prior works. We demonstrate positive results for the cases when x and r are vectors over $\mathbb{F}_2$ and over large finite fields. As a consequence, we obtain a conceptually different proof of indistinguishability-secure unclonable encryption.

*AI disclosure statement can be found on the last page of the paper.

†prabhanjan@cs.ucsb.edu

Contents

1 Introduction **3**
1.1 Technical Overview 6
1.1.1 Search-to-Decision over $\mathbb{F}_2$ 6
1.1.2 Search-to-Decision over Large Finite Fields 13
1.2 Related Work 16

2 Preliminaries **18**
2.1 Basic notation 18
2.2 Quantum Information 18
2.3 Characters 19

3 Non-Local Search-to-Decision over $\mathbb{F}_2$ **19**
3.1 Start from the desired expectation and decompose the state 20
3.2 The high–high component and all terms involving omitted vectors 22
3.3 Diagonal and cross terms of the retained vector 24
3.4 Expand the effective cross scalar into common-challenge contributions 31
3.5 Each common-challenge contribution defines an extractor 36
3.6 Complete the top-down estimate 40

4 Extensions over arbitrary finite fields **41**
4.1 Model and normalization 41
4.2 Independent challenges over arbitrary finite fields 42
4.3 Identical Challenge Setting over Arbitrary Finite Fields 49

1 Introduction

Search-to-Decision Reductions in TCS. Search-to-decision reductions use a procedure for a decision problem as an oracle to solve its associated search problem, and they have been studied in many areas of theoretical computer science. A classic example is the self-reduction for satisfiability: a satisfying assignment can be recovered using polynomially many queries to an oracle that decides whether a formula is satisfiable. In cryptography, the seminal Goldreich–Levin theorem [GL89] gives, in this broad sense, a prediction-to-reconstruction reduction and shows that every one-way function admits a hard-core predicate. Another example is Regev's reduction, which turns an efficient algorithm for the decisional learning with errors problem, in some parameter regimes, into an efficient algorithm for the search learning with errors problem [Reg05].

Search-to-Decision Reductions in Unclonable Cryptography. Unclonable cryptography is an area of quantum cryptography that leverages the no-cloning principle of quantum mechanics to demonstrate the existence of cryptographic tasks that are otherwise impossible to achieve using purely classical resources. Often in unclonable cryptography, one has to deal with (non-local) adversarial parties which cannot communicate but may share entanglement. Just as search-to-decision reductions have been immensely useful for classical cryptography, search-to-decision reductions have also proven to be essential for unclonable cryptography [CLL+21, AKL+22, AKL23, AKY25]. However, entanglement presents a significant obstacle to search-to-decision reductions for non-local parties; in particular, entanglement presents technical challenges when applying a one-party reduction separately to both non-local parties. Before we move further, we first state the non-local search-to-decision problem considered in this work.

Non-Local Search-to-Decision. Consider the following problem: a dealer encodes a uniformly random string $x \in \mathbb{F}^n$, for a field $\mathbb{F}$, into a bipartite state and gives one part to Bob and the other to Charlie. The two non-communicating parties then receive the same uniformly random vector $r \in \mathbb{F}^n$ and they must *simultaneously* output the inner product $\langle r, x\rangle$. The trivial strategy that Bob and Charlie can employ is to output a uniformly random element in $\mathbb{F}$. The non-local search-to-decision question asks whether a noticeable improvement over this trivial strategy can occur unless there also exist local measurements with which Bob and Charlie both *simultaneously* recover the entire hidden string x with noticeable probability.

Formally, let us first consider the case when $\mathbb{F} = \mathbb{F}_2$. Let $x \xleftarrow{\$} \mathbb{F}_2^n$ and let a quantum channel Φ encode the computational basis state $|x\rangle\langle x|$ as the following bipartite state:

$$\rho_x := \Phi(|x\rangle\langle x|)$$

on registers $(\mathbf{B}, \mathbf{C})$ held by two non-communicating parties, Bob and Charlie.

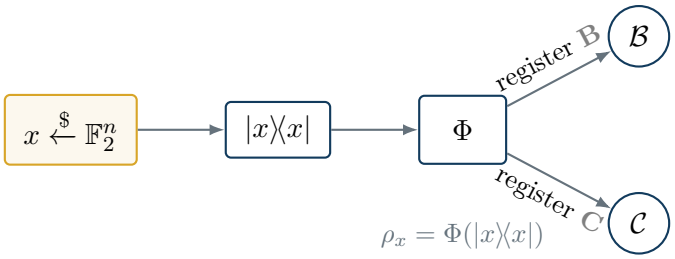


Figure 1: The splitting stage. A uniformly random string is encoded by an arbitrary channel into the bipartite state ρ_x. Bob and Charlie receive disjoint registers and cannot communicate.

We consider two experiments:

- **Search experiment**: in this experiment, Bob and Charlie succeed if they simultaneously output x. Denote the optimal success probability in this experiment by p_{srch}.
- **Identical-challenge decision experiment**: in this experiment, both Bob and Charlie are additionally given $r \xleftarrow{\$} \mathbb{F}^n$. They win if they simultaneously predict $\langle r, x\rangle$. Denote the optimal success probability in this experiment by p_{pred}.

The non-local search-to-decision problem asks whether the following holds: if p_{srch} is negligible, is p_{pred} at most $\frac{1}{2} + \varepsilon$, where ε is negligible?

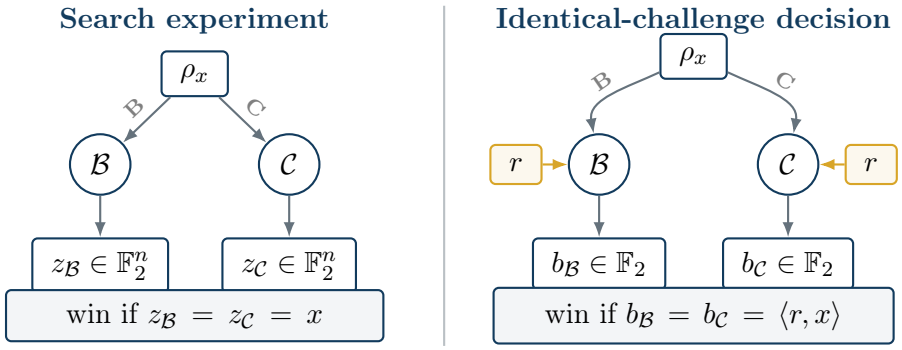


Figure 2: The search and the identical-challenge decision experiments.

Note that if, instead of two parties, there were exactly one party, then the celebrated Goldreich-Levin theorem [GL89] (and its quantum generalizations [AC02, CLL+21]) answers the above question in the affirmative. The non-triviality of the above question stems from the fact that there are two parties which are entangled with each other. If the two parties do not share any entanglement (and instead share correlated randomness) then the above question has already been answered in the affirmative [CG24b, CKN+25, CLX26a]. Indeed, the above question was left as an important open problem by many works [KT25, AB24, CHV23, AKY25, CLX26a].

Our Results. Our main theorem addresses the above non-local search-to-decision problem. We show the following:

Theorem 1.1 (Informal)**.** *If p_{srch} is the optimal success probability in the search experiment and p_{pred} is the optimal success probability in the identical-challenge decision experiment (as defined above), then:*

$$p_{\mathsf{pred}} \leq \min\left\{1, \frac{1}{2} + 5p_{\mathsf{srch}}^{1/22}\right\}$$

As a consequence, if p_{srch} is negligible (in n) then p_{pred} is at most $\frac{1}{2}+$ negligible (in n).

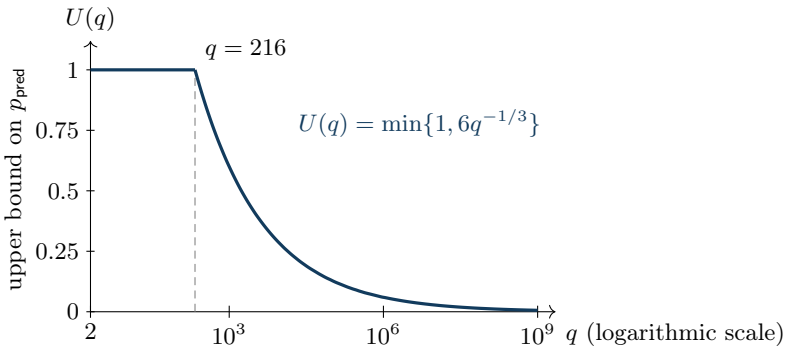
We note that the exponent $1/22$ is not optimal (as demonstrated by the concurrent work; more on this later). However, for applications, we typically consider the setting when p_{srch} is negligible and in this setting, the constant exponent does **not** significantly affect the qualitative result statement.

The above result is concerned with the non-local search-to-decision problem over $\mathbb{F}_2$. We also extend the above result to arbitrary finite fields. In this case, however, we obtain non-trivial (but sub-optimal) bounds for large-sized finite fields.

Theorem 1.2 (Informal)**.** *If p_{srch} and p_{pred} are optimal success probabilities of the extended (i.e., over a field $\mathbb{F}$) search and identical-challenge decision[1] experiments, then:*

$$p_{\mathsf{pred}} \leq \min\left\{1,\ 6\left(\frac{1}{q} + 2\left(\frac{q-1}{q}\right)^{2/3} p_{\mathsf{srch}}^{1/3}\right)^{1/3}\right\},$$

where $q = |\mathbb{F}|$. As a consequence, for $q > 216$, if p_{srch} is negligible (in n) then p_{pred} is at most $\frac{6}{q^{\frac{1}{3}}}+$ negligible (in n); see the graph below.



Here $U(q) := \min\{1, 6q^{-1/3}\}$ is the upper bound on p_{pred} after suppressing the negligible term; it equals 1 through $q = 216$ and then decreases as $q^{-1/3}$. Here, we do not treat q as the size of a finite field and in particular, allow it to take any positive value.

As a result of Theorem 1.1, we obtain the following corollary to unclonable encryption [BL20, AK21]. Informally speaking, unclonable encryption is a one-time encryption scheme in which the ciphertexts are unclonable quantum states. There are two security notions considered for unclonable encryption, namely, search security and indistinguishability security. In the search security, the non-local adversaries (after the splitting stage) are expected to predict the message encrypted in the

[1]Even though in the extended decision experiment, the outputs of $\mathcal{B}$ and $\mathcal{C}$ are no longer boolean, for consistency purposes, we abuse terminology and call this a decision experiment.

unclonable challenge ciphertext whereas in the indistinguishability security, the non-local adversaries are expected to distinguish whether the message encrypted in the challenge ciphertext is one of two adversarially chosen messages.

Corollary 1.3 (Application to unclonable encryption)**.** *Suppose a one-time secret-key scheme for unclonable encryption of uniform messages in* $\mathbb{F}_2^{n(\lambda)}$ *has information-theoretic search security. Then there is a one-time secret-key scheme for one-bit messages whose identical-challenge distinguishing success is at most* $1/2 + \mathrm{negl}(\lambda)$.

The above corollary follows in two steps:

- Ananth, Kaleoglu, and Liu [AKL23, Section 5.4.1, Corollaries 4–5] prove the above search-to-indistinguishability corollary but for a different primitive, namely one-time single-decryptor encryption secure with respect to identical challenges[2]. Roughly speaking, a single-decryptor encryption is an encryption scheme where the decryption keys are unclonable quantum states. Here, identical challenges refer to the fact that in the security experiment of single-decryptor encryption, the non-local parties receive the same challenge ciphertext.
- Moreover, Georgiou and Zhandry [GZ20, Theorem 2] show the equivalence between one-time single-decryptor encryption with respect to identical challenges and one-time unclonable encryption.

The resulting construction of indistinguishability-secure unclonable encryption retains the same ciphertext structure as the search-secure unclonable encryption. As a result, we obtain indistinguishability-secure unclonable encryption based on a family of unclonable states, including BB84 states and coset states. Prior works designed unclonable encryption based on Pauli-based states [AS26, Rag26] or based on Clifford-based states [BBC26].

1.1 Technical Overview

1.1.1 Search-to-Decision over $\mathbb{F}_2$

We follow the notation set up in Section 1. That is, we define ρ_x to be the bipartite state received by $\mathcal{B}$ and $\mathcal{C}$, p_{pred} to be the probability of simultaneously predicting the inner product and p_{srch} to be the optimal probability of simultaneously guessing the entire string.

Additionally, we define the following notation:

- Let $\{\mathcal{P}_{\mathcal{B}}^{r,0}, \mathcal{P}_{\mathcal{B}}^{r,1}\}$ be a family of POVM measurements, parameterized by $r \in \mathbb{F}_2^n$, performed by $\mathcal{B}$.
- Similarly, let $\{\mathcal{P}_{\mathcal{C}}^{r,0}, \mathcal{P}_{\mathcal{C}}^{r,1}\}$ be a family of POVM measurements, parameterized by $r \in \mathbb{F}_2^n$, performed by $\mathcal{C}$.

Thus, we can express the probability of predicting the inner product by $\mathcal{B}$ and $\mathcal{C}$ as follows:

$$p_{\mathsf{pred}} := \mathbb{E}_{x,r}\left[\mathrm{Tr}\left(\left(\mathcal{P}_{\mathcal{B}}^{r,\langle r,x\rangle} \otimes \mathcal{P}_{\mathcal{C}}^{r,\langle r,x\rangle}\right)\rho_x\right)\right], \tag{1}$$

[2]Technically, in [AKL23], the single-decryptor encryption is proven secure with respect to independent challenges. However, the same result can be readily extended to work with identical challenges as well.

We first purify ρ_x to obtain a state of the form $|\psi^x\rangle_{\mathbf{BCE}}$, where $\mathbf{B}$ is the register measured by $\mathcal{B}$, $\mathbf{C}$ is the register measured by $\mathcal{C}$ and $\mathbf{E}$ is the external reference register. Moreover, we define $\mathbf{G}_x$ to be $\mathbb{E}_r\left[\mathcal{P}_{\mathcal{B}}^{r,\langle r,x\rangle} \otimes \mathcal{P}_{\mathcal{C}}^{r,\langle r,x\rangle}\right] \otimes I_{\mathbf{E}}$. Thus, we can rewrite p_{pred} as:

$$p_{\mathsf{pred}} = \mathbb{E}_x\left[\langle\psi^x|\mathbf{G}_x|\psi^x\rangle\right]$$

Approach in a Nutshell: In the rest of the proof, we decompose $|\psi^x\rangle$ into many parts and then we analyze the action of $\mathbf{G}_x$ on both diagonal and cross terms that arise from these parts. In a bit more detail, we express $|\psi^x\rangle$ as follows:

$$|\psi^x\rangle = |\psi^x_{\mathrm{HH}}\rangle + |\psi^x_{\mathrm{HL}}\rangle + |\psi^x_{\mathrm{LH}}\rangle + |\psi^x_{\mathrm{LL}}\rangle$$

Informally speaking,

- $|\psi^x_{\mathrm{HH}}\rangle$: on this component, both $\mathcal{B}$ and $\mathcal{C}$ succeed with probability above a certain threshold τ.
- $|\psi^x_{\mathrm{LL}}\rangle$: on this component, both $\mathcal{B}$ and $\mathcal{C}$ succeed with probability below τ.
- $|\psi^x_{\mathrm{HL}}\rangle$: on this component, $\mathcal{B}$ succeeds with probability above τ whereas $\mathcal{C}$ succeeds with probability below τ.
- $|\psi^x_{\mathrm{LH}}\rangle$: on this component, $\mathcal{B}$ succeeds with probability below τ whereas $\mathcal{C}$ succeeds with probability above τ.

Thus, we can write p_{pred} as:

$$p_{\mathsf{pred}} = \underbrace{\sum_{U\in\{\mathrm{HH,HL,LH,LL}\}} \mathbb{E}_x\left[\langle\psi^x_U|\mathbf{G}_x|\psi^x_U\rangle\right]}_{\text{diagonal terms}} + \underbrace{\sum_{\substack{U,V\in\{\mathrm{HH,HL,LH,LL}\}\\ U\neq V}} \mathbb{E}_x\left[\langle\psi^x_U|\mathbf{G}_x|\psi^x_V\rangle\right]}_{\text{cross terms}}$$

The rest of the proof is to upper bound each of these terms. There are two relatively simple cases:

- A term is of the form $\langle\psi^x_U|G_x|\psi^x_V\rangle$ when either $U = \mathrm{HH}$ or $V = \mathrm{HH}$: we can upper bound these terms by designing a non-local adversary that can simultaneously predict x given $|\psi^x\rangle$ whose success probability is in turn upper bounded by[3] p_{srch}.
- The averaged norm of any of the terms $|\psi^x_{\mathrm{HL}}\rangle, |\psi^x_{\mathrm{LH}}\rangle$ or $|\psi^x_{\mathrm{LL}}\rangle$ is low: in this case, it can be shown that $\langle\psi^x_U|G_x|\psi^x_V\rangle$ has to be low as well. In this case, we can safely ignore these terms in the above analysis.

We then strip away the terms from $|\psi^x\rangle$ that capture the above cases and then we are left with the state $|\phi^x\rangle$. It now suffices to determine p_{pred} in terms of $\mathbf{G}_x$. We in turn express $|\phi^x\rangle$ as the following sum:

$$|\phi^x\rangle = |\phi^x_{\mathrm{HL}}\rangle + |\phi^x_{\mathrm{LH}}\rangle + |\phi^x_{\mathrm{LL}}\rangle,$$

where $|\phi^x_U\rangle$ is either $|\psi^x_U\rangle$ or 0, depending on the second bullet point above.

[3] We are being loose here; technically, we suffer from a multiplicative loss.

Instead of determining an upper bound for $\mathbb{E}_x\left[\langle\phi^x|\mathbf{G}_x|\phi^x\rangle\right]$, we instead define a new operator S_x: $S_x := \left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x$, where η is some function of p_{srch}. For instance, if p_{srch} is negligible then so is η. Now, the new goal[4] is to show that $\mathbb{E}_x\left[\langle\phi^x|S_x|\phi^x\rangle\right] \geq 0$. Since $\|\phi^x\|^2 \leq 1$, this implies that $\mathbb{E}_x\left[\langle\phi^x|\mathbf{G}_x|\phi^x\rangle\right] \leq \frac{1}{2}+\eta$.

Let us expand out $\mathbb{E}_x\left[\langle\phi^x|S_x|\phi^x\rangle\right]$:

$$\mathbb{E}_x\left[\langle\phi^x|S_x|\phi^x\rangle\right] = \sum_{U,V\in\{\mathrm{HL},\mathrm{LH},\mathrm{LL}\}} \mathbb{E}_x\left[\langle\phi^x_U|S_x|\phi^x_V\rangle\right]$$

Informally speaking, the terms involving $|\phi^x_{\mathrm{LL}}\rangle$ should not contribute much to the sum; this is because $|\phi^x_{\mathrm{LL}}\rangle$ is precisely the component wherein both $\mathcal{B}$ and $\mathcal{C}$ fail to succeed above a certain threshold. This suggests that we could minimize over all possible $|\phi^x_{\mathrm{LL}}\rangle$ wherein $\mathcal{B}$ and $\mathcal{C}$'s success probability is below τ. If we perform the minimization, we end up showing that:

$$\mathbb{E}_x\left[\langle\phi^x|S_x|\phi^x\rangle\right] \geq \left(\eta - \frac{\eta^3}{2}\right)\left(\underbrace{\mathbb{E}_x\left[\||\phi^x_{\mathrm{HL}}\rangle\|^2\right]}_{\text{weight of HL term}} + \underbrace{\mathbb{E}_x\left[\||\phi^x_{\mathrm{LH}}\rangle\|^2\right]}_{\text{weight of LH term}}\right) - 2|\mathbb{E}_x\left[c_x\right]|,$$

where c_x is a cross term involving $|\phi^x_{\mathrm{HL}}\rangle$ and $|\phi^x_{\mathrm{LH}}\rangle$. Thus, we effectively eliminated cross terms involving $|\phi^x_{\mathrm{LL}}\rangle$ and are left with analyzing cross terms involving $|\phi^x_{\mathrm{HL}}\rangle$ and $|\phi^x_{\mathrm{LH}}\rangle$. We clarify that c_x is of the form $\langle\phi^x_{\mathrm{HL}}|\mathbf{M}|\phi^x_{\mathrm{LH}}\rangle$, where $\mathbf{M}$ is not necessarily S_x.

Finally, it remains to show that $2|\mathbb{E}_x\left[c_x\right]| \leq \left(\eta - \frac{\eta^3}{2}\right)\left(\mathbb{E}_x\left[\||\phi^x_{\mathrm{HL}}\rangle\|^2\right] + \mathbb{E}_x\left[\||\phi^x_{\mathrm{LH}}\rangle\|^2\right]\right)$; showing this would imply that $\mathbb{E}_x\left[\langle\phi^x|S_x|\phi^x\rangle\right] \geq 0$. If not, then we use the fact that $2|\mathbb{E}_x\left[c_x\right]|$ is large enough to come up with a non-local adversary that can simultaneously predict x with probability larger than p_{srch}. This then would complete the proof.

For the rest of the overview, we break down the above template into distinct steps and also provide the appropriate pointers to the technical sections.

Step 0 [Section 3 and Section 3.1]: Expanding out $\mathbf{G}_x$. The first task is to expand out the description of the operator $\mathbf{G}_x$. Recall that:

$$\mathbf{G}_x = \mathbb{E}_r\left[\mathcal{P}^{r,\langle r,x\rangle}_{\mathcal{B}} \otimes \mathcal{P}^{r,\langle r,x\rangle}_{\mathcal{C}}\right] \otimes I_{\mathrm{E}}$$

For notational simplicity, we ignore I_{E} from here onward. We define the following operators: $\Delta^r_{\mathcal{B}}$ to be $\mathcal{P}^{r,0}_{\mathcal{B}} - \mathcal{P}^{r,1}_{\mathcal{B}}$; similarly, define $\Delta^r_{\mathcal{C}}$ to be $\mathcal{P}^{r,0}_{\mathcal{C}} - \mathcal{P}^{r,1}_{\mathcal{C}}$. Thus, we can express $\mathbf{G}_x$ as follows:

$$\begin{aligned}
\mathbf{G}_x &= \mathbb{E}_r\left[\mathcal{P}^{r,\langle r,x\rangle}_{\mathcal{B}} \otimes \mathcal{P}^{r,\langle r,x\rangle}_{\mathcal{C}}\right] \\
&= \mathbb{E}_r\left[\frac{\left(I_{\mathrm{B}} + (-1)^{\langle r,x\rangle}\Delta^r_{\mathcal{B}}\right)}{2} \otimes \frac{\left(I_{\mathrm{C}} + (-1)^{\langle r,x\rangle}\Delta^r_{\mathcal{C}}\right)}{2}\right] \\
&= \frac{1}{4}\left(I_{\mathrm{BC}} + \underbrace{\mathbb{E}_r\left[(-1)^{\langle r,x\rangle}\Delta^r_{\mathcal{B}}\right]}_{\mathbf{X}^x} \otimes I_{\mathrm{C}} + I_{\mathrm{B}} \otimes \underbrace{\mathbb{E}_r\left[(-1)^{\langle r,x\rangle}\Delta^r_{\mathcal{C}}\right]}_{\mathbf{Y}^x} + \underbrace{\mathbb{E}_r\left[\Delta^r_{\mathcal{B}} \otimes \Delta^r_{\mathcal{C}}\right]}_{\mathbf{Z}}\right) \\
&= \frac{1}{4}\left(I_{\mathrm{BC}} + \mathbf{X}^x \otimes I_{\mathrm{C}} + I_{\mathrm{B}} \otimes \mathbf{Y}^x + \mathbf{Z}\right)
\end{aligned}$$

[4] Although the approach works, the author found the step of analyzing the action of S_x instead of $\mathbf{G}_x$ to not be a natural one.

The reason behind expanding out $\mathbf{G}_x$ this way will be clear in the next step: this will help us decompose $|\psi^x\rangle$ into the four components corresponding to HH, HL, LH and LL, as described earlier.

Step 1 [Section 3.1]: Decomposing $|\psi^x\rangle$ into four components HH, HL, LH and LL. Since $\mathbf{X}^x$ is Hermitian, we can look at the spectral decomposition of $\mathbf{X}^x$. Let $\{|u_i\rangle\}$ be an orthonormal basis such that:

$$\mathbf{X}^x = \sum_{i:\lambda_i} \lambda_i \, |u_i\rangle\langle u_i|$$

We ignore the terms where $|\lambda_i| < \tau$, where τ is some function defined as a function of p_{srch}. Define the projector Π_x as follows: $\Pi_x = \sum_{|\lambda_i|\geq\tau} |u_i\rangle\langle u_i|$. Similarly, define $\Theta_x = \sum_{|\gamma_i|\geq\tau} |v_i\rangle\langle v_i|$, where $\mathbf{Y}^x = \sum_i \gamma_i \, |v_i\rangle\langle v_i|$. Intuitively, Π_x projects $\mathcal{B}$'s state onto $\mathcal{B}$'s eigenstates with absolute eigenvalues at least τ and similarly, Θ_x projects $\mathcal{C}$'s state onto $\mathcal{C}$'s eigenstates with absolute eigenvalues at least τ. Now, define the complement projectors, namely $N_x = I - \Pi_x$ and $O_x = I - \Theta_x$. Using these projectors, we can correspondingly define the four product projectors:

$$P^x_{\mathrm{HH}} = \Pi_x \otimes \Theta_x, \quad P^x_{\mathrm{HL}} = \Pi_x \otimes O_x,$$
$$P^x_{\mathrm{LH}} = N_x \otimes \Theta_x, \quad P^x_{\mathrm{LL}} = N_x \otimes O_x.$$

We are now ready to define the different components of $|\psi^x\rangle$:

- $|\psi^x_{\mathrm{HH}}\rangle = (\Pi_x \otimes \Theta_x)\,|\psi^x\rangle$,
- $|\psi^x_{\mathrm{HL}}\rangle = (\Pi_x \otimes O_x)\,|\psi^x\rangle$,
- $|\psi^x_{\mathrm{LH}}\rangle = (N_x \otimes \Theta_x)\,|\psi^x\rangle$,
- $|\psi^x_{\mathrm{LL}}\rangle = (N_x \otimes O_x)\,|\psi^x\rangle$,

Since the above orthogonal projectors sum to the identity, we have: $|\psi^x\rangle = |\psi^x_{\mathrm{HH}}\rangle + |\psi^x_{\mathrm{HL}}\rangle + |\psi^x_{\mathrm{LH}}\rangle + |\psi^x_{\mathrm{LL}}\rangle$.

Substituting this decomposition into the prediction probability produces four diagonal terms and six cross terms:

$$\begin{aligned} p_{\mathsf{pred}} = \sum_{U\in\{\mathrm{HH},\mathrm{HL},\mathrm{LH},\mathrm{LL}\}} & \mathbb{E}_x \langle\psi^x_U|\mathbf{G}_x|\psi^x_U\rangle \\ & + 2\,\mathrm{Re}\,\mathbb{E}_x\Big[\langle\psi^x_{\mathrm{HH}}|\mathbf{G}_x|\psi^x_{\mathrm{HL}}\rangle + \langle\psi^x_{\mathrm{HH}}|\mathbf{G}_x|\psi^x_{\mathrm{LH}}\rangle + \langle\psi^x_{\mathrm{HH}}|\mathbf{G}_x|\psi^x_{\mathrm{LL}}\rangle \\ & \qquad + \langle\psi^x_{\mathrm{HL}}|\mathbf{G}_x|\psi^x_{\mathrm{LH}}\rangle + \langle\psi^x_{\mathrm{HL}}|\mathbf{G}_x|\psi^x_{\mathrm{LL}}\rangle + \langle\psi^x_{\mathrm{LH}}|\mathbf{G}_x|\psi^x_{\mathrm{LL}}\rangle\Big]. \end{aligned}$$

Step 2 [Sections 3.2–3.3]: Removing the easy and negligible sectors. We first explain why the HH component must have small average weight. It can be shown that the set of operators $\{\tau^2\Pi_x\}_x$ has the following properties: each of $\tau^2\Pi_x$ is positive semidefinite; moreover, $\sum_x \tau^2\Pi_x \leq I_{\mathrm{B}}$. Thus, we can define a POVM $\Gamma_{\mathcal{B}} = \left\{\left\{\tau^2\Pi_x\right\}_x,\ I_{\mathrm{B}} - \sum_x \tau^2\Pi_x\right\}$ that can be applied to $\mathcal{B}$'s state. Similarly, we can define a POVM $\Gamma_{\mathcal{C}} = \left\{\left\{\tau^2\Theta_x\right\}_x,\ I_{\mathrm{C}} - \sum_x \tau^2\Theta_x\right\}$ that can be applied to $\mathcal{C}$'s state.

This yields a simultaneous search strategy. Bob measures his register using $\Gamma_{\mathcal{B}}$ and outputs the measurement outcome; Charlie does the same with his own measurement $\Gamma_{\mathcal{C}}$. When the hidden string is x, the probability that both conclusive outcomes equal x is

$$\begin{aligned}\langle\psi^x|\big(\tau^2\Pi_x \otimes \tau^2\Theta_x \otimes I_{\mathbb{E}}\big)|\psi^x\rangle &= \tau^4\big\|(\Pi_x \otimes \Theta_x \otimes I_{\mathbb{E}})|\psi^x\rangle\big\|^2 \\ &= \tau^4\big\||\psi^x_{\mathrm{HH}}\rangle\big\|^2.\end{aligned}$$

Thus its average success probability is at least $\tau^4\mathbb{E}_x\||\psi^x_{\mathrm{HH}}\rangle\|^2$. Since p_{srch} is the optimal success probability of any simultaneous full-string search strategy, this quantity is at most p_{srch}. Therefore:

$$\mathbb{E}_x\big\||\psi^x_{\mathrm{HH}}\rangle\big\|^2 \le \frac{p_{\mathsf{srch}}}{\tau^4}$$

We now set aside $|\psi^x_{\mathrm{HH}}\rangle$ for the remainder of the calculation and charge every diagonal or cross term containing it to an error term. We also set aside any one of the $\mathrm{HL}, \mathrm{LH}, \mathrm{LL}$ components whose average squared norm is below the same threshold τ.[5] This weight cutoff ensures that every component retained for the later extraction argument has enough average mass; otherwise normalizing that component could incur an uncontrolled loss. These are only mathematical deletions used in the proof; Bob and Charlie do not project or modify their state. Since $0 \preceq \mathbf{G}_x \preceq I$, Cauchy–Schwarz bounds the total change caused by all deleted terms in terms of their total squared norm, which is at most $p_{\mathsf{srch}}/\tau^4 + 3\tau$ in the simplified notation of this overview. Let

$$|\phi^x\rangle = |\phi^x_{\mathrm{HL}}\rangle + |\phi^x_{\mathrm{LH}}\rangle + |\phi^x_{\mathrm{LL}}\rangle$$

be the retained vector, and write

$$W_U := \mathbb{E}_x\big\||\phi^x_U\rangle\big\|^2 \qquad (U \in \{\mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}).$$

Each W_U is either zero or at least τ. If $W_{\mathrm{HL}} = 0$, every retained component is low for Bob, so $\langle\phi^x|\mathbf{G}_x|\phi^x\rangle$ is at most $((1+\tau)/2)\|\phi^x\|^2$. Averaging and using $\mathbb{E}_x\left[\|\phi^x\|^2\right] \le 1$ gives the desired bound. Similarly, if $W_{\mathrm{LH}} = 0$, Charlie's marginal gives the same conclusion. It remains to study the case $W_{\mathrm{HL}}, W_{\mathrm{LH}} \ge \tau$.

Step 3 [Section 3.3]: Accounting for the LL **cross terms.** Our target for the retained vector is

$$\mathbb{E}_x\langle\phi^x|\mathbf{G}_x|\phi^x\rangle \le \frac{1}{2} + \eta.$$

It is convenient to measure the shortfall from this target. Define

$$S_x := \left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x.$$

It suffices to prove $\mathbb{E}_x\langle\phi^x|S_x|\phi^x\rangle \ge 0$: this bounds the retained success by $(1/2+\eta)\mathbb{E}_x\|\phi^x\|^2 \le 1/2+\eta$. Moreover, whenever Charlie is in his low subspace, the probability that both parties are correct is at most Charlie's marginal probability, which is at most $(1+\tau)/2$. Thus S_x has a positive lower bound on $\mathrm{HL} \oplus \mathrm{LL}$. The analogous statement holds on $\mathrm{LH} \oplus \mathrm{LL}$ because Bob is low there.

[5] For readability, the overview uses one symbol for both cutoffs. The technical proof uses two parameters: τ cuts off the magnitude of the local eigenvalues, whereas δ cuts off the average weight of an entire sector family. Keeping them separate allows the two resulting error terms to be optimized independently.

The cross terms HL–LL and LH–LL need not be small by themselves. Bounding their absolute values separately would lose the positive LL diagonal term that controls them. Instead, fix the HL and LH components and collect every term that depends on $|\ell\rangle := |\phi^x_{\mathrm{LL}}\rangle$. It has the form

$$\langle \ell | M_x | \ell \rangle - 2\,\mathrm{Re}\langle \ell \mid u_x \rangle,$$

where M_x is the restriction of S_x to the LL subspace and $|u_x\rangle$ collects the two couplings from LL to HL and LH. The positive margin just discussed makes M_x invertible on this subspace. Completing the square gives

$$(\langle \ell| - \langle u_x|M_x^{-1})M_x(|\ell\rangle - M_x^{-1}|u_x\rangle) - \langle u_x|M_x^{-1}|u_x\rangle.$$

The first term is nonnegative. The most unfavorable possible LL component is therefore obtained at $|\ell\rangle = M_x^{-1}|u_x\rangle$. This explains both why the LL vector disappears from the lower bound and why M_x^{-1} appears. This is the matrix version of completing the square, also called taking a Schur complement.

Expanding the correction $\langle u_x|M_x^{-1}|u_x\rangle$ leaves positive diagonal contributions from HL and LH, together with one effective interaction c_x between them. Claim 3.4 makes this precise:

$$\mathbb{E}_x\langle \phi^x|S_x|\phi^x\rangle \geq \left(\eta - \frac{\eta^3}{2}\right)(W_{\mathrm{HL}} + W_{\mathrm{LH}}) - 2\,|\mathbb{E}_x c_x|\,.$$

Thus the original three-sector problem has been reduced to controlling one averaged HL–LH interaction.

Step 4 [Section 3.4]: Rewriting the remaining interaction as a sum of finite paths. Set $\alpha := \eta - \eta^3/2$. If

$$|\mathbb{E}_x c_x| \leq \alpha\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}},$$

then

$$2\,|\mathbb{E}_x c_x| \leq 2\alpha\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}} \leq \alpha(W_{\mathrm{HL}} + W_{\mathrm{LH}}),$$

where the last inequality is $2\sqrt{ab} \leq a + b$. This cancels the only negative term in the lower bound from Step 3. It therefore remains to prove the displayed estimate on c_x.

Recall that

$$\mathbf{G}_x = \frac{1}{4}\big(I + \mathbf{X}^x \otimes I + I \otimes \mathbf{Y}^x + \mathbf{Z}\big).$$

The high and low subspaces were defined from the spectral decompositions of $\mathbf{X}^x$ and $\mathbf{Y}^x$. Hence $\mathbf{X}^x$ cannot move Bob's state from his high subspace to his low subspace, and $\mathbf{Y}^x$ cannot do so on Charlie's side. The identity changes neither label. Consequently, whenever a matrix element of $\mathbf{G}_x$ connects two different sectors among $\mathrm{HL}, \mathrm{LH}, \mathrm{LL}$, the identity and the two one-party terms vanish by orthogonality; only $\mathbf{Z}/4$ remains. Thus the direct HL–LH interaction and the two couplings through LL can all be written using $\mathbf{Z}$.

The inverse M_x^{-1} from Step 3 also contains the restrictions of $\mathbf{X}^x$ and $\mathbf{Y}^x$ to the LL subspace. These restrictions have norm at most τ, precisely because LL is low on both sides. A standard inverse-perturbation estimate lets us replace them at a controlled cost. More precisely, writing

$$A_x := (1 + 4\eta)P^x_{\mathrm{LL}} - Q_x, \qquad Q_x := P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{LL}},$$

the technical proof has $4M_x = A_x - R_x$ for a small perturbation R_x and shows that M_x^{-1} is close to $4A_x^{-1}$. Here Q_x is simply the part of $\mathbf{Z}$ that starts in the LL subspace and ends there again. Since $\|\mathbf{Z}\| \le 1$, we have $\|Q_x\| \le 1$. The familiar scalar geometric series

$$\frac{1}{a-t} = \frac{1}{a} + \frac{t}{a^2} + \frac{t^2}{a^3} + \cdots \qquad (|t| < a)$$

has the following operator analogue on the LL subspace:

$$\big((1+4\eta)P_{\mathrm{LL}}^x - Q_x\big)^{-1} = \sum_{k\ge 0}(1+4\eta)^{-(k+1)}Q_x^k.$$

This is the Neumann, or operator geometric, expansion. Its purpose here is concrete: every power Q_x^k describes k successive applications of $\mathbf{Z}$, with the state required to return to LL between applications.

We record these finite paths by scalars Γ_m. The first one is the direct transition from LH to HL through one copy of $\mathbf{Z}$:

$$\Gamma_1 := \mathbb{E}_x \langle \phi_{\mathrm{HL}}^x | \mathbf{Z} | \phi_{\mathrm{LH}}^x \rangle.$$

For $m \ge 2$, define

$$\Gamma_m := \mathbb{E}_x \langle \phi_{\mathrm{HL}}^x | \mathbf{Z} P_{\mathrm{LL}}^x (P_{\mathrm{LL}}^x \mathbf{Z} P_{\mathrm{LL}}^x)^{m-2} P_{\mathrm{LL}}^x \mathbf{Z} | \phi_{\mathrm{LH}}^x \rangle.$$

This expression starts in LH, returns to LL after every intermediate application of $\mathbf{Z}$, and reaches HL after the m-th copy. Claim 3.5 says that, up to a remainder of size at most $\frac{\eta}{16}\sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}}$, the scalar $\mathbb{E}_x c_x$ is a geometrically weighted sum of the Γ_m's. We have therefore replaced a difficult inverse by finite expressions built from copies of the original common-challenge operator $\mathbf{Z}$, separated by LL projections.

Step 5 [Sections 3.5–3.6]: Bounding every finite path by search security. Fix m and consider Γ_m. Replacing every occurrence of

$$\mathbf{Z} = \mathbb{E}_r[\Delta_{\mathcal{B}}^r \otimes \Delta_{\mathcal{C}}^r]$$

by its definition expresses Γ_m as an average over a word of challenges $w = (r_1, \ldots, r_m)$. Once w is fixed, the corresponding operator factors as

$$V_{x,w}^{\mathcal{B}} \otimes V_{x,w}^{\mathcal{C}}.$$

The first factor acts only on Bob's Hilbert space and the second only on Charlie's. Each r_i is a common challenge for the two tensor factors, while the different r_i's are sampled independently. The word w is independent of the hidden string and may be fixed as shared randomness before the parties separate. The label x will become a possible measurement outcome; it is not supplied to either party.

The maps $V_{x,w}^{\mathcal{B}}$ and $V_{x,w}^{\mathcal{C}}$ are not themselves measurement effects, but their adjoint-products are positive. The high-projector estimate from Step 2, together with an expansion of the intermediate low projectors, gives

$$\sum_x (V_{x,w}^{\mathcal{B}})^\dagger V_{x,w}^{\mathcal{B}} \preceq (m+1)^2 \tau^{-2} I, \qquad \sum_x (V_{x,w}^{\mathcal{C}})^\dagger V_{x,w}^{\mathcal{C}} \preceq \tau^{-2} I.$$

The first bound contains $m+1$ because a telescoping identity removes Bob's m intermediate low projectors one at a time. This produces one base term and m correction terms; each is bounded using one high projector Π_x and the estimate from Step 2. On Charlie's side, the high projector is on the input side: $V^{\mathcal{C}}_{x,w} = K^{\mathcal{C}}_w \Theta_x$ for a contraction $K^{\mathcal{C}}_w$. Hence $(V^{\mathcal{C}}_{x,w})^\dagger V^{\mathcal{C}}_{x,w} \preceq \Theta_x$ directly, so no analogous $m+1$ loss appears.

It follows that

$$F^x_{\mathcal{B},w} := \frac{\tau^2}{(m+1)^2}(V^{\mathcal{B}}_{x,w})^\dagger V^{\mathcal{B}}_{x,w}, \qquad F^x_{\mathcal{C},w} := \tau^2 (V^{\mathcal{C}}_{x,w})^\dagger V^{\mathcal{C}}_{x,w}$$

form local sub-POVMs indexed by candidate strings x; residual effects complete them to POVMs. Search security bounds the probability that these two measurements both output the true hidden string. This controls the squared norm of the path applied to $|\psi^x\rangle$. Since Γ_m is an amplitude rather than a squared norm, a final Cauchy–Schwarz inequality takes a square root and gives

$$|\Gamma_m| \le \frac{(m+1)\sqrt{p_{\mathsf{srch}}}}{\tau^2}.$$

Substituting these bounds into the geometrically weighted sum from Step 4 controls $\mathbb{E}_x c_x$. The inequality $2\sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}} \le W_{\mathrm{HL}} + W_{\mathrm{LH}}$ then makes the lower bound in Step 3 nonnegative. Finally, adding back the error incurred when the small sectors were removed in Step 2 proves the theorem after choosing η, τ, and δ as specified in the technical proof.

1.1.2 Search-to-Decision over Large Finite Fields

The large-field result follows a different route from the four-sector argument over $\mathbb{F}_2$. We first compare the common-challenge experiment with an experiment in which Bob and Charlie receive independent challenge vectors. We then give a direct finite-field Fourier procedure that recovers the hidden vectors from predictors for the independent-challenge experiment. The Fourier procedure works over every finite field; the assumption that the field is large is needed only at the end, when the independent theorem's $1/q$ baseline passes through the cube-root common-to-independent comparison. In this subsection, p_{com} denotes the finite-field analogue of the quantity called p_{pred} above.

Step 1 [Section 4.3]: From a common challenge to independent challenges. Let p_{ind} be the success probability of the same local response measurements when Bob and Charlie receive independent uniform challenges. The correlated-to-independent inequality of Ananth, Kaleoglu, and Liu [AKL23, Lemma 18] gives

$$p_{\mathrm{com}} \le 6 p_{\mathrm{ind}}^{1/3}, \qquad \text{or equivalently} \qquad p_{\mathrm{ind}} \ge p_{\mathrm{com}}^3/216.$$

Thus sufficiently large success with a common challenge forces some success when the two challenges are sampled independently. This step changes only the challenge distribution in the analysis; the local response measurements remain the same. It is also the source of the cube-root loss in the final common-challenge bound.

Step 2 [Section 4.2]: What the independent-challenge extractor detects. It is useful to state this part in the slightly more general two-secret form used in Section 4.2, where the hidden vectors have dimension ℓ over $\mathbb{F}_q$. Suppose Bob's hidden vector is x, Charlie's is y, their independent challenges are $v, w \in \mathbb{F}_q^\ell$, and their answers are $g, h \in \mathbb{F}_q$. Define their additive errors by

$$E^{\mathsf{B}} := g - \langle x, v\rangle, \qquad E^{\mathsf{C}} := h - \langle y, w\rangle.$$

If both answers are correct, then $E^{\mathsf{B}} = E^{\mathsf{C}} = 0$, and in particular $E^{\mathsf{B}} + E^{\mathsf{C}} = 0$. The proof establishes the stronger statement that an advantage in satisfying this last relation already implies joint recovery of x and y. This strengthening is useful because the sum of the two errors is exactly what the Fourier calculation detects.

To encode a field element as a complex phase, fix a nontrivial additive character $\chi : \mathbb{F}_q \to \mathbb{C}$. It satisfies $|\chi(s)| = 1$ and $\chi(s+t) = \chi(s)\chi(t)$, and therefore plays the role of the binary phase $(-1)^s$. When $q = p^t$, one concrete choice is

$$\chi(s) = \exp\left(\frac{2\pi i}{p} \operatorname{Tr}_{\mathbb{F}_q/\mathbb{F}_p}(s)\right).$$

Fix a nonzero multiplier $a \in \mathbb{F}_q^*$. Bob prepares a uniform superposition over challenges v, runs his predictor coherently, gives a branch with answer g the phase $\chi(ag)$, runs the predictor backward, and applies the inverse additive Fourier transform. If the predictor were perfect, it would insert the phase

$$\chi(a\langle x, v\rangle) = \chi(\langle ax, v\rangle).$$

The inverse Fourier transform maps this entire phase pattern to the label ax; multiplying by a^{-1} returns x. Charlie performs the same local procedure and obtains y. This is the finite-field analogue of Hadamard-basis extraction in the quantum Goldreich–Levin argument [AC02]; the finite-field Fourier transform itself is standard [BCW02].

Step 3 [Section 4.2]: Why the multiplier is shared and averaged. When the predictors make errors, Bob's phase is the ideal phase pointing to ax, multiplied by $\chi(aE^{\mathsf{B}})$, and Charlie obtains the analogous factor $\chi(aE^{\mathsf{C}})$. Because both parties use the same multiplier, their error phases combine as

$$\chi(aE^{\mathsf{B}})\chi(aE^{\mathsf{C}}) = \chi\Big(a(E^{\mathsf{B}} + E^{\mathsf{C}})\Big).$$

For a fixed hidden pair $z = (x, y)$, let

$$A_z(s) := \Pr\Big[E^{\mathsf{B}} + E^{\mathsf{C}} = s \mid z\Big]$$

be the conditional total-error distribution. In the extraction analysis, the overlap of the all-zero-ancilla successful-output component with the original purification is

$$\Phi_z(a) := \mathbb{E}\Big[\chi\Big(a(E^{\mathsf{B}} + E^{\mathsf{C}})\Big) \mid z\Big] = \sum_{s \in \mathbb{F}_q} A_z(s)\chi(as).$$

Thus $\Phi_z(a)$ is the additive Fourier coefficient of the total-error distribution in direction a. The extraction success conditioned on (z, a) is at least $|\Phi_z(a)|^2$.

Over $\mathbb{F}_2$ there is only one nonzero multiplier. Over a larger field, however, different nonzero errors may contribute different complex phases that cancel for a particular a. We avoid depending on one such coefficient by averaging over every $a \in \mathbb{F}_q^*$. Additive-character orthogonality gives the exact identity

$$\mathbb{E}_{a\in\mathbb{F}_q^*}|\Phi_z(a)|^2 = \frac{q\sum_{s\in\mathbb{F}_q} A_z(s)^2 - 1}{q-1}.$$

The sum $\sum_s A_z(s)^2$ is the collision probability of the total error. It equals $1/q$ when that error is uniform, and increases when extra probability is placed on zero. In particular, if $A_z(0) \geq 1/q + \epsilon/2$, then the last display is at least $q^2\epsilon^2/(4(q-1)^2)$. The hypothesis that the overall probability of $E^{\mathsf{B}} + E^{\mathsf{C}} = 0$ is $1/q + \epsilon$ implies that such good hidden pairs have total probability at least $\epsilon/2$. Averaging over those pairs and over a therefore gives joint recovery probability at least

$$\frac{q^2}{8(q-1)^2}\epsilon^3.$$

In particular, some single multiplier works for the entire hidden-pair distribution and can be fixed before the parties separate. The nonzero multiplier phase kick and final multiplication by its inverse are adapted from large-modulus one-party extraction [STH+23]; using the same multiplier on both sides and averaging over it is the new two-party step. The independent-challenge framework follows Kundu and Tan [KT25].

Step 4 [Section 4.3]: Combining the two stages. Applying the independent-challenge extraction bound after Step 1 yields

$$p^\star_{\text{search}} \geq \frac{q^2}{8(q-1)^2}\left(\frac{p^3_{\text{com}}}{216} - \frac{1}{q}\right)_+^3,$$

or, equivalently,

$$p_{\text{com}} \leq 6\left(\frac{1}{q} + 2\left(\frac{q-1}{q}\right)^{2/3}(p^\star_{\text{search}})^{1/3}\right)^{1/3}.$$

Here $p^\star_{\text{search}}$ is the finite-field analogue of p_{srch}. If $p^\star_{\text{search}}$ is negligible, the explicit upper bound is $6q^{-1/3} + \text{negl}(n)$. Hence it becomes negligible when q is exponential in the security parameter. This pinpoints where the large-field assumption is used: the Fourier extraction works over every finite field, but the common-to-independent comparison leaves the residual term $6q^{-1/3}$. For fixed or moderately sized fields, the same quantitative theorem remains valid, but it does not control prediction probabilities close to the natural $1/q$ baseline.

Returning to the binary case, Figure 3 summarizes the logical flow of the binary proof and indicates where each step is carried out in the technical section.

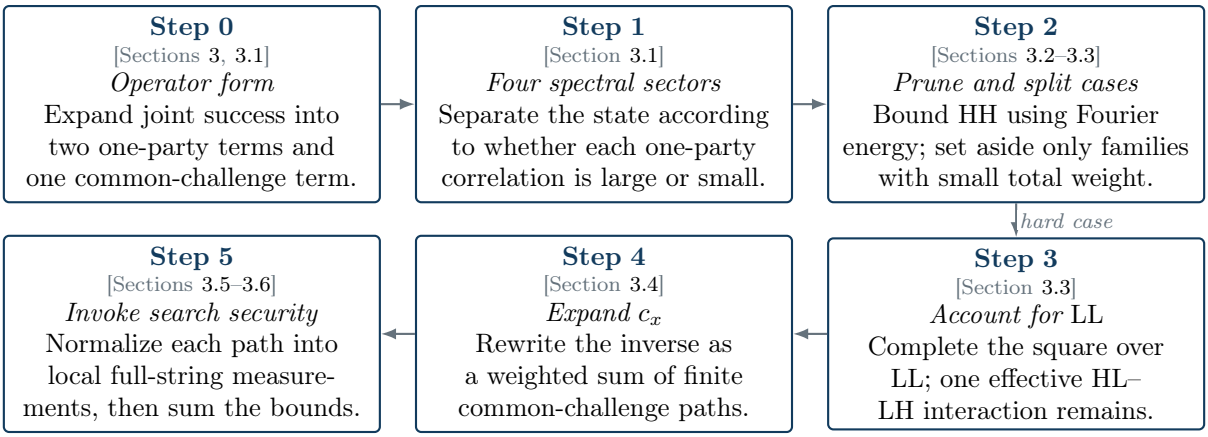


Figure 3: Flow of the proof. Steps 0–2 reformulate the prediction experiment and isolate the only hard case. Steps 3–5 use positivity of the LL block to turn the surviving interaction into common-challenge paths bounded by joint search security.

1.2 Related Work

Different Avatars of Non-Local Search-to-Decision. The non-local search-to-decision problem considered in this work can be cast as an instance of the generalized non-local search-to-decision problem, described in Figure 4. In the f-generalized non-local decision problem, $\mathcal{B}$ and $\mathcal{C}$ respectively receive $r_{\mathcal{B}}$ and $r_{\mathcal{C}}$, where $(r_{\mathcal{B}}, r_{\mathcal{C}})$ are sampled from a distribution. Moreover, the winning condition is characterized by a predicate f that takes as input $(b_{\mathcal{B}}, b_{\mathcal{C}}, \langle r_{\mathcal{B}}, x\rangle, \langle r_{\mathcal{C}}, x\rangle)$. If f outputs 0 then we say that $\mathcal{B}$ and $\mathcal{C}$ win the game.

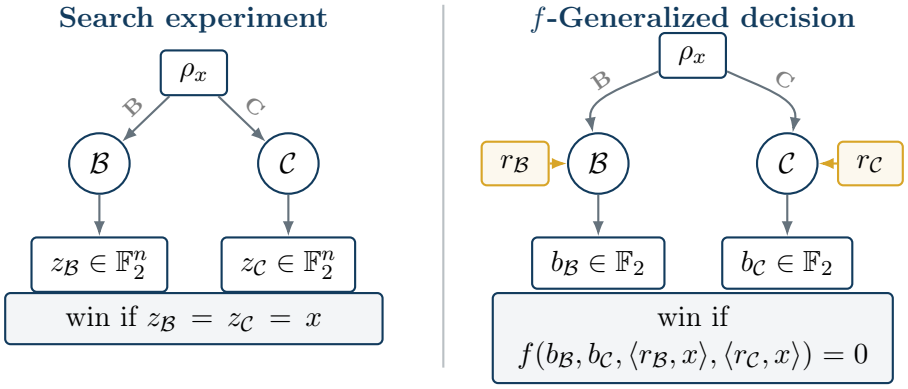


Figure 4: The search and the identical-challenge decision experiments.

We note other instances of the above generalized problem considered in the literature:

- **Independent**: in this case, $(r_{\mathcal{B}}, r_{\mathcal{C}})$ are sampled from a product distribution on $\mathbb{F}_2^n$. Both $\mathcal{B}$ and $\mathcal{C}$ win if $b_{\mathcal{B}} = \langle r_{\mathcal{B}}, x\rangle$ and $b_{\mathcal{C}} = \langle r_{\mathcal{C}}, x\rangle$. This setting was explored in the works [AKL23, KT25]. Specifically, they present a simultaneous Goldreich-Levin reduction that works for non-local adversaries in the independent challenge setting.

- **Correlated**: similar to the above case, $(r_{\mathcal{B}}, r_{\mathcal{C}})$ are sampled from a product distribution on $\mathbb{F}_2^n$. However, the winning condition is a relaxation of the above condition: $\mathcal{B}$ and $\mathcal{C}$ win if $b_{\mathcal{B}} \oplus b_{\mathcal{C}} = \langle r_{\mathcal{B}}, x\rangle \oplus \langle r_{\mathcal{C}}, x\rangle$. This setting was explored in [AKY25].

Other variants of the f-generalized decision problem have also been considered in the literature. Instead of requiring $\mathcal{B}$ and $\mathcal{C}$ to predict the inner product, they participate in a distinguishing game. That is, $\mathcal{B}$ (and similarly, $\mathcal{C}$) receives either $\langle r_{\mathcal{B}}, x\rangle$ or a uniformly random element in $\mathbb{F}_2$ and they need to distinguish. The work [CG24a] considered a multi-sample distinguishing variant, wherein both $\mathcal{B}$ and $\mathcal{C}$ receive $o(n)$ samples of the random inner product $\langle r, x\rangle$. The work [AB24] considered the distinguishing version over an arbitrary finite field, instead of just $\mathbb{F}_2$.

Applications. Search-to-decision problems have been useful in different contexts in unclonable cryptography. Apart from unclonable encryption, they have also been used in the proofs of quantum copy-protection [AB24]. Roughly speaking, copy-protection is a compiler that converts a program into a functionality-preserving unclonable quantum state. Other applications include leakage-resilience [CG24b] and telegraphing schemes [CKN+25].

Concurrent Works. Concurrent and independent work of Coladangelo, Liu, and Xie [CLX26b] proves a stronger theorem for the $\mathbb{F}_2$ setting. Writing $\Delta = 2p_{\mathsf{pred}} - 1$, their reduction gives the sharper guarantee $p_{\mathsf{srch}} \geq \Delta^4$ and, when supplied with an efficiently computable lower bound $0 < \gamma \leq \Delta$, admits a uniform implementation whose running time is polynomial in $1/\gamma$ and in the running time of the prediction strategy; without such a lower bound, they also give a uniform expected-time variant. Nonetheless, for the known unclonable cryptographic applications, our qualitatively weaker theorem suffices.

The two papers nevertheless make different technical contributions. The proof of Coladangelo, Liu, and Xie follows a filtering framework closely related to that of Ananth and Sahai [AS26]. By contrast, our proof decomposes the bipartite state received by Bob and Charlie and analyzes the action of Bob's and Charlie's measurements on these different regions using comparatively elementary ingredients. Thus, although our quantitative bound is weaker, our argument provides a structurally different and relatively direct explanation of the non-local search-to-decision phenomenon. As a result, we can view our work as providing an alternate proof of the existence of unclonable encryption using completely different techniques. We consider this alternative proof strategy as an independent conceptual contribution of the present work.

Concurrent work of Behera, Çakan, and Goyal [BCG26] takes a complementary route. They prove an identical-challenge decisional monogamy theorem for coset states and use computational variants to obtain correlated-challenge single-decryptor encryption and copy-protection schemes. Unlike our generic implication from search security to decision security, their theorem exploits the algebraic structure of coset states to obtain new constructions of a variety of unclonable primitives from indistinguishability obfuscation and other cryptographic assumptions.

The results beyond $\mathbb{F}_2$ in this paper remain complementary to both concurrent works: specifically, Theorem 1.2 is unique to our work and does not appear in concurrent works.

2 Preliminaries

2.1 Basic notation

For a positive integer n, write $[n] := \{1, \ldots, n\}$. If S is a finite set, then $s \xleftarrow{\$} S$ denotes a uniformly random element of S. An expectation whose subscript consists only of random variables, such as $\mathbb{E}_x$ or $\mathbb{E}_{x,r}$, is taken with respect to the uniform and mutually independent choices of those variables unless another distribution is stated explicitly. A function $\mu : \mathbb{N} \to \mathbb{R}_{\geq 0}$ is negligible if, for every positive polynomial p, there is an N_p such that $\mu(n) < 1/p(n)$ for every $n \geq N_p$.

For a linear operator M, its operator norm is

$$\|M\|_\infty := \sup_{\||v\rangle\|=1} \|M|v\rangle\|. \tag{2}$$

It is submultiplicative, and $|\langle u|M|v\rangle| \leq \|M\|_\infty \||u\rangle\| \||v\rangle\|$. For a Hermitian operator, the operator norm is the largest absolute value of an eigenvalue. We write $\mathrm{ran}(M)$ for its range. We use the vector Cauchy–Schwarz inequality and its averaged form

$$|\mathbb{E}_x \langle u_x|v_x\rangle| \leq \left(\mathbb{E}_x \||u_x\rangle\|^2\right)^{1/2} \left(\mathbb{E}_x \||v_x\rangle\|^2\right)^{1/2}$$

without further comment.

2.2 Quantum Information

A register $\mathbf{R}$ is associated with a finite-dimensional Hilbert space, also denoted $\mathbf{R}$. We write $\mathcal{L}(\mathbf{R})$ for the linear operators on $\mathbf{R}$ and $\mathcal{D}(\mathbf{R})$ for the density operators. The identity on $\mathbf{R}$ is $I_\mathbf{R}$; identities on registers unaffected by an operator are often suppressed. An operator M is positive semidefinite, written $M \succeq 0$, if $\langle v|M|v\rangle \geq 0$ for every vector $|v\rangle$. For Hermitian M, N, the Loewner order $M \preceq N$ means $N - M \succeq 0$.

A state $\rho \in \mathcal{D}(\mathbf{R})$ satisfies $\rho \succeq 0$ and $\mathrm{Tr}(\rho) = 1$. A unit vector $|\psi\rangle$ represents the pure state $|\psi\rangle\langle\psi|$. For a bipartite state $\rho_{\mathbf{RS}}$, its reduced state on $\mathbf{R}$ is $\rho_\mathbf{R} = \mathrm{Tr}_\mathbf{S}(\rho_{\mathbf{RS}})$. A purification of $\rho_\mathbf{R}$ is a unit vector $|\psi\rangle_{\mathbf{RE}}$ satisfying $\mathrm{Tr}_\mathbf{E}(|\psi\rangle\langle\psi|) = \rho_\mathbf{R}$.

A quantum channel $\Phi : \mathcal{L}(\mathbf{R}) \to \mathcal{L}(\mathbf{S})$ is a completely positive, trace-preserving linear map. A measurement with outcome set $\mathcal{X}$ is a POVM $\{M^z\}_{z\in\mathcal{X}}$ with $M^z \succeq 0$ and $\sum_z M^z = I$. On state ρ, the Born rule assigns outcome z probability $\mathrm{Tr}(M^z\rho)$. A sub-POVM only requires $\sum_z M^z \preceq I$; it can be completed to a POVM by assigning the positive residual effect $I - \sum_z M^z$ to any outcome. We also use cyclicity of the trace, $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$ whenever the products are defined.

A binary POVM $\{M^0, M^1\}$ is equivalently described by the Hermitian contraction $D := M^0 - M^1$, for which $-I \preceq D \preceq I$ and

$$M^b = \frac{I + (-1)^b D}{2}, \qquad b \in \{0, 1\}. \tag{3}$$

A projection P satisfies $P = P^\dagger = P^2$. When an inverse of a compressed operator PMP appears, it is always taken on $\mathrm{ran}(P)$ and, when convenient, extended by zero on the orthogonal complement. In particular, if $PMP \succeq aP$ for $a > 0$, then this restricted inverse exists and has operator norm at most $1/a$ (equivalently, $(PMP)^{-1} \preceq a^{-1}P$ on $\mathrm{ran}(P)$).

2.3 Characters

We identify $\{0,1\}^n$ with the vector space $\mathbb{F}_2^n$. Addition and the inner product

$$\langle r, x\rangle := \sum_{j=1}^{n} r_j x_j \pmod 2 \tag{4}$$

are over $\mathbb{F}_2$. For $x, r \in \mathbb{F}_2^n$, define the character

$$\chi_x(r) := (-1)^{\langle r,x\rangle}. \tag{5}$$

The identities

$$\chi_x(r+r') = \chi_x(r)\chi_x(r'), \qquad \sum_{x\in\mathbb{F}_2^n} \chi_x(s) = \begin{cases} 2^n, & s = 0, \\ 0, & s \neq 0, \end{cases} \tag{6}$$

are the character orthogonality relations used in the Parseval calculation.

3 Non-Local Search-to-Decision over $\mathbb{F}_2$

We prove the main theorem.

Theorem 3.1. *Let n be a positive integer, and let*

$$\boldsymbol{\rho} := (\rho_x)_{x\in\mathbb{F}_2^n}, \qquad \rho_x \in \mathcal{D}(\mathbf{B}\otimes\mathbf{C}),$$

be a bipartite ensemble indexed by a uniformly distributed $x \in \mathbb{F}_2^n$. Define its optimal joint local recovery probability by

$$p_{\mathsf{srch}}(\boldsymbol{\rho}) := \sup_{\Lambda_{\mathcal{B}}^{\mathsf{srch}}, \Lambda_{\mathcal{C}}^{\mathsf{srch}}} \mathbb{E}_x \left[\mathrm{Tr}\left((\mathcal{E}_{\mathcal{B}}^x \otimes \mathcal{E}_{\mathcal{C}}^x)\rho_x\right)\right]. \tag{7}$$

Here the supremum is over local full-string POVMs $\Lambda_{\mathcal{B}}^{\mathsf{srch}} = \{\mathcal{E}_{\mathcal{B}}^z\}_{z\in\mathbb{F}_2^n}$ on $\mathbf{B}$ and $\Lambda_{\mathcal{C}}^{\mathsf{srch}} = \{\mathcal{E}_{\mathcal{C}}^z\}_{z\in\mathbb{F}_2^n}$ on $\mathbf{C}$. The parties can always output the same fixed string, and hence

$$2^{-n} \le p_{\mathsf{srch}}(\boldsymbol{\rho}) \le 1.$$

For every collection of binary-outcome POVMs

$$\Lambda_{\mathcal{B}}^{\mathsf{pred}}[r] = \{\mathcal{P}_{\mathcal{B}}^{r,0}, \mathcal{P}_{\mathcal{B}}^{r,1}\}, \qquad \Lambda_{\mathcal{C}}^{\mathsf{pred}}[r] = \{\mathcal{P}_{\mathcal{C}}^{r,0}, \mathcal{P}_{\mathcal{C}}^{r,1}\},$$

one for each $r \in \mathbb{F}_2^n$, define

$$p_{\mathsf{pred}} := \mathbb{E}_{r,x}\left[\mathrm{Tr}\left((\mathcal{P}_{\mathcal{B}}^{r,\langle r,x\rangle} \otimes \mathcal{P}_{\mathcal{C}}^{r,\langle r,x\rangle})\rho_x\right)\right], \tag{8}$$

where x and r are sampled independently and uniformly from $\mathbb{F}_2^n$. Then

$$p_{\mathsf{pred}} \le \min\left\{1, \frac{1}{2} + 5p_{\mathsf{srch}}(\boldsymbol{\rho})^{1/22}\right\}. \tag{9}$$

Consequently, let $\{\boldsymbol{\rho}^{(n)} = (\rho_x^{(n)})_{x\in\mathbb{F}_2^n}\}_{n\in\mathbb{N}}$ be any family of such ensembles, where the local register dimensions may depend on n. If $p_{\mathsf{srch}}(\boldsymbol{\rho}^{(n)})$ is negligible in n, then, for every family of local prediction POVMs, there is a negligible function ν such that

$$p_{\mathsf{pred}}^{(n)} \le \frac{1}{2} + \nu(n).$$

Proof. For brevity, write $p_{\mathsf{srch}} := p_{\mathsf{srch}}(\boldsymbol{\rho})$. Define

$$\Delta_{\mathcal{B}}^{r} := \mathcal{P}_{\mathcal{B}}^{r,0} - \mathcal{P}_{\mathcal{B}}^{r,1}, \qquad \Delta_{\mathcal{C}}^{r} := \mathcal{P}_{\mathcal{C}}^{r,0} - \mathcal{P}_{\mathcal{C}}^{r,1}. \tag{10}$$

These are Hermitian contractions, and

$$\mathcal{P}_{\mathcal{B}}^{r,b} = \frac{I_{\mathrm{B}} + (-1)^b \Delta_{\mathcal{B}}^{r}}{2}, \qquad \mathcal{P}_{\mathcal{C}}^{r,b} = \frac{I_{\mathrm{C}} + (-1)^b \Delta_{\mathcal{C}}^{r}}{2}.$$

Let

$$\chi_x(r) := (-1)^{\langle r,x \rangle},$$

and define

$$\mathbf{X}^x := \mathbb{E}_r[\chi_x(r)\Delta_{\mathcal{B}}^{r}], \qquad \mathbf{Y}^x := \mathbb{E}_r[\chi_x(r)\Delta_{\mathcal{C}}^{r}], \qquad \mathbf{Z} := \mathbb{E}_r[\Delta_{\mathcal{B}}^{r} \otimes \Delta_{\mathcal{C}}^{r}]. \tag{11}$$

For every x, define the effect corresponding to the event that both parties answer the parity correctly:

$$\begin{aligned} \mathbf{G}_x &:= \mathbb{E}_r \left[\mathcal{P}_{\mathcal{B}}^{r,\langle r,x \rangle} \otimes \mathcal{P}_{\mathcal{C}}^{r,\langle r,x \rangle} \right] \\ &= \frac{1}{4} \left(I_{\mathrm{BC}} + \mathbf{X}^x \otimes I_{\mathrm{C}} + I_{\mathrm{B}} \otimes \mathbf{Y}^x + \mathbf{Z} \right). \end{aligned} \tag{12}$$

Hence

$$p_{\mathsf{pred}} = \mathbb{E}_x[\mathrm{Tr}(\mathbf{G}_x \rho_x)]. \tag{13}$$

Since joint correctness implies Bob's correctness and also Charlie's correctness,

$$0 \preceq \mathbf{G}_x \preceq \frac{I_{\mathrm{BC}} + \mathbf{X}^x \otimes I_{\mathrm{C}}}{2}, \qquad 0 \preceq \mathbf{G}_x \preceq \frac{I_{\mathrm{BC}} + I_{\mathrm{B}} \otimes \mathbf{Y}^x}{2}. \tag{14}$$

In particular, $0 \preceq \mathbf{G}_x \preceq I_{\mathrm{BC}}$.

Fix

$$0 < \eta \leq \frac{1}{4}, \qquad \tau := \eta^3, \qquad \delta := \frac{\eta^2}{3072}, \tag{15}$$

and temporarily assume

$$p_{\mathsf{srch}} \leq \frac{\eta^{22}}{3072^2}. \tag{16}$$

We will prove the conditional bound

$$p_{\mathsf{pred}} < \frac{1}{2} + 1.064\eta. \tag{17}$$

The theorem will then follow by choosing η as a function of p_{srch}.

3.1 Start from the desired expectation and decompose the state

Choose a purification $|\psi^x\rangle_{\mathrm{BCE}}$ of each ρ_x. Every operator on $\mathbf{BC}$ below is understood to act as the identity on $\mathbf{E}$. Equation (13) becomes

$$p_{\mathsf{pred}} = \mathbb{E}_x[\langle \psi^x | \mathbf{G}_x | \psi^x \rangle]. \tag{18}$$

Our entire aim is to upper-bound the expectation on the right by $1/2$ plus a negligible quantity.

Consider spectral decompositions

$$\mathbf{X}^x = \sum_i \lambda_i \, |u_i\rangle\langle u_i| \,, \qquad \mathbf{Y}^x = \sum_i \mu_i \, |v_i\rangle\langle v_i| \,.$$

Define Bob's high and low projections and Charlie's high and low projections by

$$\Pi_x := \sum_{i:\,|\lambda_i|\geq\tau} |u_i\rangle\langle u_i| \,, \qquad N_x := I_{\mathbf{B}} - \Pi_x, \qquad \Theta_x := \sum_{i:\,|\mu_i|\geq\tau} |v_i\rangle\langle v_i| \,, \qquad O_x := I_{\mathbf{C}} - \Theta_x. \tag{19}$$

Thus $N_x \mathbf{X}^x N_x$ has operator norm at most τ, while $O_x \mathbf{Y}^x O_x$ has operator norm at most τ.

For each fixed x, define four mutually orthogonal projections

$$P^x_{\mathrm{HH}} := \Pi_x \otimes \Theta_x \otimes I_{\mathbf{E}}, \qquad P^x_{\mathrm{HL}} := \Pi_x \otimes O_x \otimes I_{\mathbf{E}}, \tag{20}$$

$$P^x_{\mathrm{LH}} := N_x \otimes \Theta_x \otimes I_{\mathbf{E}}, \qquad P^x_{\mathrm{LL}} := N_x \otimes O_x \otimes I_{\mathbf{E}}. \tag{21}$$

The first letter refers to Bob and the second to Charlie. For example, HL means that Bob is in his high subspace and Charlie is in his low subspace. These projections sum to the identity. Define

$$|\psi^x_U\rangle := P^x_U |\psi^x\rangle, \qquad U \in \{\mathrm{HH}, \mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}. \tag{22}$$

Then

$$|\psi^x\rangle = |\psi^x_{\mathrm{HH}}\rangle + |\psi^x_{\mathrm{HL}}\rangle + |\psi^x_{\mathrm{LH}}\rangle + |\psi^x_{\mathrm{LL}}\rangle, \tag{23}$$

and the four vectors on the right are mutually orthogonal. For later use, put

$$\mathrm{wt}(U) := \mathbb{E}_x \| |\psi^x_U\rangle \|^2. \tag{24}$$

Their weights sum to one.

We now open the target quantity in (18) using (23):

$$\begin{aligned} p_{\mathsf{pred}} = & \sum_{U \in \{\mathrm{HH},\mathrm{HL},\mathrm{LH},\mathrm{LL}\}} \mathbb{E}_x \langle \psi^x_U | \mathbf{G}_x | \psi^x_U \rangle \\ & + 2 \, \mathrm{Re} \sum_{U<V} \mathbb{E}_x \langle \psi^x_U | \mathbf{G}_x | \psi^x_V \rangle. \end{aligned} \tag{25}$$

Here $U < V$ denotes any fixed ordering of the four labels. The first line contains the four diagonal terms. The second line contains the six cross terms.

This expansion reveals the proof strategy.

1. The diagonal terms supported on HL, LH, and LL are easy: on each of these subspaces at least one party is low, so the marginal bounds in (14) put the value below $(1+\tau)/2$ times the squared norm.

2. The HH diagonal term does not have this easy bound. Instead, we will show that the average squared norm of $|\psi^x_{\mathrm{HH}}\rangle$ is small (by reducing to search security). The same small-norm estimate will control all cross terms involving the HH component at once.

3. After the HH component and any other component of tiny average weight are removed from the estimate, the only potentially difficult interaction is among the HL, LH, and LL components. The cross terms involving LL should not be estimated one by one: their signs are uncontrolled. We instead study

$$\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x$$

and minimize over the LL component. This combines the HL–LL and LH–LL cross terms with the LL diagonal term and leaves one effective HL–LH cross scalar c_x.

4. Finally, we expand the average of c_x into contributions associated with fixed sequences of common challenges. Each such contribution yields a valid pair of local full-string extraction measurements and is therefore bounded by p_{srch}.

The rest of the proof implements these four points in this order.

3.2 The high–high component and all terms involving omitted vectors

The first lemma supplies the measurement needed to control the HH component.

Lemma 3.2 (Parseval and the average weight of the HH component)**.** *We have*

$$\sum_x \Pi_x \preceq \tau^{-2} I_{\mathrm{B}}, \qquad \sum_x \Theta_x \preceq \tau^{-2} I_{\mathrm{C}}, \tag{26}$$

and consequently

$$\mathrm{wt}(\mathrm{HH}) \le \frac{p_{\mathsf{srch}}}{\tau^4}. \tag{27}$$

Proof. Character orthogonality gives

$$\begin{aligned}\sum_x (\mathbf{X}^x)^2 &= \sum_x \mathbb{E}_{r,r'} \chi_x(r+r') \Delta_{\mathcal{B}}^r \Delta_{\mathcal{B}}^{r'} \\ &= \mathbb{E}_r (\Delta_{\mathcal{B}}^r)^2 \preceq I_{\mathrm{B}}. \end{aligned} \tag{28}$$

Likewise,

$$\sum_x (\mathbf{Y}^x)^2 \preceq I_{\mathrm{C}}. \tag{29}$$

On the range of Π_x, the positive operator $(\mathbf{X}^x)^2$ is at least $\tau^2 I$. Hence $\Pi_x \preceq \tau^{-2}(\mathbf{X}^x)^2$. Summing over x proves the first inequality in (26); the second is identical.

It follows that $\{\tau^2 \Pi_x\}_x$ and $\{\tau^2 \Theta_x\}_x$ are local sub-POVMs. Assigning each residual effect to an arbitrary fixed output string completes them to POVMs. Their joint probability of outputting the correct full string is at least

$$\begin{aligned}\mathbb{E}_x \langle \psi^x | (\tau^2 \Pi_x \otimes \tau^2 \Theta_x \otimes I_{\mathrm{E}}) | \psi^x \rangle &= \tau^4 \mathbb{E}_x \| |\psi^x_{\mathrm{HH}}\rangle \|^2 \\ &= \tau^4 \,\mathrm{wt}(\mathrm{HH}).\end{aligned}$$

By the definition of p_{srch}, this is at most p_{srch}, proving (27). □

Rather than bounding separately the HH diagonal term and the three cross terms involving HH in (25), we remove the entire HH vector from the mathematical estimate. We also remove any of the other three indexed component families whose average weight is below δ. This second removal will later ensure that whenever both HL and LH remain, their average weights are bounded away from zero.

For $U \in \{\mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}$, define

$$\kappa_U := \begin{cases} 1, & \mathrm{wt}(U) \geq \delta, \\ 0, & \mathrm{wt}(U) < \delta. \end{cases} \tag{30}$$

Set

$$|\phi_U^x\rangle := \kappa_U |\psi_U^x\rangle, \qquad |\phi^x\rangle := |\phi_{\mathrm{HL}}^x\rangle + |\phi_{\mathrm{LH}}^x\rangle + |\phi_{\mathrm{LL}}^x\rangle, \tag{31}$$

and define

$$|e^x\rangle := |\psi^x\rangle - |\phi^x\rangle, \qquad q_{\mathrm{omit}} := \mathbb{E}_x \||e^x\rangle\|^2. \tag{32}$$

No party performs this removal. It is only a decomposition used in the estimate.

Claim 3.3 (Cost of all omitted diagonal and cross terms)**.** *We have*

$$q_{\mathrm{omit}} \leq \frac{p_{\mathsf{srch}}}{\tau^4} + 3\delta, \tag{33}$$

and

$$p_{\mathsf{pred}} \leq \mathbb{E}_x \langle \phi^x | \mathbf{G}_x | \phi^x \rangle + 2\sqrt{q_{\mathrm{omit}}} + q_{\mathrm{omit}}. \tag{34}$$

Proof. The four original components are orthogonal. The omitted HH component has average weight at most p_{srch}/τ^4 by (27); each additional omitted component has average weight below δ. This proves (33).

Using $|\psi^x\rangle = |\phi^x\rangle + |e^x\rangle$,

$$p_{\mathsf{pred}} = \mathbb{E}_x \langle \phi^x | \mathbf{G}_x | \phi^x \rangle + 2\,\mathrm{Re}\,\mathbb{E}_x \langle \phi^x | \mathbf{G}_x | e^x \rangle + \mathbb{E}_x \langle e^x | \mathbf{G}_x | e^x \rangle.$$

The last term is at most q_{omit} because $0 \preceq \mathbf{G}_x \preceq I$. Moreover,

$$\begin{aligned} |\mathbb{E}_x \langle \phi^x | \mathbf{G}_x | e^x \rangle| &= \left| \mathbb{E}_x \langle \mathbf{G}_x^{1/2} \phi^x \mid \mathbf{G}_x^{1/2} e^x \rangle \right| \\ &\leq \left(\mathbb{E}_x \langle \phi^x | \mathbf{G}_x | \phi^x \rangle \right)^{1/2} \left(\mathbb{E}_x \langle e^x | \mathbf{G}_x | e^x \rangle \right)^{1/2} \\ &\leq \sqrt{q_{\mathrm{omit}}}. \end{aligned}$$

This proves (34). □

The claim has now accounted for the HH diagonal term, every cross term involving HH, and every term involving any other omitted component. It remains to show

$$\mathbb{E}_x \langle \phi^x | \mathbf{G}_x | \phi^x \rangle \leq \frac{1}{2} + \eta. \tag{35}$$

3.3 Diagonal and cross terms of the retained vector

Define the retained average weights

$$W_U := \mathbb{E}_x \| |\phi_U^x\rangle \|^2, \qquad U \in \{\mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}. \tag{36}$$

Each W_U is either zero or at least δ.

We first dispose of the case in which one of the two one-high components is absent. If $W_{\mathrm{HL}} = 0$, then $|\phi_{\mathrm{HL}}^x\rangle = 0$ for every x, and every retained vector lies in Bob's low subspace. The first inequality in (14) gives

$$\langle \phi^x | \mathbf{G}_x | \phi^x \rangle \le \frac{1+\tau}{2} \| |\phi^x\rangle \|^2 \le \left(\frac{1}{2} + \eta\right) \| |\phi^x\rangle \|^2. \tag{37}$$

If $W_{\mathrm{LH}} = 0$, then every retained vector lies in Charlie's low subspace, and the second inequality in (14) gives the same result. Thus (35) already holds whenever

$$W_{\mathrm{HL}} = 0 \quad \text{or} \quad W_{\mathrm{LH}} = 0. \tag{38}$$

We henceforth assume

$$W_{\mathrm{HL}} \ge \delta, \qquad W_{\mathrm{LH}} \ge \delta, \qquad \sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}} \ge \delta. \tag{39}$$

In this case $|\phi_{\mathrm{HL}}^x\rangle = |\psi_{\mathrm{HL}}^x\rangle$ and $|\phi_{\mathrm{LH}}^x\rangle = |\psi_{\mathrm{LH}}^x\rangle$ for every x.

The remaining diagonal and cross terms can now be seen directly by writing

$$\begin{aligned}
&\langle \phi^x | \left(\left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x \right) |\phi^x\rangle \\
&= \sum_{U \in \{\mathrm{HL},\mathrm{LH},\mathrm{LL}\}} \langle \phi_U^x | \left(\left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x \right) |\phi_U^x\rangle \\
&\quad - 2\,\mathrm{Re}\langle \phi_{\mathrm{HL}}^x | \mathbf{G}_x | \phi_{\mathrm{LH}}^x \rangle - 2\,\mathrm{Re}\langle \phi_{\mathrm{HL}}^x | \mathbf{G}_x | \phi_{\mathrm{LL}}^x \rangle - 2\,\mathrm{Re}\langle \phi_{\mathrm{LH}}^x | \mathbf{G}_x | \phi_{\mathrm{LL}}^x \rangle.
\end{aligned} \tag{40}$$

We work with the operator $\left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x$ because it is sufficient for (35) to prove that the average of (40) is nonnegative. Indeed, that nonnegativity gives

$$\mathbb{E}_x \langle \phi^x | \mathbf{G}_x | \phi^x \rangle \le \left(\frac{1}{2} + \eta\right) \mathbb{E}_x \| |\phi^x\rangle \|^2 \le \frac{1}{2} + \eta.$$

More importantly, on any subspace where Bob or Charlie is low, the quadratic form of $\left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x$ has a positive lower bound. This lets us combine the LL diagonal term with both cross terms involving LL by completing the square.

For each x, define the compressed LL-block operator

$$\boxed{M_x := P_{\mathrm{LL}}^x \left(\left(\frac{1}{2} + \eta\right) I - \mathbf{G}_x \right) P_{\mathrm{LL}}^x \Big|_{\mathrm{ran}(P_{\mathrm{LL}}^x)}.} \tag{41}$$

Here M_x is regarded as an operator from $\mathrm{ran}(P_{\mathrm{LL}}^x)$ to itself. Equation (56) establishes that $M_x \succeq (\eta - \eta^3/2) I$ on this subspace, so M_x is invertible there and M_x^{-1} denotes this inverse. By contrast, the sandwiched operator in (41), if regarded as an operator on the entire ambient space, vanishes

on $\ker(P^x_{\mathrm{LL}})$ and is therefore generally not invertible there. If $\mathrm{ran}(P^x_{\mathrm{LL}})$ is zero-dimensional, every term containing M_x^{-1} is defined to be zero.

For each x, define the scalar

$$\begin{aligned} c_x := \langle \phi^x_{\mathrm{HL}} | \mathbf{G}_x | \phi^x_{\mathrm{LH}} \rangle \\ + \langle \phi^x_{\mathrm{HL}} | \mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x | \phi^x_{\mathrm{LH}} \rangle. \end{aligned} \tag{42}$$

The first term in c_x is the direct HL–LH cross term in (40). The second is the additional HL–LH cross term created when the two cross terms involving LL are minimized together with the LL diagonal term.

Claim 3.4 (The retained expectation reduces to the average of c_x)**.**

$$\begin{aligned} \mathbb{E}_x \langle \phi^x | \left(\left(\frac{1}{2} + \eta \right) I - \mathbf{G}_x \right) | \phi^x \rangle \\ \geq \left(\eta - \frac{\eta^3}{2} \right) (W_{\mathrm{HL}} + W_{\mathrm{LH}}) - 2 \left| \mathbb{E}_x c_x \right|. \end{aligned} \tag{43}$$

Proof. Fix x. If $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{HL}} + P^x_{\mathrm{LL}})$, then the definitions of the two projections give

$$\begin{aligned} P^x_{\mathrm{HL}} + P^x_{\mathrm{LL}} &= \Pi_x \otimes O_x \otimes I_{\mathbf{E}} + N_x \otimes O_x \otimes I_{\mathbf{E}} \\ &= (\Pi_x + N_x) \otimes O_x \otimes I_{\mathbf{E}} \\ &= I_{\mathbf{B}} \otimes O_x \otimes I_{\mathbf{E}}. \end{aligned} \tag{44}$$

Consequently,

$$(I_{\mathbf{B}} \otimes O_x \otimes I_{\mathbf{E}}) |z\rangle = |z\rangle. \tag{45}$$

Since O_x is the spectral projection onto the eigenvalues of $\mathbf{Y}^x$ whose absolute values are smaller than τ,

$$\begin{aligned} &\langle z | (I_{\mathbf{B}} \otimes \mathbf{Y}^x \otimes I_{\mathbf{E}}) | z \rangle \\ &\quad = \langle z | (I_{\mathbf{B}} \otimes O_x \mathbf{Y}^x O_x \otimes I_{\mathbf{E}}) | z \rangle \\ &\quad \leq \tau \| |z\rangle \|^2. \end{aligned} \tag{46}$$

The second marginal inequality in (14) now gives

$$\begin{aligned} \langle z | \mathbf{G}_x | z \rangle &\leq \frac{1}{2} \langle z | \left(I_{\mathbf{BC}} + I_{\mathbf{B}} \otimes \mathbf{Y}^x \right) | z \rangle \\ &= \frac{1}{2} \| |z\rangle \|^2 + \frac{1}{2} \langle z | \left(I_{\mathbf{B}} \otimes \mathbf{Y}^x \right) | z \rangle \\ &\leq \frac{1+\tau}{2} \| |z\rangle \|^2 = \frac{1+\eta^3}{2} \| |z\rangle \|^2, \end{aligned} \tag{47}$$

and therefore

$$\begin{aligned}
&\langle z| \left(\left(\frac{1}{2} + \eta \right) I - \mathbf{G}_x \right) |z\rangle \\
&\quad = \left(\frac{1}{2} + \eta \right) \| |z\rangle \|^2 - \langle z|\mathbf{G}_x|z\rangle \\
&\quad \geq \left(\frac{1}{2} + \eta - \frac{1+\eta^3}{2} \right) \| |z\rangle \|^2 \\
&\quad = \left(\eta - \frac{\eta^3}{2} \right) \| |z\rangle \|^2. \qquad (48)
\end{aligned}$$

If $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LH}} + P^x_{\mathrm{LL}})$, then similarly

$$\begin{aligned}
P^x_{\mathrm{LH}} + P^x_{\mathrm{LL}} &= N_x \otimes \Theta_x \otimes I_{\mathbf{E}} + N_x \otimes O_x \otimes I_{\mathbf{E}} \\
&= N_x \otimes (\Theta_x + O_x) \otimes I_{\mathbf{E}} \\
&= N_x \otimes I_{\mathbf{C}} \otimes I_{\mathbf{E}}. \qquad (49)
\end{aligned}$$

Thus

$$(N_x \otimes I_{\mathbf{C}} \otimes I_{\mathbf{E}})|z\rangle = |z\rangle. \qquad (50)$$

Since N_x is the spectral projection onto the eigenvalues of $\mathbf{X}^x$ whose absolute values are smaller than τ,

$$\begin{aligned}
&\langle z|(\mathbf{X}^x \otimes I_{\mathbf{C}} \otimes I_{\mathbf{E}})|z\rangle \\
&\quad = \langle z|(N_x \mathbf{X}^x N_x \otimes I_{\mathbf{C}} \otimes I_{\mathbf{E}})|z\rangle \\
&\quad \leq \tau \| |z\rangle \|^2. \qquad (51)
\end{aligned}$$

The first marginal inequality in (14) therefore gives

$$\begin{aligned}
\langle z|\mathbf{G}_x|z\rangle &\leq \frac{1}{2} \langle z| \left(I_{\mathbf{BC}} + \mathbf{X}^x \otimes I_{\mathbf{C}} \right) |z\rangle \\
&= \frac{1}{2} \| |z\rangle \|^2 + \frac{1}{2} \langle z| \left(\mathbf{X}^x \otimes I_{\mathbf{C}} \right) |z\rangle \\
&\leq \frac{1+\tau}{2} \| |z\rangle \|^2 = \frac{1+\eta^3}{2} \| |z\rangle \|^2, \qquad (52)
\end{aligned}$$

and therefore

$$\begin{aligned}
&\langle z| \left(\left(\frac{1}{2} + \eta \right) I - \mathbf{G}_x \right) |z\rangle \\
&\quad = \left(\frac{1}{2} + \eta \right) \| |z\rangle \|^2 - \langle z|\mathbf{G}_x|z\rangle \\
&\quad \geq \left(\frac{1}{2} + \eta - \frac{1+\eta^3}{2} \right) \| |z\rangle \|^2 \\
&\quad = \left(\eta - \frac{\eta^3}{2} \right) \| |z\rangle \|^2. \qquad (53)
\end{aligned}$$

In particular, for every nonzero $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$,

$$\begin{aligned}
&\langle z|P^x_{\mathrm{LL}}\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)P^x_{\mathrm{LL}}|z\rangle \\
&\quad=\langle z|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|z\rangle \\
&\quad\geq\left(\eta-\frac{\eta^3}{2}\right)\||z\rangle\|^2>0.
\end{aligned} \tag{54}$$

The inequality before the final strict inequality also holds for $|z\rangle = 0$. Hence, for every vector $|v\rangle$ in the full space, we may apply it to $|z\rangle = P^x_{\mathrm{LL}}|v\rangle$ and obtain

$$\begin{aligned}
&\langle v|P^x_{\mathrm{LL}}\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)P^x_{\mathrm{LL}}|v\rangle \\
&\quad\geq\left(\eta-\frac{\eta^3}{2}\right)\langle v|P^x_{\mathrm{LL}}|v\rangle.
\end{aligned} \tag{55}$$

By the definition of positive-semidefinite order through quadratic forms, this is precisely the operator inequality

$$P^x_{\mathrm{LL}}\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)P^x_{\mathrm{LL}}\succeq\left(\eta-\frac{\eta^3}{2}\right)P^x_{\mathrm{LL}}. \tag{56}$$

Because $\eta-\eta^3/2=\eta(1-\eta^2/2)>0$, the restriction of the left-hand side to $\mathrm{ran}(P^x_{\mathrm{LL}})$ is strictly positive and therefore invertible there. This proves that the inverse in (42) exists on that subspace.

Using $|\phi^x\rangle = |\phi^x_{\mathrm{HL}}\rangle + |\phi^x_{\mathrm{LH}}\rangle + |\phi^x_{\mathrm{LL}}\rangle$ and the pairwise orthogonality of these three vectors,

$$\begin{aligned}
&\langle\phi^x|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|\phi^x\rangle \\
&=\langle\phi^x_{\mathrm{HL}}|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|\phi^x_{\mathrm{HL}}\rangle \\
&\quad+\langle\phi^x_{\mathrm{LH}}|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|\phi^x_{\mathrm{LH}}\rangle \\
&\quad+\langle\phi^x_{\mathrm{LL}}|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|\phi^x_{\mathrm{LL}}\rangle \\
&\quad-2\,\mathrm{Re}\langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle \\
&\quad-2\,\mathrm{Re}\langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x|\phi^x_{\mathrm{LL}}\rangle \\
&\quad-2\,\mathrm{Re}\langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x|\phi^x_{\mathrm{LL}}\rangle.
\end{aligned} \tag{57}$$

We now combine the LL diagonal term with the two cross terms involving $|\phi^x_{\mathrm{LL}}\rangle$, namely

$$\begin{aligned}
&\langle\phi^x_{\mathrm{LL}}|\left(\left(\frac{1}{2}+\eta\right)I-\mathbf{G}_x\right)|\phi^x_{\mathrm{LL}}\rangle \\
&\quad-2\,\mathrm{Re}\langle\phi^x_{\mathrm{LL}}|\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle \\
&\quad-2\,\mathrm{Re}\langle\phi^x_{\mathrm{LL}}|\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle.
\end{aligned} \tag{58}$$

Here we used the Hermiticity of $\mathbf{G}_x$ to reverse the two cross scalars inside their real parts.

Recall from (41) that M_x is the compression of $\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x$ to $\mathrm{ran}(P^x_{\mathrm{LL}})$ and hence is the coefficient of the quadratic term in $|\phi^x_{\mathrm{LL}}\rangle$. By (56), M_x is positive definite, so M_x^{-1} exists and is Hermitian on this subspace.

For this same calculation, also define

$$|u_x\rangle := P^x_{\mathrm{LL}} \mathbf{G}_x \left(|\phi^x_{\mathrm{HL}}\rangle + |\phi^x_{\mathrm{LH}}\rangle\right). \tag{59}$$

Taking the adjoint reverses the operator order. In particular,

$$(P^x_{\mathrm{LL}} \mathbf{G}_x)^\dagger = \mathbf{G}_x P^x_{\mathrm{LL}}, \tag{60}$$

and hence

$$\langle u_x| = \left(\langle \phi^x_{\mathrm{HL}}| + \langle \phi^x_{\mathrm{LH}}|\right) \mathbf{G}_x P^x_{\mathrm{LL}}. \tag{61}$$

No commutation between P^x_{LL} and $\mathbf{G}_x$ is being used.

Because $P^x_{\mathrm{LL}}|\phi^x_{\mathrm{LL}}\rangle = |\phi^x_{\mathrm{LL}}\rangle$, the three terms in (58) first satisfy the intermediate equality

$$\begin{aligned} &\langle \phi^x_{\mathrm{LL}}| \left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x \right) |\phi^x_{\mathrm{LL}}\rangle \\ &\quad - 2\,\mathrm{Re}\langle \phi^x_{\mathrm{LL}}| \mathbf{G}_x \left(|\phi^x_{\mathrm{HL}}\rangle + |\phi^x_{\mathrm{LH}}\rangle\right) \\ &= \langle \phi^x_{\mathrm{LL}}| M_x |\phi^x_{\mathrm{LL}}\rangle - 2\,\mathrm{Re}\langle \phi^x_{\mathrm{LL}}||u_x\rangle. \end{aligned} \tag{62}$$

This equality accounts for every inserted copy of P^x_{LL}.

Remark (completing the square and why M^{-1} appears). *The quadratic expression under consideration is precisely the part of* (62) *that depends on the* LL *vector. Hold the* HL *and* LH *components fixed, and let $|\ell\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$ stand for a possible* LL *component. The relevant expression is*

$$q_x(\ell) := \langle \ell| M_x |\ell\rangle - 2\,\mathrm{Re}\langle \ell||u_x\rangle.$$

Its value for the state in the proof is $q_x(\phi^x_{\mathrm{LL}})$. To obtain a lower bound, we may minimize $q_x(\ell)$ over all vectors in the LL *sector, because the actual vector $|\phi^x_{\mathrm{LL}}\rangle$ is one of the allowed choices.*

Here is the general completion-of-the-square calculation. Let M be positive definite and set

$$f(\ell) := \langle \ell| M |\ell\rangle - 2\,\mathrm{Re}\langle \ell||u\rangle.$$

Since

$$2\,\mathrm{Re}\langle \ell||u\rangle = \langle \ell||u\rangle + \langle u||\ell\rangle,$$

we add and subtract $\langle u|M^{-1}|u\rangle$. Expanding the resulting product gives

$$\begin{aligned} &\left(\langle \ell| - \langle u|M^{-1}\right) M \left(|\ell\rangle - M^{-1}|u\rangle\right) \\ &\quad = \langle \ell| M |\ell\rangle - \langle \ell||u\rangle - \langle u||\ell\rangle + \langle u|M^{-1}|u\rangle. \end{aligned}$$

Therefore

$$f(\ell) = \left(\langle \ell| - \langle u|M^{-1}\right) M \left(|\ell\rangle - M^{-1}|u\rangle\right) - \langle u|M^{-1}|u\rangle.$$

This is the completion of the square. The first term is nonnegative and vanishes exactly when $M|\ell\rangle = |u\rangle$, equivalently $|\ell\rangle = M^{-1}|u\rangle$. Thus minimizing the eliminated sector produces the correction

$-\langle u|M^{-1}|u\rangle$*; this is the corresponding Schur-complement correction. In* (63)*,* $M = M_x$*,* $|u\rangle = |u_x\rangle$*, and* $|\ell\rangle = |\phi^x_{\mathrm{LL}}\rangle$*.*

Completing the square now gives the shorter identity

$$\begin{aligned}
&\langle\phi^x_{\mathrm{LL}}|M_x|\phi^x_{\mathrm{LL}}\rangle - 2\,\mathrm{Re}\langle\phi^x_{\mathrm{LL}}||u_x\rangle \\
&= \left(\langle\phi^x_{\mathrm{LL}}| - \langle u_x|M_x^{-1}\right) M_x \left(|\phi^x_{\mathrm{LL}}\rangle - M_x^{-1}|u_x\rangle\right) - \langle u_x|M_x^{-1}|u_x\rangle.
\end{aligned} \tag{63}$$

The first term on the right is nonnegative. Therefore the three terms in (58) are at least

$$-\langle u_x|M_x^{-1}|u_x\rangle. \tag{64}$$

Only after completing the square do we expand this scalar. The definition of $|u_x\rangle$ and the adjoint formula (61) give

$$\begin{aligned}
&\langle u_x|M_x^{-1}|u_x\rangle \\
&= \langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle \\
&\quad + \langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle \\
&\quad + 2\,\mathrm{Re}\langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle.
\end{aligned} \tag{65}$$

The left copy of $\mathbf{G}_x$ in every scalar comes from $\langle u_x|$, and the right copy comes from $|u_x\rangle$; both copies are required. By the definition of M_x, the operator $P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}}$ is exactly the inverse in (42), embedded in the full space by the surrounding projections. We now combine (57), (64), (65), and the definition of c_x:

$$\begin{aligned}
&\langle\phi^x|\left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right)|\phi^x\rangle \\
&\geq \Bigg[\langle\phi^x_{\mathrm{HL}}|\left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right)|\phi^x_{\mathrm{HL}}\rangle \\
&\qquad\qquad - \langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle\Bigg] \\
&\quad + \Bigg[\langle\phi^x_{\mathrm{LH}}|\left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right)|\phi^x_{\mathrm{LH}}\rangle \\
&\qquad\qquad - \langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle\Bigg] \\
&\quad - 2\,\mathrm{Re}\, c_x.
\end{aligned} \tag{66}$$

For every $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$,

$$\begin{aligned}
&(\langle\phi^x_{\mathrm{HL}}| + \langle z|)\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)(|\phi^x_{\mathrm{HL}}\rangle + |z\rangle)\\
&= \langle\phi^x_{\mathrm{HL}}|\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)|\phi^x_{\mathrm{HL}}\rangle\\
&\quad + \langle z|M_x|z\rangle\\
&\quad - 2\,\mathrm{Re}\langle z|P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle\\
&= \left(\langle z| - \langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1}\right) M_x \left(|z\rangle - M_x^{-1}P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle\right)\\
&\quad + \langle\phi^x_{\mathrm{HL}}|\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)|\phi^x_{\mathrm{HL}}\rangle\\
&\quad - \langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle. \qquad (67)
\end{aligned}$$

Taking the minimum over $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$ in (67) yields

$$\begin{aligned}
&\langle\phi^x_{\mathrm{HL}}|\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)|\phi^x_{\mathrm{HL}}\rangle\\
&\quad - \langle\phi^x_{\mathrm{HL}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle\\
&= \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} (\langle\phi^x_{\mathrm{HL}}| + \langle z|)\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)(|\phi^x_{\mathrm{HL}}\rangle + |z\rangle)\\
&\geq \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} \left(\eta - \frac{\eta^3}{2}\right)\||\phi^x_{\mathrm{HL}}\rangle + |z\rangle\|^2\\
&= \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} \left(\eta - \frac{\eta^3}{2}\right)\left(\||\phi^x_{\mathrm{HL}}\rangle\|^2 + \||z\rangle\|^2\right)\\
&= \left(\eta - \frac{\eta^3}{2}\right)\||\phi^x_{\mathrm{HL}}\rangle\|^2. \qquad (68)
\end{aligned}$$

Likewise, for every $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$,

$$\begin{aligned}
&(\langle\phi^x_{\mathrm{LH}}| + \langle z|)\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)(|\phi^x_{\mathrm{LH}}\rangle + |z\rangle)\\
&= \langle\phi^x_{\mathrm{LH}}|\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)|\phi^x_{\mathrm{LH}}\rangle\\
&\quad + \langle z|M_x|z\rangle\\
&\quad - 2\,\mathrm{Re}\langle z|P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle\\
&= \left(\langle z| - \langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1}\right) M_x \left(|z\rangle - M_x^{-1}P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle\right)\\
&\quad + \langle\phi^x_{\mathrm{LH}}|\left(\left(\frac{1}{2}+\eta\right)I - \mathbf{G}_x\right)|\phi^x_{\mathrm{LH}}\rangle\\
&\quad - \langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle. \qquad (69)
\end{aligned}$$

Taking the minimum over $|z\rangle \in \mathrm{ran}(P^x_{\mathrm{LL}})$ in (69) yields

$$\begin{aligned}
&\langle\phi^x_{\mathrm{LH}}| \left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right) |\phi^x_{\mathrm{LH}}\rangle \\
&\quad - \langle\phi^x_{\mathrm{LH}}|\mathbf{G}_x P^x_{\mathrm{LL}} M_x^{-1} P^x_{\mathrm{LL}} \mathbf{G}_x |\phi^x_{\mathrm{LH}}\rangle \\
&= \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} (\langle\phi^x_{\mathrm{LH}}| + \langle z|) \left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right) (|\phi^x_{\mathrm{LH}}\rangle + |z\rangle) \\
&\geq \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} \left(\eta - \frac{\eta^3}{2}\right) \||\phi^x_{\mathrm{LH}}\rangle + |z\rangle\|^2 \\
&= \min_{|z\rangle\in\mathrm{ran}(P^x_{\mathrm{LL}})} \left(\eta - \frac{\eta^3}{2}\right) \left(\||\phi^x_{\mathrm{LH}}\rangle\|^2 + \||z\rangle\|^2\right) \\
&= \left(\eta - \frac{\eta^3}{2}\right) \||\phi^x_{\mathrm{LH}}\rangle\|^2. \qquad (70)
\end{aligned}$$

Combining (66), (68), and (70),

$$\begin{aligned}
&\langle\phi^x| \left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right) |\phi^x\rangle \\
&\qquad \geq \left(\eta - \frac{\eta^3}{2}\right) \left(\||\phi^x_{\mathrm{HL}}\rangle\|^2 + \||\phi^x_{\mathrm{LH}}\rangle\|^2\right) - 2\,\mathrm{Re}\, c_x. \qquad (71)
\end{aligned}$$

Finally,

$$\begin{aligned}
&\mathbb{E}_x\langle\phi^x| \left(\left(\frac{1}{2}+\eta\right) I - \mathbf{G}_x\right) |\phi^x\rangle \\
&\qquad \geq \left(\eta - \frac{\eta^3}{2}\right) \mathbb{E}_x \left[\||\phi^x_{\mathrm{HL}}\rangle\|^2 + \||\phi^x_{\mathrm{LH}}\rangle\|^2\right] - 2\mathbb{E}_x[\mathrm{Re}\, c_x] \\
&\qquad = \left(\eta - \frac{\eta^3}{2}\right) (W_{\mathrm{HL}} + W_{\mathrm{LH}}) - 2\,\mathrm{Re}\,\mathbb{E}_x[c_x] \\
&\qquad \geq \left(\eta - \frac{\eta^3}{2}\right) (W_{\mathrm{HL}} + W_{\mathrm{LH}}) - 2\,|\mathbb{E}_x c_x|\,. \qquad (72)
\end{aligned}$$

This is (43). □

By the arithmetic–geometric mean inequality, the claim shows that (35) follows once we prove

$$|\mathbb{E}_x c_x| \leq \left(\eta - \frac{\eta^3}{2}\right) \sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}}. \qquad (73)$$

Indeed, $2\sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}} \leq W_{\mathrm{HL}} + W_{\mathrm{LH}}$. Thus the entire proof has now been reduced to one averaged cross scalar.

3.4 Expand the effective cross scalar into common-challenge contributions

The definition of c_x still contains the inverse of the restriction of $((1/2+\eta)I - \mathbf{G}_x)$ to the LL subspace. The next claim rewrites this inverse as a geometric series involving the common-challenge operator $\mathbf{Z}$. This is useful because every copy of

$$\mathbf{Z} = \mathbb{E}_r[\Delta^r_{\mathcal{B}} \otimes \Delta^r_{\mathcal{C}}]$$

uses the same challenge on Bob's and Charlie's sides.

Define

$$\Gamma_1 := \mathbb{E}_x \langle \phi^x_{\mathrm{HL}} | \mathbf{Z} | \phi^x_{\mathrm{LH}} \rangle, \tag{74}$$

and, for $m \geq 2$,

$$\Gamma_m := \mathbb{E}_x \langle \phi^x_{\mathrm{HL}} | \mathbf{Z} P^x_{\mathrm{LL}} (P^x_{\mathrm{LL}} \mathbf{Z} P^x_{\mathrm{LL}})^{m-2} P^x_{\mathrm{LL}} \mathbf{Z} | \phi^x_{\mathrm{LH}} \rangle. \tag{75}$$

Claim 3.5 (Geometric expansion of the average effective cross term)**.**

$$\left| \mathbb{E}_x c_x - \frac{1}{4} \left[\Gamma_1 + \sum_{k=0}^{\infty} (1+4\eta)^{-(k+1)} \Gamma_{k+2} \right] \right| \leq \frac{\eta}{16} \sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}}. \tag{76}$$

Proof. Fix x.

We first isolate the part of $\mathbf{G}_x$ that changes a high/low label. In (12), the identity preserves each of the four high/low subspaces. Moreover, $\mathbf{X}^x$ commutes with its spectral projections Π_x, N_x and acts trivially on Charlie, while $\mathbf{Y}^x$ commutes with Θ_x, O_x and acts trivially on Bob. In particular,

$$\Pi_x N_x = N_x \Pi_x = 0, \qquad O_x \Theta_x = \Theta_x O_x = 0,$$

and the spectral subspaces reduce the corresponding Fourier operators:

$$\Pi_x \mathbf{X}^x N_x = N_x \mathbf{X}^x \Pi_x = 0, \qquad O_x \mathbf{Y}^x \Theta_x = \Theta_x \mathbf{Y}^x O_x = 0.$$

We suppress the identity on $\mathbf{E}$. Using the support identities $P^x_U |\phi^x_U\rangle = |\phi^x_U\rangle$ and the decomposition of $\mathbf{G}_x$ in (12), the HL–LH block expands as

$$\begin{aligned}
\langle \phi^x_{\mathrm{HL}} | \mathbf{G}_x | \phi^x_{\mathrm{LH}} \rangle &= \langle \phi^x_{\mathrm{HL}} | P^x_{\mathrm{HL}} \mathbf{G}_x P^x_{\mathrm{LH}} | \phi^x_{\mathrm{LH}} \rangle \\
&= \frac{1}{4} \langle \phi^x_{\mathrm{HL}} | \Big[P^x_{\mathrm{HL}} P^x_{\mathrm{LH}} + P^x_{\mathrm{HL}} (\mathbf{X}^x \otimes I_{\mathbf{C}}) P^x_{\mathrm{LH}} \\
&\qquad\qquad + P^x_{\mathrm{HL}} (I_{\mathbf{B}} \otimes \mathbf{Y}^x) P^x_{\mathrm{LH}} + P^x_{\mathrm{HL}} \mathbf{Z} P^x_{\mathrm{LH}} \Big] | \phi^x_{\mathrm{LH}} \rangle \\
&= \frac{1}{4} \langle \phi^x_{\mathrm{HL}} | \Big[\underbrace{(\Pi_x N_x) \otimes (O_x \Theta_x)}_{=0} + \underbrace{(\Pi_x \mathbf{X}^x N_x) \otimes (O_x \Theta_x)}_{=0} \\
&\qquad\qquad + \underbrace{(\Pi_x N_x) \otimes (O_x \mathbf{Y}^x \Theta_x)}_{=0} + P^x_{\mathrm{HL}} \mathbf{Z} P^x_{\mathrm{LH}} \Big] | \phi^x_{\mathrm{LH}} \rangle \\
&= \frac{1}{4} \langle \phi^x_{\mathrm{HL}} | \mathbf{Z} | \phi^x_{\mathrm{LH}} \rangle.
\end{aligned} \tag{77}$$

Likewise, the block from HL into LL expands as

$$\begin{aligned} P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{HL}}\rangle &= P^x_{\mathrm{LL}}\mathbf{G}_x P^x_{\mathrm{HL}}|\phi^x_{\mathrm{HL}}\rangle \\ &= \frac{1}{4}\Big[P^x_{\mathrm{LL}}P^x_{\mathrm{HL}} + P^x_{\mathrm{LL}}(\mathbf{X}^x\otimes I_{\mathbf{C}})P^x_{\mathrm{HL}} \\ &\qquad + P^x_{\mathrm{LL}}(I_{\mathbf{B}}\otimes\mathbf{Y}^x)P^x_{\mathrm{HL}} + P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{HL}}\Big]|\phi^x_{\mathrm{HL}}\rangle \\ &= \frac{1}{4}\Big[\underbrace{(N_x\Pi_x)\otimes O_x}_{=0} + \underbrace{(N_x\mathbf{X}^x\Pi_x)\otimes O_x}_{=0} \\ &\qquad + \underbrace{(N_x\Pi_x)\otimes(O_x\mathbf{Y}^xO_x)}_{=0} + P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{HL}}\Big]|\phi^x_{\mathrm{HL}}\rangle \\ &= \frac{1}{4}P^x_{\mathrm{LL}}\mathbf{Z}|\phi^x_{\mathrm{HL}}\rangle. \end{aligned} \tag{78}$$

Finally, the block from LH into LL expands as

$$\begin{aligned} P^x_{\mathrm{LL}}\mathbf{G}_x|\phi^x_{\mathrm{LH}}\rangle &= P^x_{\mathrm{LL}}\mathbf{G}_x P^x_{\mathrm{LH}}|\phi^x_{\mathrm{LH}}\rangle \\ &= \frac{1}{4}\Big[P^x_{\mathrm{LL}}P^x_{\mathrm{LH}} + P^x_{\mathrm{LL}}(\mathbf{X}^x\otimes I_{\mathbf{C}})P^x_{\mathrm{LH}} \\ &\qquad + P^x_{\mathrm{LL}}(I_{\mathbf{B}}\otimes\mathbf{Y}^x)P^x_{\mathrm{LH}} + P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{LH}}\Big]|\phi^x_{\mathrm{LH}}\rangle \\ &= \frac{1}{4}\Big[\underbrace{N_x\otimes(O_x\Theta_x)}_{=0} + \underbrace{(N_x\mathbf{X}^xN_x)\otimes(O_x\Theta_x)}_{=0} \\ &\qquad + \underbrace{N_x\otimes(O_x\mathbf{Y}^x\Theta_x)}_{=0} + P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{LH}}\Big]|\phi^x_{\mathrm{LH}}\rangle \\ &= \frac{1}{4}P^x_{\mathrm{LL}}\mathbf{Z}|\phi^x_{\mathrm{LH}}\rangle. \end{aligned} \tag{79}$$

Using the definition of M_x in (41), these three identities rewrite the full scalar c_x as

$$\begin{aligned} c_x &= \frac{1}{4}\langle\phi^x_{\mathrm{HL}}|\mathbf{Z}|\phi^x_{\mathrm{LH}}\rangle \\ &\quad + \frac{1}{16}\langle\phi^x_{\mathrm{HL}}|\mathbf{Z}P^x_{\mathrm{LL}}M_x^{-1}P^x_{\mathrm{LL}}\mathbf{Z}|\phi^x_{\mathrm{LH}}\rangle. \end{aligned} \tag{80}$$

We next approximate M_x^{-1} by an inverse involving only $\mathbf{Z}$. On $\mathrm{ran}(P^x_{\mathrm{LL}})$, introduce

$$\begin{aligned} Q_x &:= P^x_{\mathrm{LL}}\mathbf{Z}P^x_{\mathrm{LL}}, \\ R_x &:= P^x_{\mathrm{LL}}\big(\mathbf{X}^x\otimes I_{\mathbf{C}} + I_{\mathbf{B}}\otimes\mathbf{Y}^x\big)P^x_{\mathrm{LL}}, \\ A_x &:= (1+4\eta)P^x_{\mathrm{LL}} - Q_x. \end{aligned} \tag{81}$$

Here and below the identity on $\mathbf{E}$ is suppressed. Multiplying the definition of M_x by four and

using (12) gives the exact operator identity

$$\begin{aligned}
4M_x &= P_{\mathrm{LL}}^x\Big[(2+4\eta)I_{\mathrm{BC}} - 4\mathbf{G}_x\Big]P_{\mathrm{LL}}^x \\
&= P_{\mathrm{LL}}^x\Big[(2+4\eta)I_{\mathrm{BC}} - \big(I_{\mathrm{BC}} + \mathbf{X}^x\otimes I_{\mathrm{C}} + I_{\mathrm{B}}\otimes\mathbf{Y}^x + \mathbf{Z}\big)\Big]P_{\mathrm{LL}}^x \\
&= (1+4\eta)P_{\mathrm{LL}}^x - P_{\mathrm{LL}}^x\mathbf{Z}P_{\mathrm{LL}}^x \\
&\quad - P_{\mathrm{LL}}^x\big(\mathbf{X}^x\otimes I_{\mathrm{C}} + I_{\mathrm{B}}\otimes\mathbf{Y}^x\big)P_{\mathrm{LL}}^x \\
&= (1+4\eta)P_{\mathrm{LL}}^x - Q_x - R_x = A_x - R_x.
\end{aligned} \tag{82}$$

Since $\Delta_{\mathcal{B}}^r$ and $\Delta_{\mathcal{C}}^r$ are contractions and $\mathbf{Z}$ is Hermitian, we have

$$\begin{aligned}
\|\mathbf{Z}\|_\infty \le \mathbb{E}_r\|\Delta_{\mathcal{B}}^r\otimes\Delta_{\mathcal{C}}^r\|_\infty &= \mathbb{E}_r[\|\Delta_{\mathcal{B}}^r\|_\infty\|\Delta_{\mathcal{C}}^r\|_\infty] \le 1 \\
\implies -P_{\mathrm{LL}}^x \preceq Q_x &= P_{\mathrm{LL}}^x\mathbf{Z}P_{\mathrm{LL}}^x \preceq P_{\mathrm{LL}}^x, \\
A_x = (1+4\eta)P_{\mathrm{LL}}^x - Q_x &\succeq (1+4\eta)P_{\mathrm{LL}}^x - P_{\mathrm{LL}}^x = 4\eta P_{\mathrm{LL}}^x, \\
0 \preceq A_x^{-1} \preceq \frac{1}{4\eta}P_{\mathrm{LL}}^x, &\qquad \big\|A_x^{-1}\big\|_\infty \le \frac{1}{4\eta}.
\end{aligned} \tag{83}$$

By the definition of the low spectral projectors, $\|N_x\mathbf{X}^xN_x\|_\infty \le \tau$ and $\|O_x\mathbf{Y}^xO_x\|_\infty \le \tau$. Since $P_{\mathrm{LL}}^x = N_x\otimes O_x$ (with the identity on $\mathbf{E}$ suppressed), the two compressed operators satisfy

$$\begin{aligned}
P_{\mathrm{LL}}^x(\mathbf{X}^x\otimes I_{\mathrm{C}})P_{\mathrm{LL}}^x &= (N_x\mathbf{X}^xN_x)\otimes O_x, \\
\|P_{\mathrm{LL}}^x(\mathbf{X}^x\otimes I_{\mathrm{C}})P_{\mathrm{LL}}^x\|_\infty &= \|N_x\mathbf{X}^xN_x\|_\infty\|O_x\|_\infty \le \tau = \eta^3, \\
P_{\mathrm{LL}}^x(I_{\mathrm{B}}\otimes\mathbf{Y}^x)P_{\mathrm{LL}}^x &= N_x\otimes(O_x\mathbf{Y}^xO_x), \\
\|P_{\mathrm{LL}}^x(I_{\mathrm{B}}\otimes\mathbf{Y}^x)P_{\mathrm{LL}}^x\|_\infty &= \|N_x\|_\infty\|O_x\mathbf{Y}^xO_x\|_\infty \le \tau = \eta^3,
\end{aligned} \tag{84}$$

so

$$\|R_x\|_\infty \le 2\eta^3. \tag{85}$$

Finally, (56) and (82) imply

$$A_x - R_x = 4M_x \succeq (4\eta - 2\eta^3)P_{\mathrm{LL}}^x, \qquad \big\|(A_x-R_x)^{-1}\big\|_\infty \le \frac{1}{4\eta-2\eta^3}. \tag{86}$$

Remark (the resolvent identity used here). *Let $\mathcal{H}_x := \mathrm{ran}(P_{\mathrm{LL}}^x)$, whose identity operator is the restriction of P_{LL}^x. For an operator T on $\mathcal{H}_x$ and a complex number z for which $T - zI_{\mathcal{H}_x}$ is invertible, use the convention*

$$\mathcal{R}_T(z) := (T - zI_{\mathcal{H}_x})^{-1}.$$

The second resolvent identity in this convention is standard; see, for example, [Tes14, Lemma 6.5, Eq. (6.4)]:

$$\mathcal{R}_S(z) - \mathcal{R}_T(z) = \mathcal{R}_S(z)(T-S)\mathcal{R}_T(z).$$

The bounds in (83) and (86) show that A_x and $A_x - R_x$ are invertible on $\mathcal{H}_x$. Thus

$$A_x^{-1} = \mathcal{R}_{A_x}(0), \qquad (A_x-R_x)^{-1} = \mathcal{R}_{A_x-R_x}(0).$$

Taking $S = A_x - R_x$, $T = A_x$, and $z = 0$ in the second resolvent identity gives

$$(A_x-R_x)^{-1} - A_x^{-1} = (A_x-R_x)^{-1}R_xA_x^{-1}.$$

Thus the two inverse operators to which the second resolvent identity is applied, A_x^{-1} and $(A_x - R_x)^{-1}$, are resolvents at zero. Neither A_x nor R_x is itself a resolvent; R_x is the perturbation relating $A_x - R_x$ to A_x.

Using (82) together with this identity gives

$$\begin{aligned} M_x^{-1} - 4A_x^{-1} &= 4\big((A_x - R_x)^{-1} - A_x^{-1}\big) \\ &= 4(A_x - R_x)^{-1} R_x A_x^{-1}. \end{aligned} \tag{87}$$

Taking norms and substituting (83)–(86),

$$\big\|M_x^{-1} - 4A_x^{-1}\big\|_\infty \le \frac{8\eta^3}{(4\eta)(4\eta - 2\eta^3)} = \frac{\eta}{2 - \eta^2} \le \eta, \tag{88}$$

where the last inequality holds in particular for $0 < \eta \le 1/4$.

Define the scalar obtained by replacing M_x^{-1} with $4A_x^{-1}$:

$$\begin{aligned} \widetilde{c}_x := \frac{1}{4}\langle \phi_{\mathrm{HL}}^x | \mathbf{Z} | \phi_{\mathrm{LH}}^x \rangle \\ + \frac{1}{4}\langle \phi_{\mathrm{HL}}^x | \mathbf{Z} P_{\mathrm{LL}}^x A_x^{-1} P_{\mathrm{LL}}^x \mathbf{Z} | \phi_{\mathrm{LH}}^x \rangle. \end{aligned} \tag{89}$$

Equations (80) and (88), together with $\|\mathbf{Z}\|_\infty \le 1$, give the pointwise estimate

$$\begin{aligned} |c_x - \widetilde{c}_x| &\le \frac{1}{16} \||\phi_{\mathrm{HL}}^x\rangle\| \big\|M_x^{-1} - 4A_x^{-1}\big\|_\infty \||\phi_{\mathrm{LH}}^x\rangle\| \\ &\le \frac{\eta}{16} \||\phi_{\mathrm{HL}}^x\rangle\| \||\phi_{\mathrm{LH}}^x\rangle\|. \end{aligned} \tag{90}$$

Averaging and applying Cauchy–Schwarz therefore yields

$$|\mathbb{E}_x(c_x - \widetilde{c}_x)| \le \frac{\eta}{16}\sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}}. \tag{91}$$

Remark (Neumann expansion). *The Neumann expansion is the operator analogue of the scalar geometric series. If B is an operator with $\|B\|_\infty < 1$, then*

$$(I - B)^{-1} = \sum_{k=0}^{\infty} B^k.$$

Indeed, its finite partial sums satisfy

$$(I - B)\sum_{k=0}^{m} B^k = I - B^{m+1},$$

and $\big\|B^{m+1}\big\|_\infty \le \|B\|_\infty^{m+1} \to 0$. Thus the series converges in operator norm and its limit is the inverse of $I - B$. In our setting the underlying space is $\mathrm{ran}(P_{\mathrm{LL}}^x)$, its identity is P_{LL}^x, and

$$B = \frac{Q_x}{1 + 4\eta}.$$

The first line of (92) *verifies the required condition* $\|B\|_\infty < 1$. *See, for example, Horn and Johnson [HJ12, p. 365].*

It remains to expand A_x^{-1}. On $\mathrm{ran}(P_{\mathrm{LL}}^x)$, the projection P_{LL}^x is the identity. The norm condition and the resulting Neumann expansion are

$$\begin{aligned}
\left\| \frac{Q_x}{1+4\eta} \right\|_\infty &\le \frac{\|Q_x\|_\infty}{1+4\eta} \le \frac{\|\mathbf{Z}\|_\infty}{1+4\eta} \le \frac{1}{1+4\eta} < 1, \\
A_x^{-1} &= \Big[(1+4\eta) P_{\mathrm{LL}}^x - Q_x \Big]^{-1} \\
&= \frac{1}{1+4\eta} \left[P_{\mathrm{LL}}^x - \frac{Q_x}{1+4\eta} \right]^{-1} \\
&= \frac{1}{1+4\eta} \sum_{k=0}^{\infty} \left(\frac{Q_x}{1+4\eta} \right)^k \\
&= \frac{P_{\mathrm{LL}}^x}{1+4\eta} + \frac{Q_x}{(1+4\eta)^2} + \frac{Q_x^2}{(1+4\eta)^3} + \cdots \\
&= \sum_{k=0}^{\infty} (1+4\eta)^{-(k+1)} Q_x^k. \qquad (92)
\end{aligned}$$

Substituting (92) into (89), and recalling $Q_x = P_{\mathrm{LL}}^x \mathbf{Z} P_{\mathrm{LL}}^x$, shows that

$$\mathbb{E}_x \widetilde{c}_x = \frac{1}{4} \left[\Gamma_1 + \sum_{k=0}^{\infty} (1+4\eta)^{-(k+1)} \Gamma_{k+2} \right]. \qquad (93)$$

The exchange of the expectation and the series is justified by norm convergence; more explicitly, the absolute value of the kth averaged matrix element, including its coefficient, is at most $\sqrt{W_{\mathrm{HL}} W_{\mathrm{LH}}}(1+4\eta)^{-(k+1)}$, whose sum is finite. Combining (91) and (93) proves (76). □

The scalar Γ_m has a direct algebraic interpretation. Starting with the LH component, one applies m copies of $\mathbf{Z}$. After each of the first $m-1$ copies, one keeps only the LL component; after the final copy, one takes the overlap with the HL component. The projections in (75) only select vector components in the calculation. They are not measurements performed in the original prediction protocol.

3.5 Each common-challenge contribution defines an extractor

We now prove the estimate on Γ_m that is needed to sum the geometric series. Expanding the m copies of $\mathbf{Z}$ produces an average over a challenge sequence $w = (r_1, \ldots, r_m)$. For each fixed sequence, the local products that appear on Bob's and Charlie's sides can be normalized into Kraus operators of valid local measurements whose outcome is a candidate string x. If the corresponding scalar contribution were large, these measurements would jointly recover the full string with probability larger than p_{srch}.

Lemma 3.6 (Every common-challenge contribution is controlled by full-string extraction)**.** *For every integer* $m \ge 1$,

$$|\Gamma_m| \le \frac{(m+1)\sqrt{p_{\mathsf{srch}}}}{\tau^2}. \qquad (94)$$

Proof. Let $r_1, \dots, r_m$ be independent and uniform in $\mathbb{F}_2^n$, and write $w = (r_1, \dots, r_m)$ and $\mathbb{E}_w = \mathbb{E}_{r_1,\dots,r_m}$. For a fixed word w, define

$$V^B_{x,w} := \Pi_x \Delta_{\mathcal{B}}^{r_m} N_x \Delta_{\mathcal{B}}^{r_{m-1}} N_x \cdots \Delta_{\mathcal{B}}^{r_1} N_x, \tag{95}$$

$$V^C_{x,w} := O_x \Delta_{\mathcal{C}}^{r_m} O_x \Delta_{\mathcal{C}}^{r_{m-1}} O_x \cdots \Delta_{\mathcal{C}}^{r_1} \Theta_x. \tag{96}$$

For $m = 1$, these mean $V^B_{x,(r_1)} = \Pi_x \Delta_{\mathcal{B}}^{r_1} N_x$ and $V^C_{x,(r_1)} = O_x \Delta_{\mathcal{C}}^{r_1} \Theta_x$.

We verify explicitly how these products arise. Because we are in the case (39), the retained HL and LH vectors equal the corresponding components of $|\psi^x\rangle$. For $m \geq 2$, expand each copy of $\mathbf{Z} = \mathbb{E}_r[\Delta_{\mathcal{B}}^r \otimes \Delta_{\mathcal{C}}^r]$ in (75) and insert the endpoint projections. This gives

$$\Gamma_m = \mathbb{E}_w \mathbb{E}_x \langle \psi^x | P^x_{\mathrm{HL}} (\Delta_{\mathcal{B}}^{r_m} \otimes \Delta_{\mathcal{C}}^{r_m}) P^x_{\mathrm{LL}} \cdots P^x_{\mathrm{LL}} (\Delta_{\mathcal{B}}^{r_1} \otimes \Delta_{\mathcal{C}}^{r_1}) P^x_{\mathrm{LH}} | \psi^x \rangle. \tag{97}$$

For $m = 1$, the same formula has no P^x_{LL} and follows from (74). Substituting the tensor-product formulas for $P^x_{\mathrm{HL}}, P^x_{\mathrm{LH}}, P^x_{\mathrm{LL}}$ and collecting the Bob and Charlie factors yields

$$P^x_{\mathrm{HL}} (\Delta_{\mathcal{B}}^{r_m} \otimes \Delta_{\mathcal{C}}^{r_m}) P^x_{\mathrm{LL}} \cdots P^x_{\mathrm{LL}} (\Delta_{\mathcal{B}}^{r_1} \otimes \Delta_{\mathcal{C}}^{r_1}) P^x_{\mathrm{LH}} = V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbb{E}}. \tag{98}$$

Reading from right to left confirms the order: Bob starts low, remains low after each intermediate challenge operator, and ends high; Charlie starts high, enters the low subspace after the first challenge, and remains low. Therefore (97) becomes

$$\Gamma_m = \mathbb{E}_w \mathbb{E}_x \langle \psi^x | (V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbb{E}}) | \psi^x \rangle. \tag{99}$$

The maps $V^B_{x,w}, V^C_{x,w}$ need not be positive and are not themselves POVM effects. Their adjoint-products will become effects after the normalization below.

We next show that the maps can be normalized into measurements. For every fixed w,

$$\sum_x (V^B_{x,w})^\dagger V^B_{x,w} \preceq (m+1)^2 \tau^{-2} I_{\mathrm{B}}. \tag{100}$$

To prove this, fix a unit vector $|\zeta\rangle$. Repeatedly insert $N_x = I - \Pi_x$ to obtain

$$\begin{aligned} &\Delta_{\mathcal{B}}^{r_m} N_x \cdots \Delta_{\mathcal{B}}^{r_1} N_x - \Delta_{\mathcal{B}}^{r_m} \cdots \Delta_{\mathcal{B}}^{r_1} \\ &\quad = -\sum_{j=1}^{m} \Delta_{\mathcal{B}}^{r_m} N_x \cdots \Delta_{\mathcal{B}}^{r_{j+1}} N_x \Delta_{\mathcal{B}}^{r_j} \Pi_x \Delta_{\mathcal{B}}^{r_{j-1}} \cdots \Delta_{\mathcal{B}}^{r_1}. \end{aligned} \tag{101}$$

Multiplying (101) on the left by Π_x gives the operator identity

$$\begin{aligned} V^B_{x,w} = \Pi_x &\Delta_{\mathcal{B}}^{r_m} \cdots \Delta_{\mathcal{B}}^{r_1} \\ &- \sum_{j=1}^{m} \Pi_x \Delta_{\mathcal{B}}^{r_m} N_x \cdots \Delta_{\mathcal{B}}^{r_{j+1}} N_x \Delta_{\mathcal{B}}^{r_j} \Pi_x \Delta_{\mathcal{B}}^{r_{j-1}} \cdots \Delta_{\mathcal{B}}^{r_1}. \end{aligned} \tag{102}$$

For a unit vector $|\zeta\rangle$, define

$$|\zeta_0\rangle := \Delta_{\mathcal{B}}^{r_m} \cdots \Delta_{\mathcal{B}}^{r_1} |\zeta\rangle, \qquad |\zeta_j\rangle := \Delta_{\mathcal{B}}^{r_{j-1}} \cdots \Delta_{\mathcal{B}}^{r_1} |\zeta\rangle \quad (1 \leq j \leq m), \tag{103}$$

where the product defining $|\zeta_1\rangle$ is empty and therefore equals the identity. These vectors do not depend on x, and each has norm at most one. In the jth summand of (102), every operator to the left of the displayed inner copy of Π_x is a contraction. The triangle inequality therefore gives

$$\left\|V^B_{x,w}|\zeta\rangle\right\| \le \sum_{j=0}^{m} \|\Pi_x|\zeta_j\rangle\|, \tag{104}$$

Remark (Minkowski's inequality). *For any nonnegative numbers $a_{j,x}$, Minkowski's inequality for the ℓ_2 norm states that*

$$\left(\sum_x \left(\sum_{j=0}^{m} a_{j,x}\right)^2\right)^{1/2} \le \sum_{j=0}^{m} \left(\sum_x a_{j,x}^2\right)^{1/2}.$$

This is simply the triangle inequality in Euclidean space. For each j, regard $a_j = (a_{j,x})_x$ as a vector indexed by x; the displayed inequality is $\left\|\sum_j a_j\right\|_2 \le \sum_j \|a_j\|_2$. In our application, $a_{j,x} = \|\Pi_x|\zeta_j\rangle\|$. Equation (104) *first bounds each $\left\|V^B_{x,w}|\zeta\rangle\right\|$ by $\sum_j a_{j,x}$; monotonicity of the ℓ_2 norm and the displayed inequality then give the first line below.*

Taking the ℓ_2 norm over x, applying Minkowski's inequality, and using (26),

$$\begin{aligned}\left(\sum_x \left\|V^B_{x,w}|\zeta\rangle\right\|^2\right)^{1/2} &\le \sum_{j=0}^{m} \left(\sum_x \|\Pi_x|\zeta_j\rangle\|^2\right)^{1/2} \\ &= \sum_{j=0}^{m} \left(\langle\zeta_j| \left(\sum_x \Pi_x\right) |\zeta_j\rangle\right)^{1/2} \\ &\le \frac{m+1}{\tau}.\end{aligned} \tag{105}$$

Squaring (105) shows that every unit vector $|\zeta\rangle$ satisfies

$$\langle\zeta| \left(\sum_x (V^B_{x,w})^\dagger V^B_{x,w}\right) |\zeta\rangle = \sum_x \left\|V^B_{x,w}|\zeta\rangle\right\|^2 \le \frac{(m+1)^2}{\tau^2}. \tag{106}$$

The quadratic-form characterization of positive-semidefinite order now proves (100).

For Charlie, write

$$V^C_{x,w} = K^C_{x,w}\Theta_x, \qquad K^C_{x,w} := O_x \Delta_{\mathcal{C}}^{r_m} O_x \Delta_{\mathcal{C}}^{r_{m-1}} O_x \cdots \Delta_{\mathcal{C}}^{r_1}. \tag{107}$$

Every factor in $K^C_{x,w}$ is a contraction, so $\left\|K^C_{x,w}\right\|_\infty \le 1$. Hence

$$(V^C_{x,w})^\dagger V^C_{x,w} = \Theta_x (K^C_{x,w})^\dagger K^C_{x,w} \Theta_x \preceq \Theta_x. \tag{108}$$

Summing this inequality and invoking (26) gives

$$\sum_x (V^C_{x,w})^\dagger V^C_{x,w} \preceq \tau^{-2} I_{\mathbb{C}}. \tag{109}$$

For this fixed w, define effects indexed by candidate output strings:

$$F^x_{\mathcal{B},w} := \frac{\tau^2}{(m+1)^2}(V^B_{x,w})^\dagger V^B_{x,w}, \qquad F^x_{\mathcal{C},w} := \tau^2 (V^C_{x,w})^\dagger V^C_{x,w}. \tag{110}$$

Equations (100) and (109) imply

$$\sum_x F^x_{\mathcal{B},w} \preceq I_{\mathbf{B}}, \qquad \sum_x F^x_{\mathcal{C},w} \preceq I_{\mathbf{C}}, \tag{111}$$

so the two families are local sub-POVMs. Fix any string $x_\star$. Adding the residual effects

$$I_{\mathbf{B}} - \sum_x F^x_{\mathcal{B},w}, \qquad I_{\mathbf{C}} - \sum_x F^x_{\mathcal{C},w} \tag{112}$$

to the outcomes labeled $x_\star$ completes the families to full POVMs. Every completed effect dominates the corresponding sub-POVM effect, so this completion cannot decrease their joint correct-label probability.

The completed measurements may depend on the fixed auxiliary word w, which is hardwired independently of the hidden string. This is allowed because p_{srch} upper-bounds every pair of local full-string POVMs. The index x in (110) is only a candidate output label: the constructed measurement is not given the true hidden string. Applying the definition of p_{srch} to these completed POVMs and then using the subeffects gives

$$\begin{aligned} p_{\mathsf{srch}} &\geq \mathbb{E}_x \operatorname{Tr}\left[(F^x_{\mathcal{B},w} \otimes F^x_{\mathcal{C},w})\rho_x\right] \\ &= \frac{\tau^4}{(m+1)^2}\mathbb{E}_x \operatorname{Tr}\left[\left((V^B_{x,w})^\dagger V^B_{x,w} \otimes (V^C_{x,w})^\dagger V^C_{x,w}\right)\rho_x\right] \\ &= \frac{\tau^4}{(m+1)^2}\mathbb{E}_x \left\|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle\right\|^2. \end{aligned} \tag{113}$$

Rearranging (113) yields

$$\mathbb{E}_x \left\|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle\right\|^2 \leq \frac{(m+1)^2 p_{\mathsf{srch}}}{\tau^4}. \tag{114}$$

For this fixed word, put

$$s_w := \mathbb{E}_x \langle\psi^x|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle. \tag{115}$$

Since every $|\psi^x\rangle$ is a unit vector, the triangle inequality, Cauchy–Schwarz for each inner product, and Cauchy–Schwarz over x give

$$\begin{aligned} |s_w| &\leq \mathbb{E}_x \left|\langle\psi^x|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle\right| \\ &\leq \mathbb{E}_x \left\|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle\right\| \\ &\leq \left(\mathbb{E}_x \left\|(V^B_{x,w} \otimes V^C_{x,w} \otimes I_{\mathbf{E}})|\psi^x\rangle\right\|^2\right)^{1/2} \\ &\leq \frac{(m+1)\sqrt{p_{\mathsf{srch}}}}{\tau^2}, \end{aligned} \tag{116}$$

where the last step is (114). Finally, (99) and the triangle inequality over the uniform word w yield

$$|\Gamma_m| = |\mathbb{E}_w s_w| \leq \mathbb{E}_w |s_w| \leq \frac{(m+1)\sqrt{p_{\mathsf{srch}}}}{\tau^2}. \tag{117}$$

This is (94). □

3.6 Complete the top-down estimate

We now return to the single remaining target (73). Combining (76) and (94),

$$|\mathbb{E}_x c_x| \le \frac{1}{4}\left[2+\sum_{k=0}^{\infty}(k+3)(1+4\eta)^{-(k+1)}\right]\frac{\sqrt{p_{\mathsf{srch}}}}{\tau^2} + \frac{\eta}{16}\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}}. \tag{118}$$

The geometric-series identities

$$\sum_{k=0}^{\infty}(1+4\eta)^{-(k+1)} = \frac{1}{4\eta}, \qquad \sum_{k=0}^{\infty}(k+1)(1+4\eta)^{-(k+1)} = \frac{1+4\eta}{16\eta^2}$$

give

$$\frac{1}{4}\left[2+\sum_{k=0}^{\infty}(k+3)(1+4\eta)^{-(k+1)}\right] = \frac{1}{2}+\frac{1}{64\eta^2}+\frac{3}{16\eta} \le \frac{3}{32\eta^2}, \tag{119}$$

where the final inequality uses $0<\eta\le 1/4$. Thus

$$|\mathbb{E}_x c_x| \le \frac{3\sqrt{p_{\mathsf{srch}}}}{32\tau^2\eta^2} + \frac{\eta}{16}\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}}. \tag{120}$$

Using (39), (15), and (16),

$$\frac{3\sqrt{p_{\mathsf{srch}}}}{32\tau^2\eta^2\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}}} \le \frac{3\sqrt{p_{\mathsf{srch}}}}{32\tau^2\eta^2\delta} \le \frac{3\eta}{32}. \tag{121}$$

Therefore

$$\frac{|\mathbb{E}_x c_x|}{\sqrt{W_{\mathrm{HL}}W_{\mathrm{LH}}}} \le \frac{3\eta}{32}+\frac{\eta}{16} = \frac{5\eta}{32}. \tag{122}$$

Since

$$\eta - \frac{\eta^3}{2} \ge \frac{31\eta}{32}, \tag{123}$$

the desired bound (73) holds with room to spare. The reduction (43) now gives

$$\mathbb{E}_x\langle\phi^x|\mathbf{G}_x|\phi^x\rangle \le \left(\frac{1}{2}+\eta\right)\mathbb{E}_x\||\phi^x\rangle\|^2 \le \frac{1}{2}+\eta. \tag{124}$$

This conclusion also held in the easy case (38). Therefore it is valid in all cases.

Returning to (34),

$$p_{\mathsf{pred}} \le \frac{1}{2}+\eta+2\sqrt{q_{\mathrm{omit}}}+q_{\mathrm{omit}}. \tag{125}$$

By (33), (15), and (16),

$$q_{\mathrm{omit}} \le \frac{\eta^{22}}{3072^2\eta^{12}} + \frac{3\eta^2}{3072} = \eta^2\left(\frac{\eta^8}{3072^2}+\frac{1}{1024}\right) < \frac{\eta^2}{1000}. \tag{126}$$

Hence

$$2\sqrt{q_{\mathrm{omit}}} + q_{\mathrm{omit}} < \frac{2\eta}{\sqrt{1000}} + \frac{\eta^2}{1000} < 0.064\eta. \tag{127}$$

Substituting into (125) proves (17).

Finally, since $p_{\mathsf{srch}} \geq 2^{-n} > 0$, suppose first that $p_{\mathsf{srch}}^{1/22} < 1/10$. Choose

$$\eta := \frac{5}{2} p_{\mathsf{srch}}^{1/22}. \tag{128}$$

Then $\eta < 1/4$, and (16) holds because

$$\frac{\eta^{22}}{3072^2} = p_{\mathsf{srch}} \frac{(5/2)^{22}}{3072^2} \geq p_{\mathsf{srch}}.$$

Thus

$$p_{\mathsf{pred}} < \frac{1}{2} + 1.064 \cdot \frac{5}{2} p_{\mathsf{srch}}^{1/22} < \frac{1}{2} + 5 p_{\mathsf{srch}}^{1/22}.$$

If $p_{\mathsf{srch}}^{1/22} \geq 1/10$, then $1/2 + 5p_{\mathsf{srch}}^{1/22} \geq 1$, and the theorem follows from $p_{\mathsf{pred}} \leq 1$. This proves (9). The asymptotic consequence follows because taking a fixed positive root preserves negligibility. □

4 Extensions over arbitrary finite fields

This section turns to a prime-power alphabet. After fixing the common-challenge model and its normalization, we prove a uniform cubic theorem for independent challenges over every finite field and derive from it a common-challenge consequence, including an efficient corollary for growing fields.

4.1 Model and normalization

Definition 4.1 (Common-challenge prediction and search)**.** *Let q be a prime power, let $\ell \geq 1$, and put*

$$\mathcal{X}_q := \mathbb{F}_q^\ell, \qquad N_q := |\mathcal{X}_q| = q^\ell.$$

The hidden label x is uniform in $\mathcal{X}_q$, and Bob and Charlie hold an arbitrary state $\rho_x \in \mathcal{D}(\mathbf{B} \otimes \mathbf{C})$. On a common uniform challenge $r \in \mathcal{X}_q$, they use local q-outcome POVMs

$$\{P_a^r : a \in \mathbb{F}_q\}, \qquad \{Q_a^r : a \in \mathbb{F}_q\}.$$

Here $\langle r, x\rangle := \sum_{i=1}^{\ell} r_i x_i$ is evaluated in $\mathbb{F}_q$. Their exact common-challenge prediction probability is

$$p_{\mathrm{com}} := \mathbb{E}_{x,r} \operatorname{Tr}\left[\left(P_{\langle r,x\rangle}^r \otimes Q_{\langle r,x\rangle}^r\right) \rho_x\right]. \tag{129}$$

Their optimal simultaneous full-string recovery probability is

$$p_{\mathrm{search}}^\star := \sup_{\{E_x\},\{F_x\}} \mathbb{E}_x \operatorname{Tr}[(E_x \otimes F_x)\rho_x], \tag{130}$$

where the supremum is over local POVMs indexed by $x \in \mathcal{X}_q$.

The definition samples r from all of $\mathcal{X}_q$. If $p_{\mathrm{com}}^{\neq 0}$ denotes the score conditioned on $r \neq 0$, then

$$p_{\mathrm{com}} \leq \frac{1}{N_q} + \left(1 - \frac{1}{N_q}\right) p_{\mathrm{com}}^{\neq 0}. \tag{131}$$

Thus the natural all-challenge baseline obtained from the nonzero-challenge guessing baseline $1/q$ is $N_q^{-1} + (1 - N_q^{-1})q^{-1}$. Equation (131) relates the two conventions. The independent-challenge theorem below uses a more general two-secret model; Section 4.3 then returns to Definition 4.1.

4.2 Independent challenges over arbitrary finite fields

This section presents a version of the independent-challenge non-local search-to-decision theorem over arbitrary finite fields. This can be viewed as a two-party extension of the one-party quantum Goldreich-Levin for large finite fields by Sudo, Tezuka, Hara, and Yoshida [STH+23, Theorem 1].

Let $q = p^t$ be a prime power, and let $\ell \geq 1$. We consider a classical–quantum state

$$\rho_{\mathbf{XYBC}} = \sum_{x,y \in \mathbb{F}_q^\ell} \pi(x,y)\, |x,y\rangle\langle x,y|_{\mathbf{XY}} \otimes \rho_{\mathbf{BC}}^{x,y}, \tag{132}$$

where the classical registers $\mathbf{X}, \mathbf{Y}$ store the hidden random vectors $X^{\mathsf{B}}, X^{\mathsf{C}}$, respectively, and $\mathbf{B}, \mathbf{C}$ are Bob's and Charlie's original quantum shares. The distribution π is arbitrary. Thus, in particular, the two hidden strings need not be independent or identically distributed. Bob and Charlie receive independent uniform challenges $V^{\mathsf{B}}, V^{\mathsf{C}} \in \mathbb{F}_q^\ell$. On challenges v and w, respectively, they apply local POVMs

$$\mathsf{P}^v = \{P_g^v : g \in \mathbb{F}_q\}, \qquad \mathsf{Q}^w = \{Q_h^w : h \in \mathbb{F}_q\},$$

and return $G^{\mathsf{B}} = g$ and $G^{\mathsf{C}} = h$. Define the two prediction errors by

$$E^{\mathsf{B}} = G^{\mathsf{B}} - \langle X^{\mathsf{B}}, V^{\mathsf{B}} \rangle, \qquad E^{\mathsf{C}} = G^{\mathsf{C}} - \langle X^{\mathsf{C}}, V^{\mathsf{C}} \rangle, \tag{133}$$

where $\langle x, v \rangle = \sum_{i=1}^{\ell} x_i v_i$ is evaluated in $\mathbb{F}_q$.

For later comparison, let $p_{\mathrm{search}}^{\star}$ denote the optimal simultaneous full-search probability:

$$p_{\mathrm{search}}^{\star} := \sup_{\widehat{\mathsf{P}}, \widehat{\mathsf{Q}}} \sum_{x,y} \pi(x,y)\, \mathrm{Tr}\left[\left(\widehat{P}_x \otimes \widehat{Q}_y\right) \rho_{\mathbf{BC}}^{x,y} \right], \tag{134}$$

where the supremum is over local POVMs with outcomes in $\mathbb{F}_q^\ell$.

Theorem 4.2 (Independent-challenge extraction over $\mathbb{F}_q$)**.** *Suppose that, for some $0 < \epsilon \leq 1 - 1/q$,*

$$\Pr\left[E^{\mathsf{B}} + E^{\mathsf{C}} = 0\right] \geq \frac{1}{q} + \epsilon. \tag{135}$$

Then there are local full-string measurements for Bob and Charlie such that

$$\Pr\left[\widehat{X}^{\mathsf{B}} = X^{\mathsf{B}} \ \wedge \ \widehat{X}^{\mathsf{C}} = X^{\mathsf{C}}\right] \geq \frac{q^2}{8(q-1)^2}\, \epsilon^3 \geq \frac{\epsilon^3}{8}. \tag{136}$$

Equivalently,

$$\Pr\left[E^{\mathsf{B}} + E^{\mathsf{C}} = 0\right] \leq \frac{1}{q} + 2\left(\frac{q-1}{q}\right)^{2/3} \left(p_{\mathrm{search}}^{\star}\right)^{1/3} \leq \frac{1}{q} + 2\left(p_{\mathrm{search}}^{\star}\right)^{1/3}. \tag{137}$$

The same upper bound applies to the smaller exact-prediction event

$$\left\{G^{\mathsf{B}} = \langle X^{\mathsf{B}}, V^{\mathsf{B}}\rangle \ \wedge \ G^{\mathsf{C}} = \langle X^{\mathsf{C}}, V^{\mathsf{C}}\rangle\right\}.$$

In particular, for a security-parameter-indexed family, negligible optimal simultaneous search probability implies that the advantage over $1/q$ in the relational experiment, and hence in the exact-prediction experiment, is negligible. This remains true when q depends on the security parameter, because the field-size factor in (137) is at most one.

Proof. An *additive character* of the finite field $\mathbb{F}_q$ is a map $\chi : \mathbb{F}_q \to \mathbb{C}$ satisfying

$$|\chi(s)| = 1, \qquad \chi(s+u) = \chi(s)\chi(u) \qquad (s, u \in \mathbb{F}_q).$$

Thus, a character converts addition in the field into multiplication of complex phases. Its basic properties include

$$\chi(0) = 1, \qquad \chi(-s) = \overline{\chi(s)}, \qquad \chi(s)^p = 1.$$

The last identity follows because $\mathbb{F}_q$ has characteristic p: adding any field element to itself p times gives zero. A character is *nontrivial* if it is not identically one; this does not require $\chi(s) \neq 1$ for every nonzero s.

We use a specific nontrivial additive character, namely the canonical trace-based choice [Con, Section 2]

$$\chi(s) := \exp\left(\frac{2\pi i}{p}\ \mathrm{Tr}_{\mathbb{F}_q/\mathbb{F}_p}(s)\right). \tag{138}$$

Here $q = p^t$, and the *field trace* is the additive map

$$\mathrm{Tr}_{\mathbb{F}_q/\mathbb{F}_p}(s) = s + s^p + \cdots + s^{p^{t-1}} \in \mathbb{F}_p.$$

In the exponential, its value is represented by an integer modulo p. This is a field trace, rather than the trace of a density operator. When $q = p$, the field trace is the identity, so $\chi(s) = e^{2\pi i s/p}$; in particular, over $\mathbb{F}_2$ this becomes $\chi(s) = (-1)^s$.

For a nontrivial character, the sum of its values over the field is zero. For $s \neq 0$, multiplication by s permutes $\mathbb{F}_q$, whereas for $s = 0$ every term below equals one. Consequently, character orthogonality gives

$$\sum_{a \in \mathbb{F}_q} \chi(as) = q\,\mathbf{1}[s = 0] \qquad (s \in \mathbb{F}_q). \tag{139}$$

Step 1: Identifying the good set. For a hidden pair $z = (x, y)$ in the support of π, set

$$p_z := \Pr\left[E^{\mathsf{B}} + E^{\mathsf{C}} = 0 \ \middle|\ X^{\mathsf{B}} = x,\ X^{\mathsf{C}} = y\right]. \tag{140}$$

Call z good if $p_z \geq 1/q + \epsilon/2$, and let μ be the total π-mass of the good pairs. Since $p_z \leq 1$ for every z,

$$\frac{1}{q} + \epsilon \leq (1-\mu)\left(\frac{1}{q} + \frac{\epsilon}{2}\right) + \mu.$$

It follows that

$$\mu \geq \frac{\epsilon/2}{1 - 1/q - \epsilon/2} \geq \frac{\epsilon}{2}. \tag{141}$$

The good set is used only in the analysis; the extraction measurements do not need to recognize it.

Step 2: Extraction procedure, parameterized by a. We use the standard quantum Fourier transform over the additive group $(\mathbb{F}_q^\ell, +)$, namely the coordinatewise finite-field transform of de Beaudrap, Cleve, and Watrous [BCW02, Definition 2.1]: for $u \in \mathbb{F}_q^\ell$,

$$F_{q,\ell}|u\rangle := q^{-\ell/2} \sum_{v \in \mathbb{F}_q^\ell} \chi(\langle u, v\rangle)|v\rangle. \tag{142}$$

Thus, if a uniform superposition over v is given the phase $\chi(\langle u, v\rangle)$, applying $F_{q,\ell}^\dagger$ returns $|u\rangle$.

Registers. Each register has a single-letter name, typeset in gray boldface, which we use throughout the procedure and its analysis. Bob and Charlie hold their original quantum shares in $\mathbf{B}$ and $\mathbf{C}$, respectively, with joint state $\rho_{\mathbf{BC}}^{x,y}$. Their challenge registers are $\mathbf{R}$ and $\mathbf{S}$, respectively, with computational bases indexed by $\mathbb{F}_q^\ell$; these same registers will hold the final extraction outputs. Their predictor-answer registers are $\mathbf{A}$ and $\mathbf{D}$, respectively, with computational bases indexed by $\mathbb{F}_q$, and their finite-dimensional workspace registers are $\mathbf{W}$ and $\mathbf{Z}$, respectively. The reference register $\mathbf{E}$ is used only to purify the original state; neither party acts on it. The classical registers $\mathbf{X}, \mathbf{Y}$ in (132) record the hidden labels and are not available to the extractors. We condition on their values $z = (x, y)$ in the analysis, so these two classical registers do not appear in the state evolution below.

All six added registers $\mathbf{R}, \mathbf{S}, \mathbf{A}, \mathbf{D}, \mathbf{W}, \mathbf{Z}$ start in their all-zero state. Juxtaposition denotes a joint system of the named registers, not an additional register. In particular, we write

$$|0\rangle_{\mathbf{AW}} := |0\rangle_{\mathbf{A}} \otimes |0\rangle_{\mathbf{W}}, \qquad |0\rangle_{\mathbf{DZ}} := |0\rangle_{\mathbf{D}} \otimes |0\rangle_{\mathbf{Z}}.$$

Thus Bob's answer and workspace registers together are $\mathbf{AW}$, and Charlie's are $\mathbf{DZ}$.

Fix $z = (x, y)$ and a purification $|\psi_z\rangle_{\mathbf{BCE}}$ of $\rho_{\mathbf{BC}}^{x,y}$. For the analysis, denote the initial state of all registers other than the two challenge registers by

$$|\Psi_z^0\rangle := |\psi_z\rangle_{\mathbf{BCE}} \otimes |0\rangle_{\mathbf{AW}} \otimes |0\rangle_{\mathbf{DZ}}.$$

Tensor products below act on the named registers; local operators are ordered as $\mathbf{BAW}, \mathbf{CDZ}, \mathbf{E}$. Bob and Charlie receive their original shares and add their own zero-initialized registers; they do not need to know or prepare the purification.

Coherent predictors. For each challenge v, choose a unitary U_{B}^v on $\mathbf{BAW}$ whose answer-register measurement implements P^v. Explicitly, its POVM effects obey

$$P_g^v = (I_{\mathbf{B}} \otimes \langle 0|_{\mathbf{AW}})(U_{\mathsf{B}}^v)^\dagger (I_{\mathbf{B}} \otimes |g\rangle\langle g|_{\mathbf{A}} \otimes I_{\mathbf{W}}) U_{\mathsf{B}}^v (I_{\mathbf{B}} \otimes |0\rangle_{\mathbf{AW}}).$$

Define U_{C}^w on $\mathbf{CDZ}$ in the same way for Q_h^w. Running the predictors on a superposition of challenges means applying the controlled unitaries

$$\mathbf{U}_{\mathsf{B}} := \sum_v |v\rangle\langle v|_{\mathbf{R}} \otimes U_{\mathsf{B}}^v, \qquad \mathbf{U}_{\mathsf{C}} := \sum_w |w\rangle\langle w|_{\mathbf{S}} \otimes U_{\mathsf{C}}^w.$$

In particular, these unitaries leave the challenge labels unchanged.

Fix a nonzero multiplier $a \in \mathbb{F}_q^*$, known to both parties. The following is a joint description of two entirely local procedures; every sum over a challenge or Fourier-output label ranges over $\mathbb{F}_q^\ell$.

1. *Prepare the challenges.* Starting with $|0\rangle_{\mathbf{R}}|0\rangle_{\mathbf{S}}|\Psi_z^0\rangle$, each party applies $F_{q,\ell}$ to its own challenge register. The state becomes

   $$q^{-\ell} \sum_{v,w} |v\rangle_{\mathbf{R}}|w\rangle_{\mathbf{S}} \otimes |\Psi_z^0\rangle.$$

   Each local superposition has normalization $q^{-\ell/2}$, so the joint normalization is $q^{-\ell}$.

2. *Run the predictors coherently.* Apply $\mathbf{U}_{\mathsf{B}}$ and $\mathbf{U}_{\mathsf{C}}$, without measuring either answer register. The resulting state is
$$q^{-\ell}\sum_{v,w}|v\rangle_{\mathbf{R}}|w\rangle_{\mathbf{S}}\otimes(U_{\mathsf{B}}^{v}\otimes U_{\mathsf{C}}^{w}\otimes I_{\mathbf{E}})|\Psi_{z}^{0}\rangle.$$
To display the answer registers explicitly, write
$$(U_{\mathsf{B}}^{v}\otimes U_{\mathsf{C}}^{w}\otimes I_{\mathbf{E}})|\Psi_{z}^{0}\rangle=\sum_{g,h\in\mathbb{F}_{q}}|g\rangle_{\mathbf{A}}|h\rangle_{\mathbf{D}}\otimes|\xi_{z;v,w}^{g,h}\rangle_{\mathbf{BWCZE}}.$$
These residual vectors need not be normalized or product states; their squared norms are $\|\xi_{z;v,w}^{g,h}\|^{2}=\mathrm{Tr}[(P_{g}^{v}\otimes Q_{h}^{w})\rho_{\mathbf{BC}}^{x,y}]$.

3. *Turn the answers into phases.* Bob applies the diagonal unitary
$$D_{a}^{\mathsf{B}}:=I_{\mathbf{B}}\otimes\left(\sum_{g\in\mathbb{F}_{q}}\chi(ag)|g\rangle\langle g|_{\mathbf{A}}\right)\otimes I_{\mathbf{W}},$$
and Charlie applies its counterpart D_{a}^{C} on $\mathbf{CDZ}$. These gates change the phases of the answer registers $\mathbf{A},\mathbf{D}$, respectively, and act as the identity on the original shares and workspaces. Thus a branch with answers (g,h) acquires the phase $\chi(ag)\chi(ah)=\chi(a(g+h))$. The state is now
$$q^{-\ell}\sum_{v,w}\sum_{g,h\in\mathbb{F}_{q}}\chi(a(g+h))|v\rangle_{\mathbf{R}}|w\rangle_{\mathbf{S}}\otimes|g\rangle_{\mathbf{A}}|h\rangle_{\mathbf{D}}\otimes|\xi_{z;v,w}^{g,h}\rangle_{\mathbf{BWCZE}}.$$

4. *Run the predictors backward.* Apply $\mathbf{U}_{\mathsf{B}}^{\dagger}$ and $\mathbf{U}_{\mathsf{C}}^{\dagger}$. With
$$K_{v,a}^{\mathsf{B}}:=(U_{\mathsf{B}}^{v})^{\dagger}D_{a}^{\mathsf{B}}U_{\mathsf{B}}^{v},\qquad K_{w,a}^{\mathsf{C}}:=(U_{\mathsf{C}}^{w})^{\dagger}D_{a}^{\mathsf{C}}U_{\mathsf{C}}^{w},$$
the state becomes
$$q^{-\ell}\sum_{v,w}|v\rangle_{\mathbf{R}}|w\rangle_{\mathbf{S}}\otimes(K_{v,a}^{\mathsf{B}}\otimes K_{w,a}^{\mathsf{C}}\otimes I_{\mathbf{E}})|\Psi_{z}^{0}\rangle.$$
Because a phase operation was inserted, this backward computation need not return the answer and workspace registers to zero, nor restore the original state on $\mathbf{BC}$.

5. *Decode the Fourier labels.* Apply $F_{q,\ell}^{\dagger}$ to $\mathbf{R}$ and $\mathbf{S}$. Expanding the two transforms gives the full state
$$q^{-2\ell}\sum_{b,c,v,w}\chi(-\langle b,v\rangle-\langle c,w\rangle)\,|b\rangle_{\mathbf{R}}|c\rangle_{\mathbf{S}}\otimes(K_{v,a}^{\mathsf{B}}\otimes K_{w,a}^{\mathsf{C}}\otimes I_{\mathbf{E}})|\Psi_{z}^{0}\rangle.$$
The additional factor $q^{-\ell}$ comes from the two inverse transforms; their phases have a minus sign because they are inverses.

6. *Rescale and output.* On each challenge register apply the permutation $|b\rangle \mapsto |a^{-1}b\rangle$, with multiplication performed coordinatewise. Measure only the challenge registers in the computational basis and output the resulting vectors $\widehat{X}^{\mathsf{B}}$ and $\widehat{X}^{\mathsf{C}}$. The other local registers may be discarded. Success means that the Fourier labels before rescaling were $b = ax$ and $c = ay$.

Analyzing the extraction probability. To determine when the procedure succeeds, return to the full state in step 5. The rescaling in step 6 produces the correct outputs (x, y) exactly when the Fourier labels are $b = ax$ and $c = ay$. Collecting the terms with these labels, the vector accompanying $|ax\rangle_{\mathbf{R}}|ay\rangle_{\mathbf{S}}$ is

$$|\Upsilon_{z,a}\rangle := \mathbb{E}_{v,w}\chi(-a\langle x, v\rangle - a\langle y, w\rangle) \cdot (K^{\mathsf{B}}_{v,a} \otimes K^{\mathsf{C}}_{w,a} \otimes I_{\mathbf{E}})|\Psi^0_z\rangle.$$

Here $q^{-2\ell}\sum_{v,w} = \mathbb{E}_{v,w}$. This vector retains all the registers **BAWCDZE** and is not normalized. Rescaling changes only the output labels, so Born's rule gives

$$\Pr\left[\widehat{X}^{\mathsf{B}} = x, \widehat{X}^{\mathsf{C}} = y \mid z, a\right] = \|\Upsilon_{z,a}\|^2.$$

Our task is therefore to lower-bound this squared norm.

The success probability counts every possible final state of the answer and workspace registers. For a lower bound, it suffices to count just the successful outcomes in which all four registers $\mathbf{A}, \mathbf{W}, \mathbf{D}, \mathbf{Z}$ are zero. We choose this part because, as the next calculation shows, it can be expressed directly using the prediction POVMs that appear in the hypothesis.

First consider one term with fixed challenges v, w. On Bob's side, the registers $\mathbf{A}, \mathbf{W}$ start at zero, and the operator applied in step 4 is $K^{\mathsf{B}}_{v,a} = (U^v_{\mathsf{B}})^\dagger D^{\mathsf{B}}_a U^v_{\mathsf{B}}$. For an input $|\varphi\rangle_{\mathbf{B}}|0\rangle_{\mathbf{AW}}$, the $\mathbf{B}$-vector accompanying output ancillas $|0\rangle_{\mathbf{AW}}$ is $T^{\mathsf{B}}_{v,a}|\varphi\rangle_{\mathbf{B}}$, where

$$T^{\mathsf{B}}_{v,a} = (I_{\mathbf{B}} \otimes \langle 0|_{\mathbf{AW}})K^{\mathsf{B}}_{v,a}(I_{\mathbf{B}} \otimes |0\rangle_{\mathbf{AW}}) = \sum_{g\in\mathbb{F}_q} \chi(ag)P^v_g,$$
$$T^{\mathsf{C}}_{w,a} = (I_{\mathbf{C}} \otimes \langle 0|_{\mathbf{DZ}})K^{\mathsf{C}}_{w,a}(I_{\mathbf{C}} \otimes |0\rangle_{\mathbf{DZ}}) = \sum_{h\in\mathbb{F}_q} \chi(ah)Q^w_h. \tag{143}$$

The second line is Charlie's identical calculation. To obtain the sums on the right, expand the answer phase D^{B}_a into its answer-register projectors, with coefficients $\chi(ag)$. The unitary and zero-ancilla factors around each projector give precisely P^v_g by the earlier displayed POVM identity; similarly, Charlie obtains Q^w_h. Thus these T operators arise from selecting the desired ancilla output in the actual computation, and act on $\mathbf{B}$ and $\mathbf{C}$, respectively.

Applying this calculation on both sides of each term in $|\Upsilon_{z,a}\rangle$ gives its part with all four added answer and workspace registers zero. After leaving out the explicit zero-ancilla kets, the resulting vector on **BCE** is

$$|\Theta_{z,a}\rangle := \mathbb{E}_{v,w\leftarrow\mathbb{F}_q^\ell}\chi(-a\langle x, v\rangle - a\langle y, w\rangle) \cdot \big(T^{\mathsf{B}}_{v,a} \otimes T^{\mathsf{C}}_{w,a} \otimes I_{\mathbf{E}}\big)|\psi_z\rangle. \tag{144}$$

In other words, the complete successful-output vector decomposes as

$$|\Upsilon_{z,a}\rangle = |\Theta_{z,a}\rangle_{\mathbf{BCE}} \otimes |0\rangle_{\mathbf{AW}} \otimes |0\rangle_{\mathbf{DZ}} + |\Upsilon^{\perp}_{z,a}\rangle,$$

where the joint ancillas of $|\Upsilon^{\perp}_{z,a}\rangle$ are orthogonal to their all-zero state. The two parts are orthogonal, so

$$\Pr\Big[\widehat{X}^{\mathsf{B}} = x, \widehat{X}^{\mathsf{C}} = y \mid z, a\Big] = \|\Upsilon_{z,a}\|^2 = \|\Theta_{z,a}\|^2 + \|\Upsilon^{\perp}_{z,a}\|^2 \geq \|\Theta_{z,a}\|^2.$$

This is the desired lower bound: successful outcomes with nonzero ancillas can only increase the probability. Splitting the vector in this way is an analysis of the existing procedure, not an additional measurement or postselection step performed by the parties.

For fixed v, w, z, write

$$p_z(g, h \mid v, w) := \mathrm{Tr}\big[(P^v_g \otimes Q^w_h)\rho^{x,y}_{\mathrm{BC}}\big]\,. \tag{145}$$

Taking the overlap of (144) with $|\psi_z\rangle$ gives

$$\begin{aligned}\Phi_z(a) &:= \langle \psi_z \mid \Theta_{z,a}\rangle \\ &= \mathbb{E}_{v,w} \sum_{g,h\in\mathbb{F}_q} p_z(g, h \mid v, w)\chi\big(a\big(g - \langle x, v\rangle + h - \langle y, w\rangle\big)\big)\,.\end{aligned} \tag{146}$$

By Cauchy–Schwarz, $\|\Theta_{z,a}\|^2 \geq |\Phi_z(a)|^2$. Hence

$$\Pr\Big[\widehat{X}^{\mathsf{B}} = x, \widehat{X}^{\mathsf{C}} = y \mid z, a\Big] \geq |\Phi_z(a)|^2. \tag{147}$$

Notice that this is only a lower bound: other orthogonal workspace branches may also lead to the correct two outputs.

Connection to Adcock–Cleve and intuition. The local procedure is the finite-field analogue of the single-trial subroutine in Adcock–Cleve's quantum Goldreich–Levin theorem [AC02, proof of Theorem 6, Figure 1]; its formulation with an auxiliary quantum input appears in the full version of [CLL+21, Lemma B.12]. In the binary case, a perfect parity predictor inserts the phases $(-1)^{\langle x,v\rangle}$. These encode the whole string x in the Fourier basis:

$$2^{-\ell/2}\sum_v (-1)^{\langle x,v\rangle}|v\rangle = H^{\otimes\ell}|x\rangle.$$

Applying the Hadamards again therefore returns x. For an imperfect predictor whose success on a fixed x is $1/2 + \beta$, the decoded x component has overlap with the initial purified input and zero ancillas equal to

$$\mathbb{E}_{v,g}\Big[(-1)^{g-\langle x,v\rangle}\Big] = \Pr[g = \langle x, v\rangle] - \Pr[g \neq \langle x, v\rangle] = 2\beta.$$

Here g follows the predictor's output distribution on v. Correct answers contribute $+1$ and incorrect answers -1: prediction bias becomes a coherent amplitude, and squaring yields recovery probability at least $4\beta^2$. Returning the unknown input state is only an overlap used for this bound, not a test performed by the extractor. We use this single trial, not Adcock–Cleve's further amplitude amplification, which uses an equality oracle.

In our setting, both parties perform this operation locally with the same multiplier a. Their shares may be entangled, so we do not multiply two one-party success bounds. Instead, the joint overlap in (146) is $\mathbb{E}[\chi(a(E^{\mathsf{B}} + E^{\mathsf{C}})) \mid z]$. The following step averages over the common multiplier to find useful Fourier energy over an arbitrary finite field.

Step 3: Averaging over a. Let A_z be the distribution of the total error, averaged over the two independent challenges:

$$A_z(s) := \mathbb{E}_{v,w} \sum_{\substack{g,h\in\mathbb{F}_q:\\ g-\langle x,v\rangle+h-\langle y,w\rangle=s}} p_z(g,h \mid v,w). \tag{148}$$

Then $A_z(0) = p_z$, $\sum_s A_z(s) = 1$, and

$$\Phi_z(a) = \sum_{s\in\mathbb{F}_q} A_z(s)\chi(as). \tag{149}$$

Using (139) and the fact that $\Phi_z(0) = 1$, we obtain the exact identity

$$\mathbb{E}_{a\leftarrow\mathbb{F}_q^*}|\Phi_z(a)|^2 = \frac{q\sum_{s\in\mathbb{F}_q} A_z(s)^2 - 1}{q-1}. \tag{150}$$

Suppose now that z is good and put $\eta = \epsilon/2$. For a prescribed value of $A_z(0)$, the sum of the squares of the other $q-1$ probabilities is minimized when those probabilities are equal. Moreover, the resulting lower bound is increasing for $A_z(0) \geq 1/q$. Therefore

$$\begin{aligned}\sum_{s\in\mathbb{F}_q} A_z(s)^2 &\geq \left(\frac{1}{q}+\eta\right)^2 + (q-1)\left(\frac{1}{q}-\frac{\eta}{q-1}\right)^2 \\ &= \frac{1}{q} + \frac{q}{q-1}\eta^2.\end{aligned} \tag{151}$$

Substitution into (150) yields

$$\mathbb{E}_{a\leftarrow\mathbb{F}_q^*}|\Phi_z(a)|^2 \geq \frac{q^2}{(q-1)^2}\eta^2 = \frac{q^2}{4(q-1)^2}\epsilon^2. \tag{152}$$

Step 4: One multiplier works globally. Let $p_{\mathrm{search}}(a)$ be the overall simultaneous search success of the preceding two local procedures with multiplier a hardcoded. Averaging first over the original hidden-pair distribution and then over a, and discarding the nonnegative contribution of the bad pairs, gives

$$\begin{aligned}\mathbb{E}_{a\leftarrow\mathbb{F}_q^*} p_{\mathrm{search}}(a) &\geq \mu\cdot\frac{q^2}{4(q-1)^2}\epsilon^2 \\ &\geq \frac{q^2}{8(q-1)^2}\epsilon^3.\end{aligned} \tag{153}$$

Thus there is one $a_\star \in \mathbb{F}_q^*$, independent of the hidden pair, for which the same lower bound holds. Hardcoding $a_\star$ gives valid local full-string measurements and proves (136). Rearranging this inequality gives (137). Finally, the exact-prediction event is contained in $\{E^{\mathsf{B}} + E^{\mathsf{C}} = 0\}$, which proves the last assertion. □

Remark 4.3 (Relation to Kundu–Tan)**.** *Kundu and Tan prove the corresponding independent-challenge statement for the prime fields $\mathbb{F}_2$ and $\mathbb{F}_3$ [KT25, Lemma 27]. It is not clear if their proof extends to larger fields.*

4.3 Identical Challenge Setting over Arbitrary Finite Fields

The quantitative result below applies to every finite field. It combines the independent-challenge theorem proved in Section 4.2 with the correlated-to-independent inequality of Ananth, Kaleoglu, and Liu [AKL23, Lemma 18]. Its growing-field consequence, where the reciprocal field size is negligible, is stated separately afterward.

Theorem 4.4 (Identical challenge setting over arbitrary finite fields)**.** *For every instance of the common-challenge experiment in Definition 4.1,*

$$p^{\star}_{\mathrm{search}} \geq \frac{q^2}{8(q-1)^2}\left(\frac{p_{\mathrm{com}}^3}{216} - \frac{1}{q}\right)_+^3. \tag{154}$$

Equivalently,

$$p_{\mathrm{com}} \leq 6\left(\frac{1}{q} + 2\left(\frac{q-1}{q}\right)^{2/3}\left(p^{\star}_{\mathrm{search}}\right)^{1/3}\right)^{1/3}. \tag{155}$$

Proof. Let p_{ind} be the exact success probability of the same local response algorithms when Bob and Charlie instead receive independent uniform challenges. For a fixed hidden string x, let c_x and i_x denote the corresponding common- and independent-challenge success probabilities. The correlated-to-independent inequality of Ananth, Kaleoglu, and Liu, applied to ρ_x with the correct-answer effects $B_r := P^r_{\langle r,x\rangle}$ and $C_r := Q^r_{\langle r,x\rangle}$, gives $c_x \leq 6i_x^{1/3}$ for every x. This is a pointwise analytical application; the operational predictor remains the original x-independent family of q-outcome measurements. Averaging over x and using concavity of $u \mapsto u^{1/3}$ gives

$$p_{\mathrm{com}} \leq 6p_{\mathrm{ind}}^{1/3}, \qquad \text{and hence} \qquad p_{\mathrm{ind}} \geq \frac{p_{\mathrm{com}}^3}{216}. \tag{156}$$

In the independent-challenge experiment, exact correctness on both sides implies that the two errors sum to zero. Therefore Theorem 4.2, specialized to $X^{\mathsf{B}} = X^{\mathsf{C}} = x$, gives

$$p_{\mathrm{ind}} \leq \frac{1}{q} + 2\left(\frac{q-1}{q}\right)^{2/3}\left(p^{\star}_{\mathrm{search}}\right)^{1/3}. \tag{157}$$

Combining (156) and (157) proves (155); rearranging proves (154). □

Corollary 4.5 (Efficient consequence for growing fields)**.** *Suppose that $q = q(\lambda)$ has an efficient field representation, that $\log q(\lambda)$ and $\ell(\lambda)$ are polynomially bounded, that $1/q(\lambda)$ is negligible. Consequently, computational joint search security implies*

$$p_{\mathrm{com}} \leq \frac{1}{q(\lambda)} + \mathrm{negl}(\lambda). \tag{158}$$

This applies, for example, to the infinite family $q(\lambda) = 2^{\lambda}$.

Proof. Use the original predictor circuits with independently prepared local challenge registers and then run the coherent extractor from Theorem 4.2, sampling its common multiplier as shared classical randomness before the parties separate. If p_{com} is nonnegligible while $1/q$ is negligible, then (156) shows that the independent exact success exceeds $1/q$ by a nonnegligible amount. Taking the contrapositive proves (158). □

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

References

[GL89] O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 25–32, 1989. doi:10.1145/73007.73010.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 84–93, 2005. doi:10.1145/1060590.1060603.

[AC02] M. Adcock and R. Cleve. A quantum Goldreich–Levin theorem with cryptographic applications. In *STACS 2002*, volume 2285 of *LNCS*, pages 323–334, 2002. doi:10.1007/3-540-45841-7_26.

[HJ12] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Second edition, Cambridge University Press, 2012. doi:10.1017/CBO9781139020411.

[Tes14] G. Teschl. *Mathematical Methods in Quantum Mechanics: With Applications to Schrödinger Operators*. Second edition, volume 157 of *Graduate Studies in Mathematics*, American Mathematical Society, 2014. doi:10.1090/gsm/157.

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

[CLL+21] A. Coladangelo, J. Liu, Q. Liu, and M. Zhandry. Hidden cosets and applications to unclonable cryptography. In *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 556–584, 2021. doi:10.1007/978-3-030-84242-0_20. Full version: `https://arxiv.org/abs/2107.05692`.

[KT25] S. Kundu and E. Y.-Z. Tan. Device-independent uncloneable encryption. *Quantum*, 9:1582, 2025. doi:10.22331/q-2025-01-08-1582.

[AKL23] P. Ananth, F. Kaleoglu, and Q. Liu. Cloning games: A general framework for unclonable primitives. In *Advances in Cryptology – CRYPTO 2023, Part V*, pages 66–98, 2023. doi:10.1007/978-3-031-38554-4_3.

[BKL24] A. Broadbent, M. Karvonen, and S. Lord. Uncloneable quantum advice. *IACR Communications in Cryptology*, 1(3), 2024. doi:10.62056/abe0fhbmo.

[CG24a] A. Coladangelo and S. Gunn. How to use quantum indistinguishability obfuscation. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1003–1008, 2024. doi:10.1145/3618260.3649779.

[AB24] P. Ananth and A. Behera. A modular approach to unclonable cryptography. In *Advances in Cryptology – CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 3–37, 2024. doi:10.1007/978-3-031-68394-7_1.

[AKY25] P. Ananth, F. Kaleoglu, and H. Yuen. Simultaneous Haar indistinguishability with applications to unclonable cryptography. In *ITCS 2025*, volume 325 of *LIPIcs*, article 7, 2025. doi:10.4230/LIPIcs.ITCS.2025.7.

[CHV23] C. Chevalier, P. Hermouet, and Q.-H. Vu. Towards unclonable cryptography in the plain model. Cryptology ePrint Archive, Paper 2023/1825, 2023. ePrint 2023/1825.

[BCG26] A. Behera, A. Çakan, and V. Goyal. Copy-protection with correlated challenges: Point functions and more via decisional coset monogamy. arXiv:2608.18841, 2026. `https://arxiv.org/abs/2608.18841`.

[CG24b] A. Çakan and V. Goyal. Unbounded leakage-resilient encryption and signatures. Cryptology ePrint Archive, Paper 2024/1876, 2024. ePrint 2024/1876.

[CKN+25] J. Champion, F. Kitagawa, R. Nishimaki, and T. Yamakawa. Untelegraphable encryption and its applications. In *Theory of Cryptography – TCC 2025, Part III*, volume 16270 of *LNCS*, pages 3–35, 2025. doi:10.1007/978-3-032-12296-4_1.

[CLX26a] A. Coladangelo, Q. Liu, and Z. Xie. The curious case of "XOR repetition" of monogamy-of-entanglement games. In *ITCS 2026*, volume 362 of *LIPIcs*, article 41, 2026. doi:10.4230/LIPIcs.ITCS.2026.41.

[CLX26b] A. Coladangelo, Q. Liu, and Z. Xie. Unclonable encryption from BB84 states: A simultaneous Goldreich–Levin reduction. arXiv:2608.17629, 2026. `https://arxiv.org/abs/2608.17629`.

[HMS04] T. Holenstein, U. Maurer, and J. Sjödin. Complete classification of bilinear hard-core functions. In *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *LNCS*, pages 73–91, 2004. doi:10.1007/978-3-540-28628-8_5.

[APV23] P. Ananth, A. Poremba, and V. Vaikuntanathan. Revocable cryptography from learning with errors. In *Theory of Cryptography – TCC 2023*, volume 14372 of *LNCS*, pages 93–122, 2023. doi:10.1007/978-3-031-48624-1_4.

[STH+23] K. Sudo, M. Tezuka, K. Hara, and Y. Yoshida. Quantum search-to-decision reduction for the LWE problem. In *Progress in Cryptology – AFRICACRYPT 2023*, volume 14064 of *LNCS*, pages 395–413, 2023. doi:10.1007/978-3-031-37679-5_17.

[BL20] A. Broadbent and S. Lord. Uncloneable quantum encryption via oracles. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, volume 158 of *LIPIcs*, pages 4:1–4:22, 2020. doi:10.4230/LIPIcs.TQC.2020.4.

[AK21] P. Ananth and F. Kaleoglu. Unclonable encryption, revisited. In *Theory of Cryptography – TCC 2021, Part I*, volume 13042 of *LNCS*, pages 299–329, 2021. doi:10.1007/978-3-030-90459-3_11.

[AKL+22] P. Ananth, F. Kaleoglu, X. Li, Q. Liu, and M. Zhandry. On the feasibility of unclonable encryption, and more. In *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 212–241, 2022. doi:10.1007/978-3-031-15979-4_8.

[GZ20] M. Georgiou and M. Zhandry. Unclonable decryption keys. Cryptology ePrint Archive, Paper 2020/877, 2020. ePrint 2020/877.

[AS26] P. Ananth and A. Sahai. Unconditional unclonable encryption. arXiv:2607.21551, 2026. https://arxiv.org/abs/2607.21551.

[Rag26] S. Ragavan. Efficient unclonable encryption from Pauli eigenstates. arXiv:2607.21811, 2026. https://arxiv.org/abs/2607.21811.

[BBC26] A. Bhattacharyya, A. Broadbent, and E. Culf. The uncloneable bit exists. arXiv:2603.08916v2, 2026. https://arxiv.org/abs/2603.08916.

[Con] K. Conrad. Gauss and Jacobi sums on finite fields and $\mathbb{Z}/m\mathbb{Z}$. Expository notes, Sections 1–2. https://kconrad.math.uconn.edu/blurbs/gradnumthy/Gauss-Jacobi-sums.pdf.

[BCW02] J. N. de Beaudrap, R. Cleve, and J. Watrous. Sharp quantum versus classical query complexity separations. *Algorithmica*, 34(4):449–461, 2002. https://doi.org/10.1007/s00453-002-0978-1.

[CW00] R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *FOCS 2000*, pages 526–536, 2000. https://cs.uwaterloo.ca/~watrous/Papers/ParallelCircuitsQFT.pdf.

Glossary of Notation

The table below records the notation used throughout the statement and proof. Operators on $\mathbf{BC}$ are understood to act as the identity on the purification register $\mathbf{E}$ whenever that register is present.

Notation	Meaning
Basic notation and probability	
$[n]$	The set $\{1, \ldots, n\}$.
$\mathbb{F}_2^n$	The n-dimensional vector space of binary strings.
$\langle r, x\rangle$	The mod-2 inner product of $r, x \in \mathbb{F}_2^n$.
$\chi_x(r)$	The character $(-1)^{\langle r,x\rangle}$.
$s \xleftarrow{\$} S$	A uniformly random sample from the finite set S.
$\mathbb{E}_x$	Expectation over a uniformly random x; multiple subscripts denote independent uniform choices unless stated otherwise.
$\mathcal{L}(\mathbf{R}), \mathcal{D}(\mathbf{R})$	Linear operators and density operators on register $\mathbf{R}$.
$I_{\mathbf{R}}$, Tr	Identity on $\mathbf{R}$ and the trace.
$M \preceq N$	Loewner order: $N - M$ is positive semidefinite.
$\|M\|_\infty$	Operator norm of M.
$\mathrm{ran}(P)$	Range of the operator or projection P.
Registers, states, and experiments	
$\mathbf{X}$	Classical input register containing the basis state indexed by x.
$\mathbf{B}, \mathbf{C}$	Quantum share registers held by Bob and Charlie, respectively.
$\mathbf{E}$	Auxiliary register used to purify ρ_x.
Φ	Channel mapping the classical input encoding to the bipartite state on $\mathbf{BC}$.
ρ_x	The bipartite state indexed by x; in the introductory channel formulation, $\rho_x = \Phi(\|x\rangle\langle x\|)$.
$\boldsymbol{\rho}$	The uniformly indexed ensemble $(\rho_x)_{x \in \mathbb{F}_2^n}$.
$\|\psi^x\rangle$	A fixed purification of ρ_x on $\mathbf{BCE}$.
$\Lambda_{\mathcal{B}}^{\mathsf{srch}}, \Lambda_{\mathcal{C}}^{\mathsf{srch}}$	Local full-string POVMs with effects $\{\mathcal{E}_{\mathcal{B}}^z\}_z$ and $\{\mathcal{E}_{\mathcal{C}}^z\}_z$.

Notation	Meaning
$p_{\mathsf{srch}}(\boldsymbol{\rho})$	Optimal probability that both local full-string POVMs output the correct uniformly sampled ensemble index x; it lies in $[2^{-n}, 1]$. Inside the proof this quantity is abbreviated as p_{srch}.
$\Lambda_{\mathcal{B}}^{\mathsf{pred}}[r], \Lambda_{\mathcal{C}}^{\mathsf{pred}}[r]$	Local binary POVMs used when both parties receive challenge r.
p_{pred}	Probability that both parties output $\langle r, x\rangle$ in the same prediction experiment.
Prediction and Fourier operators	
$\Delta_{\mathcal{B}}^r, \Delta_{\mathcal{C}}^r$	Difference observables $\mathcal{P}^{r,0} - \mathcal{P}^{r,1}$ for Bob and Charlie.
$\mathbf{X}^x, \mathbf{Y}^x$	Fourier coefficients $\mathbb{E}_r[\chi_x(r)\Delta_{\mathcal{B}}^r]$ and $\mathbb{E}_r[\chi_x(r)\Delta_{\mathcal{C}}^r]$.
$\mathbf{Z}$	Common-challenge correlation operator $\mathbb{E}_r[\Delta_{\mathcal{B}}^r \otimes \Delta_{\mathcal{C}}^r]$.
$\mathbf{G}_x$	Effect for the event that both parties correctly answer the parity challenge for hidden string x.
η, τ, δ	Proof parameters, with $\tau = \eta^3$ and $\delta = \eta^2/3072$.
High/low decomposition	
Π_x, N_x	Bob's high and low spectral projections for $\mathbf{X}^x$: Π_x projects onto eigenvectors whose eigenvalues λ satisfy $\lvert\lambda\rvert \geq \tau$, and $N_x = I - \Pi_x$.
Θ_x, O_x	Charlie's high and low spectral projections for $\mathbf{Y}^x$: Θ_x projects onto eigenvectors whose eigenvalues μ satisfy $\lvert\mu\rvert \geq \tau$, and $O_x = I - \Theta_x$.
P_U^x	Joint projection for $U \in \{\mathrm{HH}, \mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}$; the first letter refers to Bob and the second to Charlie.
$\lvert\psi_U^x\rangle$	Component $P_U^x\lvert\psi^x\rangle$ of the purified state.
$\mathrm{wt}(U)$	Average component weight $\mathbb{E}_x\lVert\lvert\psi_U^x\rangle\rVert^2$.
κ_U	Indicator that the average weight of component family U is at least δ.
$\lvert\phi_U^x\rangle, \lvert\phi^x\rangle$	Retained component $\kappa_U\lvert\psi_U^x\rangle$ and their sum over $U \in \{\mathrm{HL}, \mathrm{LH}, \mathrm{LL}\}$.
$\lvert e^x\rangle, q_{\mathrm{omit}}$	Omitted vector $\lvert\psi^x\rangle - \lvert\phi^x\rangle$ and its average squared norm.
W_U	Retained average weight $\mathbb{E}_x\lVert\lvert\phi_U^x\rangle\rVert^2$.
Effective cross term and extractors	
M_x	Compression $P_{\mathrm{LL}}^x((1/2+\eta)I - \mathbf{G}_x)P_{\mathrm{LL}}^x$, viewed as an operator on $\mathrm{ran}(P_{\mathrm{LL}}^x)$.
$\lvert u_x\rangle$	Proof-local vector $P_{\mathrm{LL}}^x\mathbf{G}_x(\lvert\phi_{\mathrm{HL}}^x\rangle + \lvert\phi_{\mathrm{LH}}^x\rangle)$ used in completing the square.
c_x	Effective HL–LH cross scalar after minimizing over the LL component.
Q_x, R_x, A_x	Proof-local compressed operators defined in (81); they satisfy $4M_x = A_x - R_x$.
$\widetilde{c}_x$	Approximation to c_x obtained by replacing M_x^{-1} with the geometric inverse.
Γ_m	Averaged contribution containing m copies of the common-challenge operator $\mathbf{Z}$.
$w = (r_1, \ldots, r_m)$	Auxiliary word of independent uniform challenge vectors; after fixing w, it is hardwired into the extraction measurements.
$V_{x,w}^B, V_{x,w}^C$	Fixed-word local products whose normalized adjoint-products define extraction effects.
Arbitrary-field extensions	
λ	Security parameter for asymptotic families; the field size $q = q(\lambda)$ may depend on it.
$q = p^t, \mathbb{F}_q$	A prime power and the finite field of order q; all arbitrary-field inner products are evaluated in $\mathbb{F}_q$.
χ	A fixed nontrivial additive character of $\mathbb{F}_q$, defined using the field trace to $\mathbb{F}_p$.
$X^{\mathsf{B}}, X^{\mathsf{C}}$	Possibly correlated hidden strings in the independent-challenge experiment of Section 4.2.
$\mathbf{X}, \mathbf{Y}$	Classical registers recording these two hidden strings in Section 4.2; not given to the extractors.
$\mathbf{R}, \mathbf{S}$	Bob's and Charlie's challenge registers, respectively; also hold the final extraction outputs.
$\mathbf{A}, \mathbf{D}$	Bob's and Charlie's predictor-answer registers, respectively.
$\mathbf{W}, \mathbf{Z}$	Bob's and Charlie's predictor-workspace registers, respectively.
$V^{\mathsf{B}}, V^{\mathsf{C}}$	Independent uniform challenge vectors given locally to Bob and Charlie in that experiment.
$E^{\mathsf{B}}, E^{\mathsf{C}}$	The corresponding additive prediction errors.
$p_{\mathrm{search}}^{\star}$	Optimal joint full-string recovery probability in the arbitrary-field experiments.

Notation	Meaning
$p_{\mathrm{com}}, p_{\mathrm{ind}}$	Exact joint prediction success with a common challenge and with independent local challenges, respectively, in Theorem 4.4.

AI Disclosure: We used ChatGPT using GPT-5.6 Sol Pro and Codex at ultra reasoning effort to assist with development of all the proofs as well as for the preparation of the paper.

The first 10 pages were mainly prepared by the authors with minor assistance from the above tools, especially in the generation of the figures. All the technical ideas in the paper were suggested by the above tools and were then iteratively polished.

The authors verified the correctness and originality of all content including references.