



Separating Quantum and Classical Advice with Good Codes

John Bostanci¹, Andrew Huang², and Vinod Vaikuntanathan²

¹Columbia University, *New York, NY*

²Massachusetts Institute of Technology, *Cambridge, MA*

Abstract

We show an unconditional classical oracle separation between the class of languages that can be verified using a quantum proof (QMA) and the class of languages that can be verified with a classical proof (QCMA). Compared to the recent work of Bostanci, Haferkamp, Nirkhe, and Zhandry (STOC 2026), our proof is conceptually and technically simpler, and readily extends to other oracle separations. In particular, our techniques yield the *first unconditional* classical oracle separation between the class of languages that can be decided with quantum advice (BQP/qpoly) and the class of languages that can be decided with classical advice (BQP/poly), improving on the quantum oracle separation of Aaronson and Kuperberg (CCC 2007) and the classically-accessible classical oracle separation of Li, Liu, Pelecanos and Yamakawa (ITCS 2024).

Our oracles are based on the code intersection problem introduced by Yamakawa and Zhandry (FOCS 2022), combined with codes that have extremely good list-recovery properties.

Contents

1	Introduction	1
1.1	Our Work	2
2	Technical Overview	2
3	Discussion and Open Questions	6
3.1	Structure versus randomness in classical oracle separations	7
3.2	QMA-completeness of a decoding problem	7
3.3	Simplifications to the separation	8
4	Preliminaries	9
4.1	Notation	9
4.2	Probability and Complexity Theory	9
4.3	Coding Theory	10
4.4	Yao’s Box Problem and Non-Uniform Advice	12
5	The Generalized Code Intersection Problem	13
5.1	Definitions and Basic Facts	13
5.2	Technical Lemmas	15
5.3	The Biased Yamakawa-Zhandry Algorithm	17
6	Separating QMA from QCMA	20
6.1	The QMA Proof System	20
6.2	Non-Existence of QCMA Proof Systems	21
7	Separating BQP/qpoly from BQP/poly	23
A	Duals of Multiplicity Codes	29
B	Diagonalization Arguments	31

1 Introduction

We study the fundamental question of whether quantum proofs and advice are more powerful than their classical counterparts. In the language of complexity theory, we consider whether the complexity classes QMA [Kit97] and QCMA [AN02], two quantum generalizations of NP with quantum and classical proofs respectively, are distinct; and whether the complexity classes BQP/qpoly and BQP/poly, two generalizations of BQP with quantum and classical advice respectively, are distinct.

Recall that the class QMA, first defined by Kitaev [Kit97], is the class of all languages $\mathcal{L}_{\text{YES}}, \mathcal{L}_{\text{NO}} \subseteq \{0, 1\}^*$ for which there is a quantum polynomial-time verifier V and a polynomial function t such that

$$\begin{aligned} x \in \mathcal{L}_{\text{YES}} &\implies \exists |\psi\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)} : \Pr[V(x, |\psi\rangle) = 1] \geq 2/3 \quad \text{and} \\ x \in \mathcal{L}_{\text{NO}} &\implies \forall |\psi^*\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)}, \Pr[V(x, |\psi^*\rangle) = 1] \leq 1/3. \end{aligned}$$

Replacing the quantum witness $|\psi\rangle$ by a classical witness w (but retaining the quantum verifier) results in the class QCMA, first defined by Aharonov and Naveh [AN02]. They posed the question of whether QMA and QCMA are distinct, and it has since then been a long-standing open problem in quantum complexity theory.

Analogously, the class BQP/qpoly [NY04] is the class of all languages $\mathcal{L}_{\text{YES}}, \mathcal{L}_{\text{NO}} \subseteq \{0, 1\}^*$ for which there is a quantum polynomial-time algorithm A , a polynomial function t and a family of advice states $\{|\text{adv}_\lambda\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)}\}_{\lambda \in \mathbb{N}}$ such that for every input x of length λ ,

$$\begin{aligned} x \in \mathcal{L}_{\text{YES}} &\implies \Pr[A(x, |\text{adv}_\lambda\rangle) = 1] \geq 2/3, \quad \text{and} \\ x \in \mathcal{L}_{\text{NO}} &\implies \Pr[A(x, |\text{adv}_\lambda\rangle) = 1] \leq 1/3. \end{aligned}$$

Replacing the quantum advice $|\text{adv}_\lambda\rangle$ by a classical advice string a_λ (but retaining the quantumness of the algorithm A) results in the class BQP/poly. Nishimura and Yamakami [NY04] first posed the question of whether BQP/qpoly and BQP/poly are distinct complexity classes, and it too has been a long-standing open problem in quantum complexity theory.

Given that unconditional separations of these classes seem far out of reach of current technology as they would imply breakthrough results like $\text{P} \neq \text{PSPACE}$ and $\text{PP} \subsetneq \text{P/poly}$, research attention has shifted to providing oracular evidence. The first progress towards this came from the seminal work of Aaronson and Kuperberg [AK07], who introduced the weaker notion of a *unitary*, or *quantum*, oracle as a means to separate QMA and QCMA, as well as BQP/qpoly and BQP/poly, and asked whether their quantum oracle separations could be strengthened to a classical oracle separation. After a long series of works which demonstrated either conditional or non-standard separations [FK15, LLPY23, NN24, BDK24, Zha24, LMY25], the first unconditional classical oracle separation between QMA and QCMA appeared in the recent work of Bostanci, Haferkamp, Nirkhe and Zhandry [BHNZ25]. Even so, the case of BQP/qpoly versus BQP/poly has remained wide open.

The oracle used in the separation of [BHNZ25] (introduced in [Zha24]) leverages one of the most natural properties to separate quantum from classical computers: whether two functions are related by the Hadamard transform. This idea was first exploited in the work of [Aar10], who presented the ‘‘Forrelation’’ problem, i.e., deciding whether two functions are related by the Hadamard transform, as a candidate separation between BQP and PH. The works of [Zha24] and [BHNZ25] lift this problem to QMA with a variant dubbed ‘‘spectral Forrelation’’. The separation of [BHNZ25] employs sophisticated techniques to analyze Forrelated oracles, including ideas inspired by mathematical physics and Hamiltonian learning, resulting in a rather complicated proof. It is therefore natural to wonder whether a classical oracle separation requires such heavy technical machinery.

Question 1: *Is there a simple classical oracle separation between QMA and QCMA?*

A closely related problem is to separate BQP/qpoly, problems that can be solved with quantum advice, from BQP/poly, problems that can be solved with classical advice. As opposed to the proof setting, where a verifier gets an untrusted proof that can depend on their specific instance, in the advice setting the verifier receives a *trusted* witness, but is expected to use their witness to solve all instances of the same length. Thus, the questions of separating QMA from QCMA and BQP/qpoly from BQP/poly, while related, are not formally equivalent. The original work of Aaronson and Kuperberg [AK07] also gave a unitary oracle separation for BQP/qpoly and BQP/poly and asked if this too could be lifted to the classical oracle model. Thus, we investigate the following question:

Question 2: *Is there a (standard) classical oracle separation between BQP/qpoly and BQP/poly?*

The only other advice separation follows from the work of Li, Liu, Pelecanos and Yamakawa [LLPY23], in a non-standard “classically-accessible classical oracle” model where all queries to the oracles are restricted to be classical. A true classical oracle separation between BQP/qpoly and BQP/poly remains unresolved.

1.1 Our Work

We re-examine a line of works [Liu23, LLYPY23, BDK24] which have made progress towards a full separation between QMA and QCMA, based on the code intersection problem of Yamakawa and Zhandry [YZ24]. We begin by showing that Question 1 can indeed be resolved using a modification of the code intersection problem when instantiated with codes that have extremely good list-recovery properties.

Theorem 1.1 (Informal). *There exists a classical oracle \mathcal{O} such that $\text{QMA}^{\mathcal{O}} \not\subseteq \text{QCMA}^{\mathcal{O}}$.*

Our separation, in addition to admitting a much simpler proof, has an additional benefit over the result of [BHNZ25]: our more structured oracle allows us to lift our result to provide the **first unconditional classical oracle separation** between BQP/qpoly from BQP/poly in a straightforward manner.¹

Theorem 1.2 (Informal). *There exists a classical oracle \mathcal{O} such that $\text{BQP}^{\mathcal{O}}/\text{qpoly} \not\subseteq \text{BQP}^{\mathcal{O}}/\text{poly}$.*

In fact, our oracle separates $\text{NP}^{\mathcal{O}} \cap \text{coNP}^{\mathcal{O}} \cap \text{BQP}^{\mathcal{O}}/\text{qpoly}$ from $\text{BQP}^{\mathcal{O}}/\text{poly}$, indicating a structural difference between our oracle and those based on spectral Forrelation [Zha24, BHNZ25] or expander mixing [Lut11, NN24, LMY25], which appear to originate from QMA-complete problems.

2 Technical Overview

The Yamakawa-Zhandry Algorithm. Our separation begins with the code intersection problem [YZ24]: given a code $C \subseteq \Sigma^n = (\mathbb{F}_q^s)^n$, function $H : [n] \times \Sigma \rightarrow \{0, 1\}$, and hash $x \in \{0, 1\}^n$, find a codeword $c \in C$ such that $H(i, c_i) = x_i$ for all $i \in [n]$ (we will use the shorthand $H(c) = x$ to refer to this constraint). Yamakawa and Zhandry show that when given oracle access to H , this problem has an efficient quantum algorithm but no (uniform) classical ones, giving a relativized separation between FP and FBQP. As our result will require modifying the Yamakawa-Zhandry algorithm, we begin by briefly explaining how it works.

We begin by noting that it suffices to be able to produce the state

$$|\psi\rangle \propto \sum_{\substack{\mathbf{v} \in C \\ H(\mathbf{v})=x}} |\mathbf{v}\rangle = \sum_{\mathbf{v} \in \Sigma^n} \mathbb{1}_{H,x}(\mathbf{v}) \cdot \mathbb{1}_C(\mathbf{v}) |\mathbf{v}\rangle,$$

where $\mathbb{1}_{H,x}(\cdot)$ and $\mathbb{1}_C(\cdot)$ are indicator functions for the event $H(\mathbf{v}) = x$ and for the event $\mathbf{v} \in C$, respectively. Taking inspiration from Regev’s reduction from SIS to LWE [Reg09], Yamakawa and Zhandry observe that $|\psi\rangle$ is the pointwise product of the states

$$|\phi_1\rangle := |\phi_1(x)\rangle \propto \sum_{\substack{\mathbf{v} \in \Sigma^n \\ H(\mathbf{v})=x}} |\mathbf{v}\rangle = \bigotimes_{i=1}^n \sum_{\substack{v_i \in \Sigma \\ H(i, v_i)=x_i}} |v_i\rangle \quad \text{and} \quad |\phi_2\rangle \propto \sum_{\mathbf{v} \in C} |\mathbf{v}\rangle,$$

both of which can be prepared efficiently given access to H . By the convolution theorem, we know that

$$\text{QFT}_q |\psi\rangle \propto \text{QFT}_q (|\phi_1\rangle \odot |\phi_2\rangle) = \text{QFT}_q |\phi_1\rangle \star \text{QFT}_q |\phi_2\rangle,$$

where \odot denotes the point-wise product of two vectors, and \star their convolution. Therefore, it suffices to efficiently prepare the state $|\text{goal}\rangle := \text{QFT}_q |\phi_1\rangle \star \text{QFT}_q |\phi_2\rangle$, as $\text{QFT}_q^{-1} |\text{goal}\rangle = |\psi\rangle$. If C is a \mathbb{F}_q -linear code, then $\text{QFT}_q |\phi_2\rangle$ is simply the uniform superposition over the dual code C^\perp , so we can produce the states

$$\text{QFT}_q |\phi_1\rangle \otimes \text{QFT}_q |\phi_2\rangle \otimes |0\rangle \propto \sum_{\substack{\mathbf{e} \in \Sigma^n \\ \mathbf{v} \in C^\perp}} \sqrt{\mathcal{D}_{H,x}(\mathbf{e})} |\mathbf{e}\rangle |\mathbf{v}\rangle |0\rangle$$

¹A similar idea appeared in [LLPY23], where a variant of the Yamakawa-Zhandry problem was used to give a separation between quantum and classical advice for algorithms which are only allowed classical access to all oracles. As we will see later, our separation strictly improves on this result since we rule out all BQP/poly algorithms with *quantum* oracle access while only requiring a single classical query given quantum advice.

$$\begin{aligned}
U_{\text{add}} &\mapsto \sum_{\substack{\mathbf{e} \in \Sigma^n \\ \mathbf{v} \in C^\perp}} \sqrt{\mathcal{D}_{H,x}(\mathbf{e})} |\mathbf{e}\rangle |\mathbf{v}\rangle |\mathbf{v} + \mathbf{e}\rangle \\
U_{\text{sub}} &\mapsto \sum_{\substack{\mathbf{e} \in \Sigma^n \\ \mathbf{v} \in C^\perp}} \sqrt{\mathcal{D}_{H,x}(\mathbf{e})} |0\rangle |\mathbf{v}\rangle |\mathbf{v} + \mathbf{e}\rangle,
\end{aligned}$$

while our desired state is

$$|0\rangle |0\rangle |\text{goal}\rangle = |0\rangle |0\rangle \otimes \text{QFT}_q |\phi_1\rangle \star \text{QFT}_q |\phi_2\rangle \propto \sum_{\substack{\mathbf{e} \in \Sigma^n \\ \mathbf{v} \in C^\perp}} \sqrt{\mathcal{D}_{H,x}(\mathbf{e})} |0\rangle |0\rangle |\mathbf{v} + \mathbf{e}\rangle,$$

where $\mathcal{D}_{H,x}(\cdot)$ is the density function of $\text{QFT}_q |\phi_1\rangle$ (here, we are ignoring phases for simplicity of exposition). For a completely random function H and a fixed x , we observe that $\mathcal{D}_{H,x}(\cdot)$ will have roughly half of its weight on 0 in each coordinate (since around half of all symbols in each coordinate should hash to 0 or 1), while the remaining half of its weight will be close to uniform over nonzero symbols. By taking C to be a folded Reed-Solomon (FRS) code with sufficiently high rate, Yamakawa and Zhandry show that one can efficiently decode C^\perp from errors over $\mathcal{D}_{H,x}(\cdot)$ with high probability and consequently prepare $|\text{goal}\rangle$ as desired.

Lifting the Separation. The obvious issue in using the code intersection problem to separate QMA from QCMA, however, is that both QMA and QCMA algorithms can make quantum oracle access to H , and our problem is already in BQP! Thus, we must make some modifications to the problem at hand.

First, as observed by [Liu23], the first phase of this algorithm can be made non-adaptive: the state $|\phi_2\rangle$ depends only on C , while $|\phi_1\rangle$ depends only mildly on x , since we can simply prepare all $2n$ preimage states

$$|H_1^{-1}(0)\rangle := \sum_{\substack{x \in \Sigma \\ H(1,x)=0}} |x\rangle, \dots, |H_n^{-1}(1)\rangle := \sum_{\substack{x \in \Sigma \\ H(n,x)=1}} |x\rangle,$$

before selecting $|\phi_1\rangle$ when given x . On the other hand, the second phase does not require access to H , and should work equally well for all x . Therefore, if we are given the state

$$|\text{adv}_H\rangle := \bigotimes_{i=1}^n |H_i^{-1}(0)\rangle \otimes \bigotimes_{i=1}^n |H_i^{-1}(1)\rangle \otimes |\phi_2\rangle$$

as our quantum proof or advice, we can produce $|\phi_1(x)\rangle := \bigotimes_{i=1}^n |H_i^{-1}(x_i)\rangle$ for *any* x and use it to produce a solution to the code intersection problem for x without access to H ! We can therefore replace H with the much weaker oracle $O_H(x, v)$ that simply verifies if the vector v hashes to x .

Of course, this problem is still in $\text{NP} \subseteq \text{QCMA}$, since for a fixed x , the prover can always send any codeword v that hashes to x ! Luckily, we claim that this is pretty much the *only* thing that the prover can do. To operationalize this intuition, we note that our quantum proof/advice is in some sense encoding many codewords along with their hashes in superposition. Thus, taking some inspiration from [LLPY23, BDK24], we define the *Code Intersection Subset Size* (CISS) problem as follows:

Estimate the size of a set $E \subseteq \{0, 1\}^n \times \Sigma^n$, under the promise that either $E = \{0, 1\}^n \times \Sigma^n$ or $|E| \leq t$ for some threshold $t \ll 2^n$, given access to the following oracle $O[H, E](x, v)$:

$$O[H, E](x, v) = \begin{cases} 1 & \text{if } v \in C, H(v) = x, \text{ and } (x, v) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

This problem is naturally in QMA: the prover can give as proof $|\text{adv}_H\rangle$ which depends only on H . Given $|\text{adv}_H\rangle$, the verifier can sample a random $x \in \{0, 1\}^n$ and should be able to produce with high probability $v \in C$ such that $H(v) = x$. The oracle O therefore allows the verifier to check if $(x, v) \in E$, making the set estimation problem trivial. In fact, the QMA verifier need not be concerned about malicious proofs, since a NO instance of O always outputs 0 whenever x (which is sampled solely by the verifier) is not in the support of E .

We note that our problem differs from the oracle problems defined in [LLPY23, BDK24], which take $E = F \times \Sigma^n$ for some small set F . This modification, while not impacting our QMA algorithm, will be crucial in establishing a QCMA lower bound.

An Entropic Viewpoint. To rule out QCMA proof systems, we first observe that the major difference between quantum and classical proofs lies in their *clonability*. In particular, an (oracle) algorithm which uses a *classical* witness can always be re-run with the same witness, even if it makes measurements. This simple and seemingly obvious fact, first formally identified in [Zha24], was utilized to great effect by [BHNZ25] to give their classical oracle separation, and we will take advantage of it as well.

To this end, suppose there was some Q -query QCMA verifier V which succeeded in the CISS problem. We observe that this means that V can always distinguish between $O[H, \{0, 1\}^n \times \Sigma^n]$ and $O[H, E]$ whenever E is small. But these oracles differ only at inputs (x, v) where $v \in C$, $H(v) = x$, and $(x, v) \notin E$! Thus, by the hybrid lemma [BBBV97], if we measure a random query that V makes to $O[H, E]$, we should expect to get a new pair $(x, v) \notin E$ such that $v \in C$ and $H(v) = x$ with good probability provided E is small.² This gives rise to a natural algorithm for *guessing* the hash values of codewords: starting with $E = \emptyset$, simulate a run of V with $O[H, E]$ and measure a random query before adding the measurement outcome to E ; rinse and repeat. By our previous argument, conditioned on having a good witness, each iteration of this algorithm should correctly produce a new codeword and hash with non-negligible probability. We can therefore turn V into a guesser with non-uniform advice which correctly produces many distinct codewords and their hash values without making *any* oracle queries. After guessing the classical w -bit witness, this gives rise to an unconditional no-query algorithm which guesses the hash values of ℓ codewords in C for all $\ell \leq t$ with probability

$$2^{-w} \cdot \left(\Omega(Q^{-2}) \right)^\ell = 2^{-\text{poly}(n)} \cdot \left(\frac{1}{\text{poly}(n)} \right)^\ell.$$

We now argue that this is in fact impossible. Observing that H is independently random at each coordinate i and symbol $x \in \Sigma$, we can upper bound the success probability of any sampler which produces ℓ points by $(\frac{1}{2})^{s(\ell)}$, where $s(\ell)$ is the minimum number of symbols that appear among ℓ distinct codewords in C . Taking $\ell = t = \omega(\text{poly}(n))$, if we can argue that $s(\ell) = \omega(\log n \cdot \ell)$, then we see that

$$2^{-\text{poly}(n)} \cdot \left(\frac{1}{\text{poly}(n)} \right)^\ell = \left(\frac{1}{\text{poly}(n)} \right)^\ell \gg \left(\frac{1}{2^{\omega(\log n)}} \right)^\ell = \left(\frac{1}{2} \right)^{s(\ell)},$$

which will give us our desired contradiction.

List Recovery and Code Expansion. How might we bound $s(\ell)$? For any ℓ distinct codewords $c_1, \dots, c_\ell \in C$, define the lists $S_1, \dots, S_n \subseteq \Sigma$ such that S_i consists of all symbols in Σ that appear in the i 'th coordinate of some codeword c_j for $j \in [\ell]$. Clearly, $s(\ell) = \min_{c_1, \dots, c_\ell} \sum_{i=1}^n |S_i|$, so there must be lists S_1^*, \dots, S_n^* such that

$$|C \cap (S_1^* \times \dots \times S_n^*)| \geq \ell \quad \text{and} \quad \sum_{i=1}^n |S_i^*| = s(\ell).$$

We now see that the question of how small $s(\ell)$ can be is precisely characterized by the *list-recoverability* of C . In particular, if we know that C is $(L, O(L))$ -list-recoverable for $L \lesssim \ell$, then this would mean that $\frac{s(\ell)}{n} = \frac{1}{n} \sum_{i=1}^n |S_i^*| = \Omega(\ell)$ as desired!

Note that this is a pretty strong condition; it necessitates the use of codes that have *near-optimal list recovery*, a property that in particular is not satisfied by the FRS codes used by [YZ24] (see Subsection §3.3 for further discussion). Fortunately, there is a fix: the setting of zero-error list recovery is closely linked to the notion of unbalanced expanders, and the recent work of [KTS22] shows that multiplicity codes exhibit precisely the sort of list recovery that we need. Moreover, the fact that multiplicity codes are \mathbb{F}_q -linear and that their duals have relatively good distance [RZVW24] should guarantee the success of our QMA algorithm. We note that although our dual code happens to admit efficient unique decoding, our separation only needs C^\perp to be *statistically* uniquely decodable, since we can always provide an (inefficient) decoding oracle.

A Final Complication. It seems that this rather simple argument completes our separation; after all, by switching to using multiplicity codes (rather than folded Reed-Solomon codes as in [YZ24, LLPY23, BDK24]), we have been able to rule out all possible QCMA algorithms. Sadly, we have to deal with one final and rather subtle issue, which has to

²Critically, if we have already successfully guessed some collection E of code words and hash values, we can always perfectly simulate $O[H, E]$. The same is *not* true for the oracles of [LLPY23, BDK24], where even “small” sets correspond to exponentially many codewords.

do with the parameters of the multiplicity codes: in the process of obtaining excellent list recovery/expansion from our multiplicity codes, we are forced to make the relative rate of our (primary) code sub-constant, which means our dual code now has sub-constant relative distance! Recalling that our error distribution should concentrate on vectors with Hamming weight roughly $n/2$, we observe that this level of noise is now likely intolerable as there may not even exist a unique decoding most of the time under this error distribution.

Our solution is relatively simple, and it uses the generous amount of flexibility that the [Reg09, YZ24] algorithm affords us. Instead of using a completely random function H , we will instead make our function *biased* in favor of 0 (reminiscent of recent strategies employed by [GGJL25, GGKM] in the context of communication complexity and low-depth implementations of the Yamakawa-Zhandry algorithm). That is, for each element $\sigma \in \Sigma$, $H(i, \sigma)$ will take on the value 0 with probability $p \gg \frac{1}{2}$. Thus, for $x = 0^n$, we can expect the error vectors in QFT $|\phi_1\rangle$ to have Hamming weight $\approx n(1-p)$, drastically reducing the amount of noise that we are required to decode from with respect to C^\perp .

This change does not come for free, however: unlike in [GGJL25], where the goal was to invert only $x = 0^n$, we need to be able to invert many x 's, including those with large Hamming weight. By biasing H towards 0, on inputs like $x = 1^n$, we create an error distribution which has expected Hamming weight $\approx np \gg n/2$, thereby *worsening* our ability to invert!

Our final idea is to observe that since our algorithm can only invert low Hamming-weight vectors x rather than all vectors in $\{0, 1\}^n$, we can simply *modify the problem to enforce this condition*. Instead of trying to invert all $x \in \{0, 1\}^n$, we can focus on inverting vectors of the form $x\|0^{n-n^c}$, where $x \in \{0, 1\}^{n^c}$ and $0 < c \ll 1$. That is, we will now try to differentiate $E = \{0, 1\}^{n^c} \times 0^{n-n^c} \times \Sigma^n$ from $E \subseteq F \times 0^{n-n^c} \times \Sigma^n$ where $|F| \leq t \ll 2^{n^c}$. Since $x\|0^{n-n^c}$ has Hamming weight at most n^c , the corresponding error distribution for $|\phi_1(x\|0^{n-n^c})\rangle$ will be concentrated on vectors with Hamming weight at most $n^c + (n - n^c)(1-p) \approx n(1-p)$, guaranteeing the success of our QMA algorithm provided our dual code has distance at least $O(n(1-p))$ and $p \lesssim 1 - 1/n^{1-c}$.

A General Recipe for a Classical Oracle Separation. Before describing the advice separation, we summarize all of the steps we have taken so far to provide a general recipe for getting a QMA versus QCMA oracle separation. We start with an infinite family of linear codes $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ over a large alphabet Σ (so that there are many solutions to the code intersection problem) such that:

1. Codewords of C_λ consist of $n = \text{poly}(\lambda)$ many symbols from Σ ,
2. C_λ^\perp can be (efficiently or combinatorially) decoded from up to $\Omega(\lambda^{1+c})$ errors for some constant c ,
3. C_λ has near optimal list recovery for sufficiently large λ ; i.e. $s(\ell) = \Omega(n \cdot \ell)$ for $\ell = \lambda^{\omega(1)}$.

Then for every λ , we can sample hash functions H_1, \dots, H_n to be biased so that roughly a $\lambda/n(\lambda)$ -fraction of symbols are pre-images of 1. We will ask for pre-images of $x\|0^{n-\lambda}$ for $x \in \{0, 1\}^\lambda$. The bias of the H_i , together with the fact that we ask for a hash that has at most λ many 1's, ensures that the dual decoding problem encounters an error with $O(\lambda)$ Hamming weight with high probability, which falls under our dual decoding distance of $\Omega(\lambda^{1+c})$. Thus, the Yamakawa-Zhandry algorithm works and the problem stays in QMA. At the same time, a QCMA verifier will imply a sampler that outputs v codewords of the code with probability $\text{poly}(\lambda)^{-v}$, and list recovery will enforce that this corresponds to $\Omega(v \cdot n)$ symbols. The bias of the H_i will mean that the probability of guessing all symbols correctly will be $(1 - \lambda/n)^{\Omega(v \cdot n)} \approx \exp(-\lambda v) \ll \text{poly}(\lambda)^{-v}$, giving us a contradiction.

By instantiating this recipe with multiplicity codes, we arrive at our classical oracle separation.

Moving to the Advice Setting: BQP/qpoly vs. BQP/poly. As mentioned previously, a separation in the proof setting does not immediately imply an advice-style separation. In fact, the CISS problem is easy with *trusted* advice: a single classical bit suffices to indicate whether the set in question is large or small.

The general paradigm we follow begins with an idea of [AK07]: after sampling a random binary language $\mathcal{L} \subseteq \{0, 1\}^n$, the [AK07] oracle essentially outputs whether an instance x is in the language if it is also given as input a specific quantum state $|\psi\rangle$. Intuitively, if it is hard to find $|\psi\rangle$ given classical advice, then it should be difficult to decide \mathcal{L} . In this manner, [AK07] transform a hard quantum search problem into a hard decision problem (relative to a quantum oracle). The classical oracle separation of [BHNZ25], being a kind of ‘‘classicalization’’ of the Aaronson-Kuperberg oracle, has a similar quantum search problem associated with it. However, since a classical oracle can no longer directly check the answer to a quantum search problem, obtaining a classical oracle separation between BQP/qpoly and BQP/poly seems to necessitate the use of a hard *classical search problem* instead.

Here, the relative simplicity of our separation and our use of a classical search/TFNP problem allows us to straightforwardly extend our results to the question of BQP/qpoly and BQP/poly. In contrast, it is not at all obvious (to us) how to construct even a candidate separating language based on the (decisional) spectral Forrelation problem of [BHNZ25].

To be more specific, the code intersection problem gives rise to the following BQP/qpoly language: begin by sampling some random binary language $\mathcal{L} \subseteq \{0, 1\}^n$. On input $x \in \{0, 1\}^n$, our oracle O will return whether $x \in \mathcal{L}$ or $x \notin \mathcal{L}$ provided it is also given a codeword v which hashes to x . On the one hand, our original QMA algorithm still works as a BQP/qpoly machine, since $|\text{adv}_H\rangle$ is agnostic of x and allows us to produce some v for any x . On the other hand, as we have argued earlier, even trusted classical advice should not help a BQP machine to produce many valid codewords, so for most x , any BQP machine with classical advice will not be able to receive the output of the oracle indicating whether $x \in \mathcal{L}$. Since the advice is bounded, it cannot itself describe many elements of \mathcal{L} , so the BQP machine will fail to decide whether $x \in \mathcal{L}$ for most x . The problem of guessing the value of a random function H given bounded-size advice and without querying H can be made precise by appealing to results on Yao’s box problem due to Chung, Guo, Liu and Qian [CGLQ20]. We show that by modifying ideas used by [LLPY23] for analyzing *classical-access* BQP/poly algorithms, we can reduce a decisional advice separation to a search-based separation for bona fide BQP/qpoly machines. Finally, by adapting the techniques we used to separate QMA from QCMA, we can also rule out the existence of BQP/poly machines which solve the code intersection search problem, completing the advice separation.

3 Discussion and Open Questions

Before describing the technical details of our construction and proof, we take some time to discuss parts of our result that might be interesting to a reader. One question a reader might have is why this result, and similar results, did not appear sooner. One of the main ideas of recent works on separating QMA from QCMA has been to start from a refined reduction from a QCMA algorithm to a sampler. The idea of using sampling to separate QMA from QCMA, although not new to this work, nor [BHNZ25], only started appearing in recent works [NZ23, Zha24]. Before this recent line of work, the common strategy for starting a separation was to take the “most popular witness” approach outlined by [AK07], which is used in [FK15, NN24, BDK24, LMY25]. Thus, one might expect that it would have been difficult to find this oracle separation prior to that. However, the related idea of reducing bounds on advice to bounds on a “multi-instance” variant of some problem has previously appeared in literature on the auxiliary-input random oracle model (AI-ROM), in particular being used to rule out advice proofs for function inversion [CGLQ20], and appeared in works as early as Aaronson’s result on the power of one-way quantum communication [Aar04].

These auxiliary-input random oracle model proofs bear some resemblance to the sampler idea. For example, if a T -query algorithm \mathcal{A} inverts a function f using an S -bit *classical* witness w with probability δ , then it is not hard to see that there is a gT -query algorithm \mathcal{B} which outputs preimages for g random images of f with no advice and probability $2^{-S} \cdot \delta^g$. The approach is (more or less) the one taken by [Aar04, CGLQ20] for classical advice. Clearly, such an approach suffices for classical advice states, since such algorithms can always be rewound. Here, the no-cloning theorem appears to present a barrier for quantum advice, but there is a workaround: using multiple copies of advice, the adversary’s success probability can be boosted to $1 - \text{negl}(\lambda)$, allowing its output to be copied and the algorithm to be rewound without causing damage to the advice state by the gentle measurement lemma. Thus, the quantum witness can be re-used *without* cloning it, and so the same upper bound should apply.

So why does such an argument not actually apply to our situation (and therefore fail to produce a separation)? A more careful examination reveals that the approach with quantum advice relies on the assumption that the aforementioned measurement produces a unique output with high probability! This is fine if one only needs to determine *if* an adversary has produced a preimage of a function, for example, but not *what* that preimage is. For example, an adversary could use its advice state to produce a superposition over many preimages before eventually measuring to produce one of exponentially many possible preimages (which is in fact what our QMA verifier does!). Then, although the measurement which verifies if the adversary can invert is gentle, the measurement which outputs a preimage as well is not. Thus, relative to the works on the AI-ROM, the main conceptual shift in recent sampler results is to rely on natural measurements which are inherently “destructive”. This turns out to be enough to force an algorithm solving a multi-instance problem to clone their advice state, thereby distinguishing between quantum and classical advice. Readers familiar with previous literature may also note that unlike works in the AI-ROM, our reduction (as well as the reduction of [BHNZ25]) produces a multi-instance adversary which makes *no* oracle queries.³

³Technically speaking, the reduction of [BHNZ25] begins with an adversary \mathcal{A}^U whose description depends on an oracle U and makes queries

Beyond the sampler idea, one might ask whether other approaches to separating QMA from QCMA using a code intersection oracle might have been successful. We emphasize that while the code intersection problem has previously been considered in the context of proof and advice separations [Liu23, LLPY23, BDK24], all existing works employed FRS codes, which (as mentioned earlier) do not enjoy strong enough list-recovery properties. In particular, running our argument with FRS codes gives an upper bound on the sampling success probability which is far too weak to derive *any* meaningful conclusions! On the other hand, multiplicity codes do not appear to have strong enough decoding properties to handle the large amounts of noise that would be incurred by perfectly random functions H , which necessitates our use of biased oracles and restriction to low Hamming-weight vectors (an idea which did not appear in earlier, more limited, separations).

Other remarks. As in [BHNZ25], it is not hard to see that our oracles also separate the clonable variants of QMA and BQP/qpoly from their regular counterparts.

Finally, although we primarily employed biased random oracles to make decoding possible for low-distance codes, we note that (as first observed by [GGKM]) biasing H has the additional benefit of making the Yamakawa-Zhandry algorithm *much more efficient*. Concretely, if we tweak H to have bias $p = 2/3$ or $p = 3/4$, this already decreases the noise level sufficiently that we can rely on (significantly faster and simpler) *unique decoders* rather than list decoders, which currently appears to be the major bottleneck in runtime.

We conclude this section with a discussion of some broader open questions left open by this work.

3.1 Structure versus randomness in classical oracle separations

As stated in the introduction, one major barrier in lifting the oracle separation of [BHNZ25] to the advice setting is that the oracle is most naturally associated with a hard quantum search problem, instead of a classical one. In their classical oracle separation, the authors identify a way of sampling instances of the spectral Forrelation problem such that the pair of functions that seems completely random, except that they are related by the Fourier transform (and, in the case of [BHNZ25], one of the two oracles being sparse). In some informal sense, our oracle separation enables an efficient quantum verifier to extract more information (namely, the solution to a hard search problem) from its witness, but doing so seems to require additional structure in the oracle.

A natural question to ask is how much information can be encoded into a quantum state before it becomes clonable, and whether our ideas are useful in encoding information into other kinds of quantum states. The cryptographic analogy of a separation between QMA and QCMA (or, really UnclonableQMA) is a primitive called quantum money [Aar09]. These are states that can not be cloned, but can be verified, similar to witnesses for QMA problems that are not in QCMA. An extremal form of this primitive (and this idea of encoding information in a quantum state) is known as “copy-protected software”, wherein an efficient quantum party can extract the input-output behavior of an entire classical function from a quantum state. Currently, there are several candidates for quantum money constructions in the plain model [FGH⁺12, BNZ25, Zha25], but less is known about copy-protected software. We hope that ideas from our separation might be useful in finding such constructions. To make progress towards cryptographic instantiations, one concrete direction to explore is a precise characterization of which witnesses cause our QMA verifier to accept with high probability.

3.2 QMA-completeness of a decoding problem

One interesting difference between our oracle separation and the oracle separation of [BHNZ25] is that their oracle separation can be seen as an obfuscation of a QMA-complete problem. To elaborate further, just as problems involving random sparse functions might model the difficulty of constructing a SAT solver which ignores the structure of the SAT instance it receives, the spectral Forrelation problem models an algorithm for solving a two-basis local Hamiltonian problem that does not look at the structure of the two local Hamiltonians it receives. Note that despite the connection to a QMA-complete problem, this property is not actually needed to achieve a separation between QMA and QCMA, as highlighted in the actual oracle separation of [BHNZ25]. In particular, the “YES” and “NO” instances can be taken to have sets S of a fixed size ℓ or $\leq \ell/10$. Such *distributions* of oracles can be easily distinguished by an AM protocol [GS86], but it appears unlikely that solving the actual spectral Forrelation problem can be done in AM, because another way to sample “NO” instances of the spectral Forrelation problem would be to re-sample sets S' independent of U . Distinguishing such pairs from Forrelated pairs (S, U) seems to truly require a QMA verifier.

to an oracle S , and produces an adversary \mathcal{B}^U which no longer queries S but still depends in some manner on the oracle U .

In contrast, the problem we construct really is in AM, as it directly involves distinguishing between large and small sets, and any reasonable variant of the code intersection problem would likely remain in the polynomial hierarchy. Of course, an oracle separation between QCMA and AM already exists (and in fact, comes from early work on separating QMA from QCMA [FK15]!), but we find it intriguing that our separation relies on a problem which appears to be of only intermediate difficulty and does not need the “full” power of QMA in some sense. In fact, as our problem appears on its face to be completely unrelated to quantum algorithms, we believe it remains an extremely interesting question to find a QMA-complete variant of the code intersection problem.

3.3 Simplifications to the separation

Naturally, one may ask if our separations can be made *even simpler*; here, we outline a few directions to consider.

On Round-Reduction Arguments. The work of [BDK24] can be thought of as a round-reduction argument as follows: in the style of [AK07], we will, given any QCMA proof system, fix some classical witness which corresponds to the largest set of NO instances. The polynomial bound on the size of the proof means that the collection of functions H which are consistent with this witness remains substantial, and in particular the distribution over all consistent H must have large min-entropy. We conclude that the set of symbols S which have low entropy conditioned on this classical proof is also bounded.

We now look at the verifier’s inputs to the oracle in the first round of queries. Observe that any inputs corresponding to non-codewords, as well as ones which have low overlap with S will be correct with negligible probability. By the hybrid lemma, it suffices to restrict our attention to codewords which have high overlap with S (so called “dangerous” inputs) – but this set is bounded precisely by the list recovery of our code! At this stage, we can simply give away the values of H on dangerous inputs *for free*, allowing the verifier to simulate its first round of queries. As a consequence, we can peel off a round of queries at the cost of requiring more advice. Repeating this “peeling and bloating” routine with FRS codes results in a $o(\log n / \log \log n)$ -round bound.

Sadly, this approach seems fairly doomed if we stick to the Yamakawa-Zhandry problem: even with nearly optimal list recovery, our advice will certainly increase by some constant factor with each round, so we would remain stuck at an $o(n)$ -round bound. One could imagine a better argument that uses some special property of any collection of recovered codewords that ensures these lists do not grow by too much iteratively, but this seems quite difficult (and would definitely be much more complicated!).

Other Codes? One might wonder why we need multiplicity codes here – after all, there are other codes which give rise to unbalanced expanders besides those of [KTS22], namely the constructions of [GUV09], but examining these constructions in closer detail presents some unexpected issues:

1. [GUV09] construct two unbalanced expanders with near-optimal expansion, one based on Parvaresh-Vardy codes and the other based on a subcode of FRS codes. Alas, neither of these instantiations are linear, which means our quantum proof/advice-based algorithm will fail.
2. [GUV09] also considers unbalanced expanders based on plain FRS codes, which *are* linear in some sense, but such expanders have expansion which is too weak to show meaningful classical hardness.

On the Necessity of Near-Optimal List Recovery. Our proof relies on fairly strong expansion/list recovery properties of the underlying code C . It is not hard to see that we can tolerate slightly suboptimal expansion, i.e. if $s(\ell) = \Omega(\ell/\text{poly}(n))$. Extending our sampling-based argument to work for polynomial or super-polynomial expansion, i.e. $s(\ell) = \ell^{O(1)}$, would open the door to a much larger class of usable codes.

What About Random Linear Codes? It is well known that random linear codes (RLCs) and their duals have good distance with high probability, and they are by definition linear, which means our quantum proof/advice-based algorithm will succeed. However, to show classical hardness, we need near-optimal list recovery, which remains a challenging open problem to show for RLCs [LS25]. We believe that a proof of such a result for RLCs is the clearest way to conceptually simplify our separation.

4 Preliminaries

4.1 Notation

We say that a function $\delta : \mathbb{N} \mapsto [0, 1]$ is inverse polynomial if there exists a polynomial p such that $\delta(n) \leq 1/p(n)$ for sufficiently large n . A function $\epsilon : \mathbb{N} \mapsto [0, 1]$ is negligible if for every polynomial p , for all sufficiently large n , $\epsilon(n) < 1/p(n)$.

We use the notation id to denote the identity operator. We will occasionally concatenate superscripts when it is clear from context, so $\text{QFT}^{-1, \otimes n}$ denotes $(\text{QFT}^{-1})^{\otimes n}$. We will also sometimes abbreviate the tensor product state $|0\rangle^{\otimes n}$ as $|0^n\rangle$.

A register R is a named finite-dimensional complex Hilbert space. If A, B, C are registers, for example, then the concatenation ABC denotes the tensor product of the associated Hilbert spaces. For a linear transformation L and register R , we write L_R to indicate that L acts on R , and similarly we write σ_R to indicate that a state σ is in the register R .

4.2 Probability and Complexity Theory

Definition 4.1 (Modified from [BHNZ25]). *We use the phrase “an oracle” to refer to a function $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$. A quantum query algorithm is a quantum circuit that interacts with an oracle $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$ via a query gate $|x, b\rangle \rightarrow |x, b \oplus \mathcal{O}(x)\rangle$, which acts on an $|x| + 1$ -qubit query register for various x . The algorithm is described by an alternating sequence of unitaries (drawn from any fixed gate set) and query gates. After all gates are applied, a designated qubit is measured in the standard basis to determine acceptance. The circuit has ancilla qubits which are initialized to $|0\rangle$ and all intermediate unitaries may act on an arbitrary (but finite) number of qubits.*

A quantum query algorithm may also receive an auxiliary witness or advice as input. A quantum witness/advice state is a state $|\psi\rangle$ on some (finite) number of qubits, while a classical witness/advice string is a (finite) bitstring w , treated as a computational basis state. The algorithm’s acceptance probability may depend on both the oracle and the witness.

Definition 4.2 ([BHNZ25]). *We denote a family of quantum oracle circuits/algorithms by $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, where the index λ corresponds to the length of the explicit input to the computational problem. $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ is P-uniform if there exists a deterministic polynomial-time Turing machine M that, on input 1^λ , outputs a full classical description of the circuit \mathcal{A}_λ . The runtime of M implies that \mathcal{A}_λ has at most $\text{poly}(\lambda)$ gates (oracle or elementary), queries \mathcal{O} at lengths of at most $\text{poly}(\lambda)$, and receives witnesses/advice of length at most $\text{poly}(\lambda)$.*

We can now define the complexity classes that we will consider in this work.

Definition 4.3 (Oracle QCMA). *A promise language $\mathcal{L}^\mathcal{O} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})^\mathcal{O} \subseteq \{0, 1\}^*$ is in $\text{QCMA}^\mathcal{O}$ if there exists a P-uniform family of quantum oracle circuits \mathcal{A}_λ with \mathcal{A}_λ accepting a witness of length $t(\lambda)$, such that for every input x of length $\lambda = |x|$,*

- $x \in \mathcal{L}_{\text{yes}} \implies \exists w \in \{0, 1\}^{t(\lambda)} \text{ s.t. } \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, w) = 1] \geq \frac{2}{3},$
- $x \in \mathcal{L}_{\text{no}} \implies \forall \tilde{w} \in \{0, 1\}^{t(\lambda)}, \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, \tilde{w}) = 1] \leq \frac{1}{3}.$

Definition 4.4 (Oracle QMA). *A promise language $\mathcal{L}^\mathcal{O} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})^\mathcal{O} \subseteq \{0, 1\}^*$ is in $\text{QMA}^\mathcal{O}$ if there exists a P-uniform family of quantum oracle circuits \mathcal{A}_λ with \mathcal{A}_λ accepting a witness of length $t(\lambda)$, such that for every input x of length $\lambda = |x|$,*

- $x \in \mathcal{L}_{\text{yes}} \implies \exists |\psi\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)} \text{ s.t. } \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, |\psi\rangle) = 1] \geq \frac{2}{3},$
- $x \in \mathcal{L}_{\text{no}} \implies \forall |\tilde{\psi}\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)}, \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, |\tilde{\psi}\rangle) = 1] \leq \frac{1}{3}.$

Definition 4.5 (Oracle BQP/poly). *A promise language $\mathcal{L}^\mathcal{O} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})^\mathcal{O} \subseteq \{0, 1\}^*$ is in $\text{BQP}^\mathcal{O}/\text{poly}$ if there exists a P-uniform family of quantum oracle circuits \mathcal{A}_λ with \mathcal{A}_λ accepting advice of length $t(\lambda)$ and an advice family $\{\text{adv}_\lambda\}_{\lambda \geq 1}$ where $|\text{adv}_\lambda| = t(\lambda)$, such that for every input x of length $\lambda = |x|$,*

- $x \in \mathcal{L}_{\text{yes}} \implies \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, \text{adv}_\lambda) = 1] \geq \frac{2}{3},$

- $x \in \mathcal{L}_{\text{no}} \implies \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, \text{adv}_\lambda) = 1] \leq \frac{1}{3}$.

If $t(\lambda) = 0$ then we say $\mathcal{L}^\mathcal{O} \in \text{BQP}^\mathcal{O}$.

Definition 4.6 (Oracle BQP/qpoly). *A promise language $\mathcal{L}^\mathcal{O} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})^\mathcal{O} \subseteq \{0, 1\}^*$ is in $\text{BQP}^\mathcal{O}/\text{qpoly}$ if there exists a P-uniform family of quantum oracle circuits \mathcal{A}_λ with \mathcal{A}_λ accepting advice of length $t(\lambda)$ and an advice family $\{|\text{adv}_\lambda\rangle\}_{\lambda \geq 1}$ where $|\text{adv}_\lambda\rangle \in (\mathbb{C}^2)^{\otimes t(\lambda)}$, such that for every input x of length $\lambda = |x|$,*

- $x \in \mathcal{L}_{\text{yes}} \implies \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, |\text{adv}_\lambda\rangle) = 1] \geq \frac{2}{3}$,
- $x \in \mathcal{L}_{\text{no}} \implies \Pr[\mathcal{A}_\lambda^\mathcal{O}(x, |\text{adv}_\lambda\rangle) = 1] \leq \frac{1}{3}$.

Given a quantum query algorithm, we can define the query mass of the algorithm on a particular set of inputs.

Definition 4.7 (Query mass). *For an oracle circuit A making Q quantum queries to an oracle \mathcal{O} with input domain D , let $\sum_x \alpha_x^{(i)} |x\rangle |\psi_x\rangle$ be the state of the algorithm immediately before their i 'th query to \mathcal{O} , where the first register is the input register to the oracle, and let $M_x(i) = |\alpha_x^{(i)}|^2$ be the query mass of x in the i 'th query. For a subset $V \subseteq [Q] \times D$, let $M_V = \sum_{(i,x) \in V} M_x(i)$ be the total query mass of points in V .*

The following theorem was proven in [BBBV97], using the hybrid method.

Theorem 4.8 (Hybrid method [BBBV97]). *Let A be an oracle circuit which makes Q queries to an oracle \mathcal{O} with input domain D . If we modify \mathcal{O} into an oracle \mathcal{O}' which differs only on a set of time-input pairs $V \subseteq [Q] \times D$, then*

$$|\Pr[A^\mathcal{O}(\cdot) = 1] - \Pr[A^{\mathcal{O}'}(\cdot) = 1]| \leq 4\sqrt{QM_V}.$$

Finally, we will also use some basic probability lemmas.

Lemma 4.9 (Chernoff Bound). *Let X_1, \dots, X_n be independent random variables taking values in $\{0, 1\}$, $X := \sum_{i=1}^n X_i$, and $\mu := \mathbb{E}[X]$. For any $\delta \geq 0$, it holds that $\Pr[X \geq (1 + \delta)\mu] \leq e^{-\delta^2\mu/(2+\delta)}$.*

Lemma 4.10 (Borel–Cantelli, [Bor09, Can17]). *Let $\{X_\lambda\}_{\lambda \in \mathbb{N}}$ be a sequence of (not necessarily independent) random variables with values in $\{0, 1\}$. If $\sum_{\lambda=1}^\infty \mathbb{E}[X_\lambda] < \infty$, then $\Pr[\sum_{\lambda=1}^\infty X_\lambda = \infty] = 0$.*

4.3 Coding Theory

For a prime power q , we denote by \mathbb{F}_q the finite field of order q and denote by $(\mathbb{F}_q)_{<k}[X]$ the set of univariate polynomials over \mathbb{F}_q with degree less than k .

Definition 4.11. *A code of length $n \in \mathbb{N}$ over an alphabet Σ is a subset $C \subseteq \Sigma^n$. $C \subseteq \Sigma^n$ is said to be \mathbb{F}_q -linear if its alphabet $\Sigma = \mathbb{F}_q^s$ for some field \mathbb{F}_q and a positive integer $s \geq 1$ and C is an \mathbb{F}_q -linear subspace of Σ^n . Equivalently, this means that for any two codewords $x, y \in C$ and scalar $\alpha \in \mathbb{F}_q$, both $x + y$ and $\alpha \cdot x$ are in C .*

For an \mathbb{F}_q -linear code C , the dual code of C is the code $C^\perp \subseteq (\mathbb{F}_q^s)^n$ containing all strings $c' \in (\mathbb{F}_q^s)^n$ which satisfy

$$\sum_{i=1}^n \sum_{j=1}^s (c'_i)_j \cdot (c_i)_j = 0$$

for all $c \in C$. Observe that C^\perp is always \mathbb{F}_q -linear, and that $|C| \cdot |C^\perp| = |\Sigma|^n$ if and only if C is \mathbb{F}_q -linear.

For any vector $x \in \Sigma^n$, define $\text{hw}(x) \in [0, n]$ as the Hamming weight of x , i.e. the number of nonzero elements in x . We say that $C \subseteq \Sigma^n$ has distance d if for any two distinct codewords $c_1, c_2 \in C$, $\text{hw}(c_1 - c_2) \geq d$.

Definition 4.12 (Formal and Hasse derivatives). *Let $f(X) = \sum_{j=0}^n a_j X^j \in \mathbb{F}_q[X]$ be a univariate polynomial over \mathbb{F}_q . We define the i 'th formal and Hasse derivatives of $f(X)$ as the linear operators which take $f(X)$ to the polynomials*

$$f^{[i]}(X) = \sum_{j=i}^n \frac{j!}{(j-i)!} a_j X^{j-i} \quad \text{and} \quad f^{(i)}(X) = \sum_{j=i}^n \binom{j}{i} a_j X^{j-i},$$

respectively. Note that $f^{[i]}(X) = i! f^{(i)}(X)$ for all i and f .

Definition 4.13 (Univariate multiplicity codes, from [RT97, Nie01, KSY14]). Let \mathbb{F}_q be a finite field and let s be a positive integer. Let $\alpha_1, \dots, \alpha_n$ be distinct points in \mathbb{F}_q , and let $k < sn$ be a positive integer. The univariate multiplicity code $\text{Mult}_{s, \mathbb{F}_q}(\alpha_1, \dots, \alpha_n; k)$ is the code over the alphabet $\Sigma = \mathbb{F}_q^s$ of length n which associates each polynomial $f(X) \in (\mathbb{F}_q)_{<k}[X]$ to the codeword $c \in \Sigma^n$ such that for $i \in [n]$,

$$c_i = (f^{(0)}(\alpha_i), f^{(1)}(\alpha_i), \dots, f^{(s-1)}(\alpha_i)),$$

where $f^{(j)}$ is the j 'th Hasse derivative of f . Let $\text{Mult}_{s, \mathbb{F}_q, k} := \text{Mult}_{s, \mathbb{F}_q}(1, \dots, q; k)$; note that $\text{Mult}_{s, \mathbb{F}_q, k}$ is \mathbb{F}_q -linear.

Definition 4.14 (List recoverability). A code $C \subseteq \Sigma^n$ is (ℓ, L) -list recoverable if for all $S_1, \dots, S_n \subseteq \Sigma$ such that $\frac{1}{n} \sum_{i=1}^n |S_i| \leq \ell$,

$$|\{(x_1, \dots, x_n) \in C : \forall i \in [n], x_i \in S_i\}| \leq L.$$

Definition 4.15 (Expanders [GUV09]). A bipartite graph with N left-vertices, M right-vertices, and left-degree D is specified by a function $\Gamma : [N] \times [D] \rightarrow [M]$, where $\Gamma(x, y)$ denotes the y 'th neighbor of x . For a set $X \subseteq [N]$, we write $\Gamma(X)$ to denote its set of neighbors $\bigcup_{x \in X, y \in [D]} \Gamma(x, y)$. For a set $T \subseteq [M]$, we write $\text{LIST}_\Gamma(T) = \{x : \Gamma(x) \subseteq T\}$.

We say that Γ is a (K, A) -expander if for every set $X \subseteq [N]$ of size at most K , $|\Gamma(X)| \geq A \cdot |X|$. Note that if Γ is a (K, A) -expander then for all $B \leq K$ and all sets T such that $|T| < AB$, $|\text{LIST}_\Gamma(T)| < B$.

Our separation will utilize expanders based on multiplicity codes as constructed in [KTS22].⁴

Theorem 4.16 ([KTS22]). For every field \mathbb{F}_q , $k, s \in \mathbb{N}$ such that $15 \leq s+1 \leq k \leq \text{char}(\mathbb{F}_q)$, identify the elements of \mathbb{F}_q^k with univariate polynomials of degree less than k . Define the graph $\Gamma : \mathbb{F}_q^k \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{s+1}$ by

$$\Gamma(f, y) = (y, f^{[0]}(y), f^{[1]}(y), \dots, f^{[s-1]}(y)),$$

where $f^{[i]}$ is the i 'th formal derivative of f in $\mathbb{F}_q[X]$. For every $K > 0$, Γ is a (K, A) -expander where

$$A = q - \frac{k(s+1)}{2} \cdot (qK)^{\frac{1}{s+1}}.$$

Corollary 4.17. For each security parameter $\lambda \in \mathbb{N}$, let $k = \lambda^3$, $\lambda^5 < q \leq 2\lambda^5$ be any prime, and $s = \lambda$. Identify the elements of \mathbb{F}_q^k with univariate polynomials of degree less than k . Define the code $C'_\lambda \subseteq \Sigma^q = (\mathbb{F}_q^s)^q$ with encoding map

$$\text{Enc}(f) = \{(f^{[0]}(y), f^{[1]}(y), \dots, f^{[s-1]}(y))\}_{y \in \mathbb{F}_q},$$

where $f^{[i]}$ is the i 'th formal derivative of f . Then, for sufficiently large λ , C'_λ is $(\ell, 2\ell)$ -list recoverable if $\ell \leq 2^s = 2^\lambda$.

Proof. Set $K = 2^{s+1}$ and fix $\ell \leq 2^s$; Theorem 4.16 implies that $\Gamma : \mathbb{F}_q^k \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{s+1}$ is a (K, A) -expander where

$$A \geq q - \frac{k(s+1)}{2} \cdot (qK)^{\frac{1}{s+1}} = q - \frac{k(s+1)}{2} \cdot (q \cdot 2^{s+1})^{\frac{1}{s+1}} > q/2,$$

for sufficiently large λ .

Fix any lists $S_1, \dots, S_q \subseteq \Sigma$ such that $\frac{1}{q} \sum_{i=1}^q |S_i| \leq \ell$ and consider the set

$$R = \bigcup_{i=1}^q R_i, \quad R_i := \{(i, w_0, \dots, w_{s-1}) : (w_0, \dots, w_{s-1}) \in S_i\}.$$

Observe that R is a set of right vertices of Γ and that $|R| = \sum_i |S_i| \leq q\ell < A \cdot 2\ell \leq AK$. Now consider the set X of all polynomials f such that for all $i \in [q]$, $\text{Enc}(f)_i \in S_i$; our goal is to bound $|X|$. By construction, $X = \text{LIST}_\Gamma(R)$, so it follows that $|X| = |\text{LIST}_\Gamma(R)| < 2\ell$. \square

In our parameter regime, it is straightforward to see that the code $C_\lambda = \text{Mult}_{s, \mathbb{F}_q, k}$ has the same list-recoverability as C'_λ , since for fields \mathbb{F}_q where $\text{char}(\mathbb{F}_q) > s$, C_λ and C'_λ are identical up to scalar factors.

⁴We note that the proof of expansion extends straightforwardly to subgraphs of Γ defined by taking edges corresponding to subsets $S \subseteq \mathbb{F}_q$, although it suffices for us to take $S = \mathbb{F}_q$.

Corollary 4.18. For each security parameter $\lambda \in \mathbb{N}$, let $k = \lambda^3$, $\lambda^5 < q \leq 2\lambda^5$ be any prime, and $s = \lambda$. Then, for sufficiently large λ , $C_\lambda = \text{Mult}_{s, \mathbb{F}_q, k}$ is $(\ell, 2\ell)$ -list recoverable for all $\ell \leq 2^\lambda$.

Finally, we will use a result about duals of univariate multiplicity codes which we reprove in Section §A.⁵

Theorem 4.19 ([RZVW24]). For all parameters s, q , and $k < sq$, $(\text{Mult}_{s, \mathbb{F}_q, k})^\perp$ has distance at least $\frac{k+1}{s}$.

4.4 Yao's Box Problem and Non-Uniform Advice

We will need the following results on non-uniform advice for our separation between BQP/qpoly and BQP/poly.

Theorem 4.20 ([CGLQ20]). Let $G : [N] \rightarrow \{0, 1\}$ be a random function. Let \mathcal{A} be an unbounded-time algorithm, with S bits of classical advice z_G . For an index $x \in [N]$, let $G^{|x} : N \rightarrow \{0, 1\}$ denote the function that results from removing x from G ; in other words, on inputs $x' \neq x$, $G(x') = G^{|x}(x')$ and $G^{|x}(x) = 0$. The probability that \mathcal{A} computes $G(x)$ while making Q quantum queries to $G^{|x}$ for a random index x is at most

$$\Pr_{G,x} [\mathcal{A}^{G^{|x}}(z_G, x) = G(x)] \leq \frac{1}{2} + O\left(\frac{(S + \log N)Q}{N}\right)^{1/3}.$$

Lemma 4.21. Let $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ be a uniformly random function. For an algorithm \mathcal{A} that makes $Q(\lambda) = \text{poly}(\lambda)$ quantum queries to G and a family of $t(\lambda) = \text{poly}(\lambda)$ -bit classical advice $\{z_G\}_G$, suppose that

$$\Pr_{G,x \leftarrow \{0,1\}^\lambda} [\mathcal{A}^G(z_G, x) = G(x)] > \frac{3}{5}.$$

Then, for sufficiently large λ , for a $\frac{1}{4000Q^2}$ fraction of $x \in \{0, 1\}^\lambda$, measuring a random query of \mathcal{A}^G (for randomly sampled G) will produce x with probability at least $\frac{1}{3200Q^2}$.

Proof. The proof closely follows [LLPY23], but considers quantum queries instead of classical ones. The only way for \mathcal{A} to distinguish G from $G^{|x}$ is to have nontrivial query mass at x . Denote by $M_{G,x}$ the total query mass that x is placed by \mathcal{A} when querying G . For each G and x we have that

$$|\Pr[\mathcal{A}^G(z_G, x) = G(x)] - \Pr[\mathcal{A}^{G^{|x}}(z_G, x) = G(x)]| \leq 4\sqrt{QM_{G,x}}.$$

Now we consider the case when we uniform randomly choose $x \leftarrow \{0, 1\}^\lambda$, and require $\mathcal{A}^{G^{|x}}(z_G, x)$ to output $G(x)$. This is exactly Yao's box problem, so by Theorem 4.20,

$$\begin{aligned} \Pr_{G,x} [\mathcal{A}^{G^{|x}}(z_G, x) = G(x)] &\leq \frac{1}{2} + O\left(\frac{(t + \lambda)Q}{2^\lambda}\right)^{1/3} = \frac{1}{2} + \text{negl}(\lambda) \\ \implies \Pr_{G,x} [\mathcal{A}^G(z_G, x) = G(x)] - \Pr_{G,x} [\mathcal{A}^{G^{|x}}(z_G, x) = G(x)] &\geq \frac{1}{10} - \text{negl}(\lambda). \end{aligned}$$

Thus,

$$\begin{aligned} \mathbb{E}_{G,x} [M_{G,x}] &\geq \mathbb{E}_{G,x} \left[\frac{1}{16Q} (\Pr[\mathcal{A}^G(z_G, x) = G(x)] - \Pr[\mathcal{A}^{G^{|x}}(z_G, x) = G(x)])^2 \right] \\ &\geq \frac{1}{16Q} \left(\mathbb{E}_{G,x} [\Pr[\mathcal{A}^G(z_G, x) = G(x)] - \Pr[\mathcal{A}^{G^{|x}}(z_G, x) = G(x)]] \right)^2 \geq \frac{1}{1600Q} - \text{negl}(\lambda), \end{aligned}$$

by Jensen's inequality. Finally, by a Markov inequality, we see that

$$\Pr_x \left[\mathbb{E}_G [M_{G,x}] \geq \frac{1}{3200Q} \right] \geq \frac{1}{3200Q^2} - \text{negl}(\lambda).$$

Thus, for sufficiently large λ , for a $\frac{1}{3200Q^2} - \text{negl}(\lambda) \geq \frac{1}{4000Q^2}$ fraction of $x \in \{0, 1\}^\lambda$, measuring a random query of \mathcal{A}^G (for randomly sampled G) will produce x with probability at least $\frac{1}{3200Q^2}$. \square

⁵We note that although the work of [RZVW24] was recently retracted, the particular theorem we use remains correct. For completeness, an entirely self-contained proof of this fact is given in the appendix.

5 The Generalized Code Intersection Problem

5.1 Definitions and Basic Facts

We begin by recalling the definitions and basic results from [YZ24]. Much of this section will be taken directly from [YZ24], with only minor modifications. We first define the code intersection relation, which is essentially the problem of finding codewords over n symbols whose symbols have a particular hash value.

Definition 5.1 (Code intersection relation, adapted from [YZ24, LLPY23, BDK24]). *For a function $H : [n] \times \Sigma \rightarrow \{0, 1\}$ and a code $C \subseteq \Sigma^n$, define the code intersection relation $R_{C,H} \subseteq \{0, 1\}^n \times \Sigma^n$ by*

$$R_{C,H} = \{(\mathbf{x}, \mathbf{v}) = (x_1, \dots, x_n, v_1, \dots, v_n) : ((v_1, \dots, v_n) \in C) \wedge (\forall i \in [n], H(i, v_i) = x_i)\}.$$

Remark 5.2. *We can view $H : [n] \times \Sigma \rightarrow \{0, 1\}$ as a collection of n many functions, $H(i, \cdot) : \Sigma \rightarrow \{0, 1\}$, and we will at times use the notation $H_i : \Sigma \rightarrow \{0, 1\}$ when referring to the function corresponding to the i 'th output coordinate of H .*

Definition 5.3 (Trace over a finite field [YZ24]). *For any prime power $q = r^m$ where r is prime, we define the trace function $\text{Tr}(x) := \sum_{i=0}^{m-1} x^{r^i}$ which maps elements of \mathbb{F}_q to \mathbb{F}_r . The trace function is \mathbb{F}_r -linear: for all $a, b \in \mathbb{F}_r$ and $x, y \in \mathbb{F}_q$, $\text{Tr}(ax + by) = a \text{Tr}(x) + b \text{Tr}(y)$. In addition, for any $x \in \mathbb{F}_q$, $\sum_{y \in \mathbb{F}_q} \omega_r^{\text{Tr}(x \cdot y)} = 0$, where $\omega_r := e^{2\pi i/r}$.*

Definition 5.4 (Quantum Fourier transform over a finite field [YZ24]). *For a finite field \mathbb{F}_q where $q = r^m$ and r is prime, the quantum Fourier transform over \mathbb{F}_q is the unitary denoted by QFT_q such that for any $x \in \mathbb{F}_q$*

$$\text{QFT}_q |x\rangle = \frac{1}{\sqrt{q}} \sum_{z \in \mathbb{F}_q} \omega_r^{\text{Tr}(x \cdot z)} |z\rangle.$$

The QFT over an alphabet $\Sigma = \mathbb{F}_q^s$ is the s -wise tensor product of QFT_q : for $\mathbf{x} = (x_1, \dots, x_s) \in \Sigma$,

$$\text{QFT}_\Sigma |\mathbf{x}\rangle := \text{QFT}_q^{\otimes s} |x_1\rangle \dots |x_s\rangle = \frac{1}{\sqrt{|\Sigma|}} \sum_{\mathbf{z} \in \Sigma} \omega_r^{\text{Tr}(\mathbf{x} \cdot \mathbf{z})} |\mathbf{z}\rangle.$$

Similarly, for any positive integer n and $\mathbf{x} \in \Sigma^n$, we have

$$\text{QFT}_\Sigma^{\otimes n} |\mathbf{x}\rangle = \frac{1}{|\Sigma|^{n/2}} \sum_{\mathbf{z} \in \Sigma^n} \omega_r^{\text{Tr}(\mathbf{x} \cdot \mathbf{z})} |\mathbf{z}\rangle.$$

The unitary QFT_q can be approximated within error ε in operator norm in time $\text{poly}(\log q, \log 1/\varepsilon)$ [CW02, vDHI06].

Definition 5.5 (Fourier transform of a function [YZ24]). *For functions $f, g : \Sigma^n \rightarrow \mathbb{C}$, we define*

$$\hat{f}(\mathbf{z}) := \frac{1}{|\Sigma|^{n/2}} \sum_{\mathbf{x} \in \Sigma^n} f(\mathbf{x}) \omega_r^{\text{Tr}(\mathbf{x} \cdot \mathbf{z})}, \quad (f \cdot g)(\mathbf{x}) := f(\mathbf{x}) \cdot g(\mathbf{x}), \quad \text{and} \quad (f \star g)(\mathbf{x}) := \sum_{\mathbf{y} \in \Sigma^n} f(\mathbf{y}) \cdot g(\mathbf{x} - \mathbf{y}).$$

Note that

$$\text{QFT}_\Sigma^{\otimes n} \sum_{\mathbf{x} \in \Sigma^n} f(\mathbf{x}) |\mathbf{x}\rangle = \sum_{\mathbf{z} \in \Sigma^n} \hat{f}(\mathbf{z}) |\mathbf{z}\rangle.$$

Fact 5.6 ([YZ24]). *The following properties hold for the Fourier transform:*

1. (Parseval's equality) *For all functions $f : \Sigma^n \rightarrow \mathbb{C}$, $\sum_{\mathbf{x} \in \Sigma^n} |f(\mathbf{x})|^2 = \sum_{\mathbf{z} \in \Sigma^n} |\hat{f}(\mathbf{z})|^2$.*
2. (Pointwise transform) *Suppose that we have $f_i : \Sigma \rightarrow \mathbb{C}$ for $i \in [n]$ and $f : \Sigma^n \rightarrow \mathbb{C}$ is defined by $f(\mathbf{x}) := \prod_{i=1}^n f_i(x_i)$. Then $\hat{f}(\mathbf{z}) = \prod_{i=1}^n \hat{f}_i(z_i)$.*
3. (Convolution theorem) *For all functions $f, g, h : \Sigma^n \rightarrow \mathbb{C}$, $\widehat{f \cdot g} = \frac{1}{|\Sigma|^{n/2}} (\hat{f} \star \hat{g})$, $\widehat{f \star g} = |\Sigma|^{n/2} (\hat{f} \cdot \hat{g})$, and $f \cdot \widehat{(g \star h)} = (\hat{f} \star (\hat{g} \cdot \hat{h}))$.*

Lemma 5.7 (Fourier transform of a linear code). *Let $C \subseteq \Sigma^n = (\mathbb{F}_q^s)^n$ be any \mathbb{F}_q -linear code. Then,*

$$f(\mathbf{u}) = \begin{cases} \frac{1}{\sqrt{|C|}} & \text{if } \mathbf{u} \in C, \\ 0 & \text{otherwise.} \end{cases} \iff \hat{f}(\mathbf{u}) = \begin{cases} \frac{1}{\sqrt{|C^\perp|}} & \text{if } \mathbf{u} \in C^\perp, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Since C is \mathbb{F}_q -linear, $|C| \cdot |C^\perp| = |\Sigma|^n$, and thus for any $\mathbf{z} \in C^\perp$,

$$\hat{f}(\mathbf{z}) = \frac{1}{|\Sigma|^{n/2}} \sum_{\mathbf{u} \in \Sigma^n} f(\mathbf{u}) \omega_r^{\text{Tr}(\mathbf{u} \cdot \mathbf{z})} = \frac{1}{|\Sigma|^{n/2}} \sum_{\mathbf{u} \in C} \frac{1}{\sqrt{|C|}} = \frac{1}{\sqrt{|C^\perp|}}.$$

Finally, $\hat{f}(\mathbf{z}) = 0$ for $\mathbf{z} \notin C^\perp$ by Parseval's equality. The reverse direction follows by an identical argument. \square

Lemma 5.8. *Let $|\psi\rangle, |\phi\rangle$ be states such that $\| |\psi\rangle \| = 1$ and $\| |\psi\rangle - |\phi\rangle \| \leq \varepsilon < 1$. Then, $\left\| |\psi\rangle - \frac{|\phi\rangle}{\| |\phi\rangle \|} \right\| \leq 2\varepsilon$.*

Proof. By the reverse/inverse triangle inequality,

$$\| |\phi\rangle \| - 1 = \| |\phi\rangle \| - \| |\psi\rangle \| \leq \| |\phi\rangle - |\psi\rangle \| = \varepsilon \implies \| |\phi\rangle \| \geq 1 - \varepsilon > 0.$$

We can thus define the normalized state $|\phi'\rangle := \frac{|\phi\rangle}{\| |\phi\rangle \|}$. By the regular triangle inequality,

$$\| |\psi\rangle - |\phi'\rangle \| \leq \| |\psi\rangle - |\phi\rangle \| + \| |\phi\rangle - |\phi'\rangle \| \leq \varepsilon + \| |\phi\rangle \| - 1 \leq 2\varepsilon. \quad \square$$

Lemma 5.9 ([BV93]). *Let $|\psi\rangle, |\phi\rangle$ be states such that $\| |\psi\rangle \| = \| |\phi\rangle \| = 1$ and $\| |\psi\rangle - |\phi\rangle \| \leq \varepsilon$. Then the total variation distance between the probability distributions resulting from measurements of $|\phi\rangle$ and $|\psi\rangle$ is at most 4ε .*

Now we state the main algorithmic result of [YZ24], namely that a quantum algorithm can approximately implement the convolution trick for some families of functions.

Lemma 5.10 ([YZ24]). *Let $|\psi\rangle$ and $|\phi\rangle$ be quantum states on a quantum system over an alphabet $\Sigma = \mathbb{F}_q^s$ written as*

$$|\psi\rangle = \sum_{\mathbf{u} \in \Sigma^n} V(\mathbf{u}) |\mathbf{u}\rangle \quad \text{and} \quad |\phi\rangle = \sum_{\mathbf{e} \in \Sigma^n} W(\mathbf{e}) |\mathbf{e}\rangle,$$

for functions $V, W : \Sigma^n \rightarrow \mathbb{C}$. Let $F : \Sigma^n \rightarrow \Sigma^n$ be a function and let $\text{GOOD} \subseteq \Sigma^n \times \Sigma^n$ be a subset such that for any $(\mathbf{u}, \mathbf{e}) \in \text{GOOD}$, we have $F(\mathbf{u} + \mathbf{e}) = \mathbf{u}$. Define $\text{BAD} := (\Sigma^n \times \Sigma^n) \setminus \text{GOOD}$. Suppose that

$$\sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}} |\widehat{V}(\mathbf{u}) \widehat{W}(\mathbf{e})|^2 \leq \varepsilon \quad \text{and} \quad \sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}: \mathbf{u} + \mathbf{e} = \mathbf{z}} \widehat{V}(\mathbf{u}) \widehat{W}(\mathbf{e}) \right|^2 \leq \delta.$$

Define the unitaries U_{add} and U_F as follows:

$$U_{\text{add}} = \sum_{(\mathbf{u}, \mathbf{e}) \in \Sigma^n \times \Sigma^n} |\mathbf{u} + \mathbf{e}, \mathbf{u}\rangle \langle \mathbf{e}, \mathbf{u}| \quad \text{and} \quad U_F = \sum_{(\mathbf{u}, \mathbf{w}) \in \Sigma^n \times \Sigma^n} |\mathbf{w}, \mathbf{u} - F(\mathbf{w})\rangle \langle \mathbf{w}, \mathbf{u}|.$$

Then,

$$\left\| \left(\text{QFT}_{\Sigma}^{-1, \otimes n} \otimes \text{id} \right) \cdot U_F \cdot U_{\text{add}} \cdot \left(\text{QFT}_{\Sigma}^{\otimes n} \otimes \text{QFT}_{\Sigma}^{\otimes n} \right) |\phi\rangle |\psi\rangle - |\Sigma|^{n/2} \sum_{\mathbf{z} \in \Sigma^n} (V(\mathbf{z}) \cdot W(\mathbf{z})) |\mathbf{z}\rangle |0\rangle \right\| \leq \sqrt{\varepsilon} + \sqrt{\delta}.$$

Lemmas 5.8 to 5.10 imply the following corollary (as the QFT, U_{add} , and U_F are unitaries).

Corollary 5.11. *Let V and W be functions, and ε and δ be the corresponding error parameters from Lemma 5.10. For any property $\mathcal{P} : \Sigma^n \rightarrow \{0, 1\}$, if $\| |\phi\rangle \| = \| |\psi\rangle \| = 1$ and measuring (in the standard basis) the normalization of*

$$|\Sigma|^{n/2} \sum_{\mathbf{z} \in \Sigma^n} (V(\mathbf{z}) \cdot W(\mathbf{z})) |\mathbf{z}\rangle |0\rangle,$$

produces an outcome \mathbf{z} such that $\mathcal{P}(\mathbf{z}) = 1$ with probability p , then measuring (in the standard basis)

$$\left(\text{QFT}_{\Sigma}^{-1, \otimes n} \otimes \text{id} \right) \cdot U_F \cdot U_{\text{add}} \cdot \left(\text{QFT}_{\Sigma}^{\otimes n} \otimes \text{QFT}_{\Sigma}^{\otimes n} \right) |\phi\rangle |\psi\rangle,$$

produces an outcome \mathbf{z} such that $\mathcal{P}(\mathbf{z}) = 1$ with probability at least $p - 8(\sqrt{\varepsilon} + \sqrt{\delta})$.

5.2 Technical Lemmas

In this section, we state some technical lemmas to extend the Yamakawa-Zhandry algorithm to work with biased oracles. We first introduce a pair of functions V and W that represent normalized indicators for a code C and the preimages of any output \mathbf{b} for a function H . For any \mathbb{F}_q -linear code $C \subseteq \Sigma^n = (\mathbb{F}_q^s)^n$, function $H : [n] \times \Sigma \rightarrow \{0, 1\}$, and string $\mathbf{b} \in \{0, 1\}^n$, let $V : \Sigma^n \rightarrow \mathbb{C}$, $W_i^{H_i, b_i} : \Sigma \rightarrow \mathbb{C}$, and $W^{H, \mathbf{b}} : \Sigma^n \rightarrow \mathbb{C}$ be defined as follows:

$$V(\mathbf{u}) = \begin{cases} \frac{1}{\sqrt{|C|}} & \mathbf{u} \in C \\ 0 & \text{otherwise} \end{cases},$$

$$W_i^{H_i, b_i}(e) = \begin{cases} \frac{1}{\sqrt{|T_i^{H_i, b_i}|}} & e \in T_i^{H_i, b_i} \\ 0 & \text{otherwise} \end{cases},$$

$$W^{H, \mathbf{b}}(e_1, \dots, e_n) = \prod_{i=1}^n W_i^{H_i, b_i}(e_i),$$

where $T_i^{H_i, b_i} \subseteq \Sigma$ is the subset consisting of $e_i \in \Sigma$ such that $H_i(e_i) = b_i$.

Definition 5.12 (p -biased distribution). *For any $p \in [0, 1]$ and set Σ , let $\text{Bias}_{p, \Sigma}$ denote the distribution over functions from Σ to $\{0, 1\}$ that samples $F : \Sigma \rightarrow \{0, 1\}$ with probability $p^{|F^{-1}(1)|}(1-p)^{|F^{-1}(0)|}$. Let $\text{Bias}_{n, p, \Sigma}$ denote the distribution over functions $G : [n] \times \Sigma \rightarrow \{0, 1\}$ that samples G with probability $\text{Bias}_{n, p, \Sigma}(G) := \prod_{i=1}^n \text{Bias}_{p, \Sigma}(G_i)$.*

The following claim follows immediately from the definition of $\text{Bias}_{n, p, \Sigma}$.

Claim 5.13. *Let π be any permutation over Σ (resp. Σ^n). Then, the distributions $\text{Bias}_{p, \Sigma}$ and $\text{Bias}_{p, \Sigma} \circ \pi$ (resp. $\text{Bias}_{n, p, \Sigma}$ and $\text{Bias}_{n, p, \Sigma} \circ \pi$) are identical.*

The following lemma shows that when we take the Fourier transform of the preimage state of H_i sampled from $\text{Bias}_{p, \Sigma}$, the resulting Fourier coefficients are in expectation uniform over all non-zero elements of Σ and fixed (either p or $1-p$ depending on if we are taking the preimage of 0 or 1) on 0.

Claim 5.14. *Fix any string $\mathbf{b} \in \{0, 1\}^n$. For all $i \in [n]$ and $\sigma, \sigma' \in \Sigma \setminus \{0\}$, it holds that*

$$\mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(0)|^2] = \begin{cases} p & \text{if } b_i = 1 \\ 1-p & \text{if } b_i = 0 \end{cases} \quad \text{and} \quad \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(\sigma)|^2] = \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(\sigma')|^2].$$

Proof. We can directly compute the expected Fourier weight on 0 as follows:

$$\mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(0)|^2] = \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i^{H_i, b_i}(z) \right|^2 \right] = \frac{\mathbb{E}_{H_i} [|\widehat{T}_i^{H_i, b_i}|]}{|\Sigma|} = \begin{cases} p & \text{if } b_i = 1 \\ 1-p & \text{if } b_i = 0 \end{cases}.$$

Since $\sigma \neq 0$ (resp. $\sigma' \neq 0$), for any $w \in \mathbb{F}_q$, the number of $z \in \Sigma$ such that $\sigma \cdot z = w$ is $|\Sigma|/q$. Therefore, there is a permutation $\pi_{\sigma, \sigma'} : \Sigma \rightarrow \Sigma$ such that $\sigma \cdot z = \sigma' \cdot \pi_{\sigma, \sigma'}(z)$ for all $z \in \Sigma$. Thus, by Claim 5.13,

$$\begin{aligned} \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(\sigma)|^2] &= \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i^{H_i, b_i}(z) \cdot \omega_r^{\text{Tr}(\sigma \cdot z)} \right|^2 \right] \\ &= \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i^{H_i \circ \pi_{\sigma, \sigma'}^{-1}, b_i}(\pi_{\sigma, \sigma'}(z)) \cdot \omega_r^{\text{Tr}(\sigma' \cdot \pi_{\sigma, \sigma'}(z))} \right|^2 \right] \\ &= \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i^{H_i \circ \pi_{\sigma, \sigma'}^{-1}, b_i}(z) \cdot \omega_r^{\text{Tr}(\sigma' \cdot z)} \right|^2 \right] \\ &= \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i^{H_i, b_i}(z) \cdot \omega_r^{\text{Tr}(\sigma' \cdot z)} \right|^2 \right] \\ &= \mathbb{E}_{H_i \leftarrow \text{Bias}_{p, \Sigma}} [|\widehat{W}_i^{H_i, b_i}(\sigma')|^2]. \quad \square \end{aligned}$$

For any function $\text{Dec}_{C^\perp} : \Sigma^n \rightarrow \Sigma^n$, define the sets $\mathcal{G} := \{\mathbf{e} \in \Sigma^n : \forall \mathbf{u} \in C^\perp, \text{Dec}_{C^\perp}(\mathbf{u} + \mathbf{e}) = \mathbf{u}\}$, $\mathcal{B} := \Sigma^n \setminus \mathcal{G}$, $\text{GOOD} := C^\perp \times \mathcal{G}$ and $\text{BAD} := (\Sigma^n \times \Sigma^n) \setminus \text{GOOD}$. By construction, $\text{Dec}_{C^\perp}(\mathbf{u} + \mathbf{e}) = \mathbf{u}$ for all $(\mathbf{u}, \mathbf{e}) \in \text{GOOD}$. Applying the definition of \widehat{V} from Lemma 5.7, we see that for all $\mathbf{b} \in \{0, 1\}^n$ and functions H ,

$$\sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}} |\widehat{V}(\mathbf{u}) \widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2 = \sum_{\mathbf{u} \in C^\perp} \sum_{\mathbf{e} \in \mathcal{B}} \left| \frac{1}{\sqrt{|C^\perp|}} \widehat{W}^{H, \mathbf{b}}(\mathbf{e}) \right|^2 = \sum_{\mathbf{e} \in \mathcal{B}} |\widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2.$$

We can apply the same logic to get that

$$\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}, \mathbf{u} + \mathbf{e} = \mathbf{z}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H, \mathbf{b}}(\mathbf{e}) \right|^2 = \sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H, \mathbf{b}}(\mathbf{e}) \right|^2.$$

Definition 5.15. Let $\mathcal{D}_{p, b}$ be the distribution over Σ that takes 0 with probability $1 - p$ if $b = 0$ (resp. probability p if $b = 1$) and otherwise takes a uniformly random element of $\Sigma \setminus \{0\}$. For any bitstring $\mathbf{b} \in \{0, 1\}^n$, define the distribution $\mathcal{D}_{p, \mathbf{b}}$ over Σ^n to be the Cartesian product of the distributions \mathcal{D}_{p, b_i} .

Definition 5.16. Fix any \mathbb{F}_q -linear code $C \subseteq \Sigma^n$, function $\text{Dec}_{C^\perp} : \Sigma^n \rightarrow \Sigma^n$, set $S \subseteq \{0, 1\}^n$, $p \in [0, 1]$, and real-valued function $\mu : \mathbb{N} \rightarrow \mathbb{R}$. $(C, \text{Dec}_{C^\perp})$ is said to be (p, μ, S) -good if for all $\lambda \in \mathbb{N}$ and $\mathbf{b} \in S$, $\Pr_{\mathbf{e} \leftarrow \mathcal{D}_{p, \mathbf{b}}}[\mathbf{e} \in \mathcal{B}] \leq \mu(\lambda)$.

Lemma 5.17. Suppose that $(C, \text{Dec}_{C^\perp})$ is (p, μ, S) -good. Then, for all $\lambda \in \mathbb{N}$ and $\mathbf{b} \in S$,

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[\sum_{\mathbf{e} \in \mathcal{B}} |\widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2 \right] \leq \mu(\lambda).$$

Proof. By Claim 5.14, $\mathcal{D}_{p, b_i}(e_i) = \mathbb{E}_{H_i} [|\widehat{W}_i^{H_i, b_i}(e_i)|^2]$ for all $e_i \in \Sigma$ where (slightly abusing notation) $\mathcal{D}_{p, b_i}(\cdot)$ is the density function of the distribution \mathcal{D}_{p, b_i} . Moreover, for any $\mathbf{e} = (e_1, \dots, e_n) \in \Sigma^n$, string \mathbf{b} and function H , since $W^{H, \mathbf{b}}(\mathbf{e}) = \prod_{i=1}^n W_i^{H_i, b_i}(e_i)$, by Fact 5.6, we have that $\widehat{W}^{H, \mathbf{b}}(\mathbf{e}) = \prod_{i=1}^n \widehat{W}_i^{H_i, b_i}(e_i)$. Thus, $\mathcal{D}_{p, \mathbf{b}}(\mathbf{e}) = \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} [|\widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2]$ for all $\mathbf{e} \in \Sigma^n$. By linearity of expectation, we see that for all $\lambda \in \mathbb{N}$ and $\mathbf{b} \in S$,

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[\sum_{\mathbf{e} \in \mathcal{B}} |\widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2 \right] = \sum_{\mathbf{e} \in \mathcal{B}} \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[|\widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2 \right] = \sum_{\mathbf{e} \in \mathcal{B}} \mathcal{D}_{p, \mathbf{b}}(\mathbf{e}) = \Pr_{\mathbf{e} \leftarrow \mathcal{D}_{p, \mathbf{b}}}[\mathbf{e} \in \mathcal{B}] \leq \mu(\lambda). \quad \square$$

We define the function $B : \Sigma^n \rightarrow \mathbb{C}$ to be the inverse Fourier transform of the indicator function $\widehat{B}(\mathbf{e}) = \mathbf{1}_{\mathbf{e} \in \mathcal{B}}$.

Claim 5.18. Suppose that $(C, \text{Dec}_{C^\perp})$ is (p, μ, S) -good. Then for all $\lambda \in \mathbb{N}$, $\mathbf{b} \in S$, and $\mathbf{z} \in \Sigma^n$,

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[|(B \star W^{H, \mathbf{b}})(\mathbf{z})|^2 \right] \leq \mu(\lambda).$$

Proof. For $\mathbf{z}_0, \mathbf{z}_1 \in \Sigma^n$, define the permutation $\pi_{\mathbf{z}_0, \mathbf{z}_1} : \Sigma^n \rightarrow \Sigma^n$ as $\pi_{\mathbf{z}_0, \mathbf{z}_1}(\mathbf{z}) := \mathbf{z} + \mathbf{z}_0 - \mathbf{z}_1$. By Claim 5.13,

$$\begin{aligned} \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[|(B \star W^{H, \mathbf{b}})(\mathbf{z}_0)|^2 \right] &= \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[\left| \sum_{\mathbf{x} \in \Sigma^n} B(\mathbf{x}) \cdot W^{H, \mathbf{b}}(\mathbf{z}_0 - \mathbf{x}) \right|^2 \right] \\ &= \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[\left| \sum_{\mathbf{x} \in \Sigma^n} B(\mathbf{x}) \cdot W^{H \circ \pi_{\mathbf{z}_0, \mathbf{z}_1}, \mathbf{b}}(\mathbf{z}_1 - \mathbf{x}) \right|^2 \right] \\ &= \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[\left| \sum_{\mathbf{x} \in \Sigma^n} B(\mathbf{x}) \cdot W^{H, \mathbf{b}}(\mathbf{z}_1 - \mathbf{x}) \right|^2 \right] \\ &= \mathbb{E}_{H \leftarrow \text{Bias}_{n, p, \Sigma}} \left[|(B \star W^{H, \mathbf{b}})(\mathbf{z}_1)|^2 \right]. \end{aligned}$$

Thus, by Lemma 5.17, we have that for all $\lambda \in \mathbb{N}$, $\mathbf{b} \in S$, and $\mathbf{z} \in \Sigma^n$,

$$\begin{aligned}
\mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[|(B \star W^{H,\mathbf{b}})(\mathbf{z})|^2 \right] &= \frac{1}{|\Sigma|^n} \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} |(B \star W^{H,\mathbf{b}})(\mathbf{z})|^2 \right] \\
&= \frac{1}{|\Sigma|^n} \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} |\Sigma|^{n/2} \widehat{B}(\mathbf{z}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{z})|^2 \right] \\
&= \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} |\widehat{B}(\mathbf{z}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{z})|^2 \right] \\
&= \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \mathcal{B}} |\widehat{W}^{H,\mathbf{b}}(\mathbf{z})|^2 \right] \leq \mu(\lambda). \quad \square
\end{aligned}$$

Claim 5.19. For any function H and string $\mathbf{b} \in \{0,1\}^n$, it holds that

$$\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) \right|^2 = \sum_{\mathbf{z} \in \Sigma^n} |(V \cdot (B \star W^{H,\mathbf{b}}))(\mathbf{z})|^2.$$

Proof. For any $\mathbf{z} \in \Sigma^n$, we use the fact that $V(\mathbf{x}) = 0$ for $\mathbf{x} \notin C^\perp$ to show that

$$\sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) = \sum_{\substack{\mathbf{u} \in \Sigma^n, \mathbf{e} \in \Sigma^n: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{B}(\mathbf{e}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) = (\widehat{V} \star (\widehat{B} \cdot \widehat{W}^{H,\mathbf{b}}))(\mathbf{z}) = (V \cdot (\widehat{B \star W}^{H,\mathbf{b}}))(\mathbf{z}).$$

The claim then follows from Parseval's equality. \square

Corollary 5.20. Suppose that $(C, \text{Dec}_{C^\perp})$ is (p, μ, S) -good. Then for all $\lambda \in \mathbb{N}$ and $\mathbf{b} \in S$,

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) \right|^2 \right] \leq \mu(\lambda).$$

Proof. By Claims 5.18 and 5.19, we have that for all $\lambda \in \mathbb{N}$ and $\mathbf{b} \in S$,

$$\begin{aligned}
\mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) \right|^2 \right] &= \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} |(V \cdot (B \star W^{H,\mathbf{b}}))(\mathbf{z})|^2 \right] \\
&= \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in C} \frac{1}{|C|} |(B \star W^{H,\mathbf{b}})(\mathbf{z})|^2 \right] \\
&= \frac{1}{|C|} \sum_{\mathbf{z} \in C} \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[|(B \star W^{H,\mathbf{b}})(\mathbf{z})|^2 \right] \\
&\leq \frac{1}{|C|} \sum_{\mathbf{z} \in C} \mu(\lambda) = \mu(\lambda). \quad \square
\end{aligned}$$

5.3 The Biased Yamakawa-Zhandry Algorithm

We are now ready to present our modified algorithm for handling biased oracles.

Definition 5.21 (The Yamakawa-Zhandry advice state). For a code $C \subseteq \Sigma^n$ and function $H : [n] \times \Sigma \rightarrow \{0,1\}$, define the sets $S_{i,b} := \{e \in \Sigma : H(i, e) = b\}$ for $(i, b) \in [n] \times \{0,1\}$. Let $|\phi_{i,b}\rangle$ and $|\psi\rangle$ denote the following states:

$$|\phi_{i,b}\rangle = \begin{cases} \frac{1}{\sqrt{|S_{i,b}|}} \sum_{e \in S_{i,b}} |e\rangle & \text{if } |S_{i,b}| \neq 0 \\ |\perp\rangle & \text{otherwise} \end{cases} \quad \text{and} \quad |\psi\rangle = \frac{1}{\sqrt{|C|}} \sum_{\mathbf{u} \in C} |\mathbf{u}\rangle.$$

We define the advice state for (C, H) , denoted by $|\text{adv}_{C,H}\rangle$ (or $|\text{adv}_H\rangle$ when the code C is implicit) as follows:

$$|\text{adv}_H\rangle := \begin{cases} (\bigotimes_{i=1}^n |\phi_{i,0}\rangle \otimes |\phi_{i,1}\rangle) \otimes |\psi\rangle & \text{if } |\phi_{i,b}\rangle \neq |\perp\rangle \text{ for all } (i,b) \in [n] \times \{0,1\}, \\ |\perp\rangle & \text{otherwise.} \end{cases}$$

Theorem 5.22. Fix any \mathbb{F}_q -linear code $C \subseteq \Sigma^n = (\mathbb{F}_q^s)^n$, function $\text{Dec}_{C^\perp} : \Sigma^n \rightarrow \Sigma^n$, set $S \subseteq \{0,1\}^n$, $p \in [0, 1/2]$, and function $\mu : \mathbb{N} \rightarrow [0, 1]$ such that $(C, \text{Dec}_{C^\perp})$ is (p, μ, S) -good. Let $\text{BiasedYZ} = \text{BiasedYZ}_{\text{Dec}_{C^\perp}}$ be the quantum algorithm described in Figure 1 and $|\text{adv}_H\rangle$ be the state described in Definition 5.21. Then, for all $\lambda \in \mathbb{N}$:

1. BiasedYZ runs in time $\text{poly}(n, s, \log q, \lambda) + T_{\text{Dec}}$, where T_{Dec} is the time required to compute Dec_{C^\perp} .
2. For all C and H , $|\text{adv}_H\rangle$ is a $O(sn \log q)$ -qubit state.
3. For all strings $\mathbf{b} \in S$,

$$\Pr_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\Pr_{\text{BiasedYZ}} [(\mathbf{b}, \mathbf{v}) \in R_{C,H} : \mathbf{v} \leftarrow \text{BiasedYZ}(|\text{adv}_H\rangle, \mathbf{b})] \geq 1 - 16\mu(\lambda)^{1/4} - 2^{-4\lambda} \right] \geq 1 - (2n(1-p)^{|\Sigma|} + 2\mu(\lambda)^{1/2}).$$

BiasedYZ_{Dec_{C[⊥]}} ($|\text{adv}_H\rangle, \mathbf{b}$):

1. If $|\text{adv}_H\rangle \neq |\perp\rangle$, construct the state $(\bigotimes_{i=1}^n |\phi_{i,\mathbf{b}_i}\rangle)_A \otimes |\psi\rangle_B$ by re-arranging $|\text{adv}_H\rangle$; else, return \perp .
2. Apply $(\text{QFT}_{\Sigma}^{\otimes n})_A \otimes (\text{QFT}_{\Sigma}^{\otimes n})_B$.
3. Controlled on register B, add the value of register B to register A.
4. Apply Dec_{C^\perp} to uncompute register B given the value of register A.
5. Apply $(\text{QFT}_{\Sigma}^{-1,\otimes n})_A \otimes \text{id}_B$ and measure register A to get an outcome $\mathbf{v} \in \Sigma^n$. Output \mathbf{v} .

Figure 1: Biased Yamakawa-Zhandry Algorithm $\text{BiasedYZ}_{\text{Dec}_{C^\perp}}(|\text{adv}_H\rangle, \mathbf{b})$

Proof. Throughout this proof, we assume that $2\mu(\lambda)^{1/4} < 1$, since the theorem holds trivially otherwise. Upon inspection, we see that $|\text{adv}_H\rangle$ is a $O(2n \log(|\Sigma|) + \log(|\Sigma|^n)) = O(sn \log q)$ -qubit state.

We now analyze the runtime of BiasedYZ. First, observe that step 1 runs in time $O(sn \log q)$. To implement steps 2 and 5, note that $\text{QFT}_{\Sigma}^{\otimes n}$ (and its inverse) can be implemented with total error at most $2^{-4\lambda}/4$ in $\text{poly}(n, s, \log q, \log 2^{4\lambda}) = \text{poly}(n, s, \log q, \lambda)$ time. Finally, step 3 consists of adding in Σ , which can be done in $\text{poly}(n, s, \log q)$ time, and step 4 takes time T_{Dec} by definition. In total, we have that BiasedYZ runs in time $O(sn \log q) + \text{poly}(n, s, \log q, \lambda) + \text{poly}(n, s, \log q) + T_{\text{Dec}} = \text{poly}(n, s, \log q, \lambda) + T_{\text{Dec}}$.

We finish by analyzing the correctness of BiasedYZ. For fixed $(i, b) \in [n] \times \{0,1\}$, $|\phi_{i,b}\rangle \neq |\perp\rangle$ with probability at least $1 - (1-p)^{|\Sigma|}$ (as long as at least one symbol hashes to b under $H(i, \cdot)$). By a union bound, $|\text{adv}_H\rangle \neq |\perp\rangle$ with probability at least $1 - 2n(1-p)^{|\Sigma|}$. Thus, using the definition of V and W from Section 5.2, we will assume for the remainder of the proof that $|\text{adv}_H\rangle = \sum_{\mathbf{u}, \mathbf{e} \in \Sigma^n} V(\mathbf{u})W(\mathbf{e})|\mathbf{u}\rangle|\mathbf{e}\rangle$. Then, by Lemma 5.17 and Corollary 5.20, we have that for all $\mathbf{b} \in S$,

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}} |\widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e})|^2 \right] = \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{e} \in \mathcal{B}} |\widehat{W}^{H,\mathbf{b}}(\mathbf{e})|^2 \right] \leq \mu(\lambda)$$

and

$$\mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{(\mathbf{u}, \mathbf{e}) \in \text{BAD}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) \right|^2 \right] = \mathbb{E}_{H \leftarrow \text{Bias}_{n,p,\Sigma}} \left[\sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{\substack{\mathbf{u} \in C^\perp, \mathbf{e} \in \mathcal{B}: \\ \mathbf{u} + \mathbf{e} = \mathbf{z}}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H,\mathbf{b}}(\mathbf{e}) \right|^2 \right] \leq \mu(\lambda).$$

Fixing $\mathbf{b} \in S$, Markov's inequality and the union bound implies that H satisfies both

$$\sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}} |\widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H, \mathbf{b}}(\mathbf{e})|^2 \leq \mu(\lambda)^{1/2} \quad \text{and} \quad \sum_{\mathbf{z} \in \Sigma^n} \left| \sum_{(\mathbf{u}, \mathbf{e}) \in \text{BAD}: \mathbf{u} + \mathbf{e} = \mathbf{z}} \widehat{V}(\mathbf{u}) \cdot \widehat{W}^{H, \mathbf{b}}(\mathbf{e}) \right|^2 \leq \mu(\lambda)^{1/2}$$

with probability at least $1 - 2\mu(\lambda)^{1/2}$ over $\text{Bias}_{n, p, \Sigma}$. For these H , by Lemma 5.10,

$$\left\| \left(\text{QFT}_{\Sigma}^{-1, \otimes n} \otimes \text{id} \right) \cdot U_{\text{Dec}_{C_{\lambda}^{\perp}}} \cdot U_{\text{add}} \cdot \left(\text{QFT}_{\Sigma}^{\otimes n} \otimes \text{QFT}_{\Sigma}^{\otimes n} \right) \bigotimes_{i=1}^n |\phi_{i, \mathbf{b}_i}\rangle \otimes |\psi\rangle - |\Sigma|^{n/2} \sum_{\mathbf{z} \in \Sigma^n} (V(\mathbf{z}) \cdot W^{H, \mathbf{b}}(\mathbf{z})) |\mathbf{z}\rangle |0\rangle \right\| \leq 2\mu(\lambda)^{1/4}.$$

Since $|\text{adv}_H\rangle \neq |\perp\rangle$ and $2\mu(\lambda)^{1/4} < 1$, the state $|\text{tgt}\rangle = |\Sigma|^{n/2} \sum_{\mathbf{z} \in \Sigma^n} (V(\mathbf{z}) \cdot W^{H, \mathbf{b}}(\mathbf{z})) |\mathbf{z}\rangle$ is nonzero. But measuring (the normalization of) $|\text{tgt}\rangle$ in the standard basis always produces vectors \mathbf{v} such that $(\mathbf{b}, \mathbf{v}) \in R_{C, H}$. Thus, by Corollary 5.11, the output of BiasedYZ will be a vector \mathbf{v} such that $(\mathbf{b}, \mathbf{v}) \in R_{C_{\lambda}, H}$ with probability at least $1 - 16\mu(\lambda)^{1/4} - 2^{-4\lambda}$ (where the $2^{-4\lambda}$ term comes from approximating the QFT).

We conclude that with probability $1 - (2n(1-p)^{|\Sigma|} + 2\mu(\lambda)^{1/2})$ over $\text{Bias}_{n, p, \Sigma}$, BiasedYZ succeeds with probability at least $1 - 16\mu(\lambda)^{1/4} - 2^{-4\lambda}$ for any given $\mathbf{b} \in S$, as desired. \square

Finally, we show that our choice of code satisfies the required conditions of Theorem 5.22.

Corollary 5.23. *For each security parameter $\lambda \in \mathbb{N}$, define the code $C_{\lambda} = \text{Mult}_{s, \mathbb{F}_q, k}$, where $\lambda^5 < q \leq 2\lambda^5$ is a prime, $k = \lambda^3$, and $s = \lambda$. Then, there exists an efficient/uniform quantum algorithm BiasYZ and a family of $\text{poly}(\lambda)$ -qubit states $\{|\text{adv}_H\rangle\}_H$ such that the following holds for sufficiently large λ : for all strings $x \in \{0, 1\}^{\lambda}$,*

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\Pr[(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda}, H} : \mathbf{v} \leftarrow \text{BiasYZ}(|\text{adv}_H\rangle, x)] \geq 1 - 2^{-\lambda} \right] \geq 1 - 2^{-2\lambda}.$$

Proof. We begin by arguing that there exists a deterministic algorithm $\text{Dec}_{C_{\lambda}^{\perp}}$ and function μ that for sufficiently large λ satisfies $\mu(\lambda) \leq 2^{-8\lambda}$, such that $T_{\text{Dec}} = \text{poly}(\lambda)$ and $(C_{\lambda}, \text{Dec}_{C_{\lambda}^{\perp}})$ is $(\frac{1}{\lambda^4}, \mu, \{0, 1\}^{\lambda} \times 0^{q-\lambda})$ -good. Define the subset $\mathcal{G} := \{\mathbf{e} \in \Sigma^q : \text{hw}(\mathbf{e}_{[\lambda+1:q]}) \leq \frac{11}{\lambda^4}(q-\lambda)\}$. By the Chernoff bound (Lemma 4.9), for any $x \in \{0, 1\}^{\lambda}$,

$$\Pr_{\mathbf{e} \leftarrow \mathcal{D}_{1/\lambda^4, x \| 0^{q-\lambda}}} [\mathbf{e} \notin \mathcal{G}] = \Pr_{\mathbf{e} \leftarrow \mathcal{D}_{1/\lambda^4, x \| 0^{q-\lambda}}} \left[\text{hw}(\mathbf{e}_{[\lambda+1:q]}) > 11 \cdot \frac{(q-\lambda)}{\lambda^4} \right] \leq e^{-\frac{100\lambda}{12} + 10} \leq 2^{-8\lambda},$$

assuming $\lambda \geq 10$, so it suffices to consider decoding errors $\mathbf{e} \in \mathcal{G}$. Theorem 4.19 implies that for all $\mathbf{e} \in \mathcal{G}$,

$$\text{hw}(\mathbf{e}) = \text{hw}(\mathbf{e}_{[1:\lambda]}) + \text{hw}(\mathbf{e}_{[\lambda+1:q]}) < \lambda + \frac{11}{\lambda^4}(q-\lambda) \leq 25\lambda \leq \lambda^2/2 \leq \text{dist}(C_{\lambda}^{\perp})/2,$$

assuming $\lambda \geq 50$. It therefore suffices to *uniquely* decode C_{λ}^{\perp} in a deterministic and efficient manner. We set μ to be the maximum probability (across all x) that unique decoding for C_{λ}^{\perp} fails with error distribution $\mathcal{D}_{1/\lambda^4, x \| 0^{q-\lambda}}$; by our previous argument, $\mu(\lambda) \leq 2^{-8\lambda}$ as long as $\lambda \geq 50$.

By Theorem A.5, we know that $C_{\lambda}^{\perp} = \text{GM}_{s, \mathbb{F}_q}(U_1, \dots, U_q; 1, \dots, q; sq - k)$. From Lemma A.3, it is easy to see that $A_{i,j}(X)$ can be computed in $\text{poly}(s, q) = \text{poly}(\lambda)$ time, and thus $a_{i,j}$ (and consequently U_i and U_i^{-1}) can be computed in $\text{poly}(\lambda)$ time. Finally, it remains to efficiently uniquely decode $\text{Mult}_{s, \mathbb{F}_q, sq-k}$, which we can do deterministically in $\text{poly}(s, q) = \text{poly}(\lambda)$ time [Nie01, KSY14, Kop15].⁶

Applying Theorem 5.22 gives a family of $\text{poly}(\lambda)$ -qubit states $\{|\text{adv}_H\rangle\}_H$ and a $\text{poly}(\lambda)$ -time uniform algorithm \mathcal{B} such that for $\lambda \geq 50$ and all strings $x \in \{0, 1\}^{\lambda}$,

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \Sigma}} \left[\Pr[(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda}, H} : \mathbf{v} \leftarrow \mathcal{B}(|\text{adv}_H\rangle, x \| 0^{q-\lambda})] \geq 1 - 2^{-2\lambda+4} - 2^{-4\lambda} \right] \geq 1 - 2q(1-p)^{|\Sigma|} - 2^{-4\lambda+1}$$

⁶In fact, a simple extension of the Berlekamp-Welch algorithm [WB86] gives efficient unique decoding for univariate multiplicity codes.

$$\implies \Pr_{H \leftarrow \text{Bias}_{q,1/\lambda^4, \mathbb{F}_q^s}} \left[\Pr[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \mathcal{B}(|\text{adv}_H\rangle, x\|0^{q-\lambda})] \geq 1 - 2^{-\lambda} \right] \geq 1 - 2^{-2\lambda}.$$

The algorithm BiasYZ simply runs $\mathcal{B}(|\text{adv}_H\rangle, x\|0^{q-\lambda})$ given advice $|\text{adv}_H\rangle$ and input $x \in \{0, 1\}^\lambda$. \square

We therefore have the following corollary by a simple union bound over all $x \in \{0, 1\}^\lambda$.

Corollary 5.24. *For each security parameter $\lambda \in \mathbb{N}$, define the code $C_\lambda = \text{Mult}_{s, \mathbb{F}_q, k}$, where $\lambda^5 < q \leq 2\lambda^5$ is a prime, $k = \lambda^3$, and $s = \lambda$. Then, there exists an efficient/uniform quantum algorithm BiasYZ and a family of $\text{poly}(\lambda)$ -qubit states $\{|\text{adv}_H\rangle\}_H$ such that for sufficiently large λ ,*

$$\Pr_{H \leftarrow \text{Bias}_{q,1/\lambda^4, \mathbb{F}_q^s}} \left[\forall x \in \{0, 1\}^\lambda, \Pr[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \text{BiasYZ}(|\text{adv}_H\rangle, x)] \geq 1 - 2^{-\lambda} \right] \geq 1 - 2^{-\lambda}.$$

6 Separating QMA from QCMA

For the rest of the paper, we fix a code family $C_\lambda := \text{Mult}_{s, \mathbb{F}_q, k}$ for $\lambda^5 < q \leq 2\lambda^5$ a prime,⁷ $k = \lambda^3$, and $s = \lambda$. For any subset $E \subseteq \{0, 1\}^\lambda \times \Sigma^q$ and function $H : [q] \times \Sigma \rightarrow \{0, 1\}$, we define the oracle

$$O[H, E](x, \mathbf{v}) = \begin{cases} 1 & \text{if } (x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} \wedge (x, \mathbf{v}) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Our proofs in this section are fairly standard and follow [BHNZ25], but we include them for completeness.

We begin by defining an oracle-input problem based on the code intersection subset size checking problem. Note that our NO instances are slightly different than those considered in [LLPY23, BDK24], which will affect the QCMA lower bound, but not the QMA containment.

Definition 6.1 (YES and NO instances of the code intersection subset size problem). *Let Good denote the set of functions $H : [q] \times \mathbb{F}_q^s \rightarrow \{0, 1\}$ such that the algorithm in Corollary 5.24 succeeds on all $x \in \{0, 1\}^\lambda$ with probability at least $2/3$. Our oracle-input separation between QMA and QCMA involves distinguishing between $O[H, E]$ with H and E being the following:*

1. YES instances: $H \in \text{Good}$ and $E = \{0, 1\}^\lambda \times \Sigma^q$.
2. NO instances: $H \in \text{Good}$ and subsets $E \subseteq \{0, 1\}^\lambda \times \Sigma^q$ such that $|E| \leq 2^\lambda/3$.

Remark 6.2. *One can also embed $O[H, E] : \{0, 1\}^\lambda \times \Sigma^q \rightarrow \{0, 1\}$ into a binary input domain oracle with $\lambda + q \log |\Sigma| \leq \lambda^7$ -bit inputs, in such a way that for sufficiently large $\lambda \geq \lambda_0$ (where $\lambda_0 \in \mathbb{N}$ is some constant), there is at most one security parameter associated with each input length.*

6.1 The QMA Proof System

We first show that there is a QMA proof system that distinguishes between YES and NO instances of the code intersection subset size checking problem, as defined in Definition 6.1.

Lemma 6.3 (A QMA proof system). *There exists a polynomial-time uniform quantum query algorithm V which makes one query to the oracle $O[H, E]$, such that for sufficiently large λ , the following holds:*

1. **Completeness.** *For all $H \in \text{Good}$, when $E = \{0, 1\}^\lambda \times \Sigma^q$, there exists a $\text{poly}(\lambda)$ -qubit state $|\text{adv}_H\rangle$ such that*

$$\Pr[V^{O[H, E]}(|\text{adv}_H\rangle) = 1] \geq \frac{2}{3}.$$

2. **Soundness.** *For all H , sets $E \subseteq \{0, 1\}^\lambda \times \Sigma^q$ where $|E| \leq 2^\lambda/3$, and quantum states $|\text{adv}_H^*\rangle$,*

$$\Pr[V^{O[H, E]}(|\text{adv}_H^*\rangle) = 1] \leq \frac{1}{3}.$$

⁷To be concrete, we can take q to be the smallest prime larger than λ^5 (which is always at most $2\lambda^5$).

Proof. By definition, for $H \in \text{Good}$ as defined in Corollary 5.24, there exists a $\text{poly}(\lambda)$ -qubit state $|\text{adv}_H\rangle$ and efficient algorithm BiasYZ such that for any $x \in \{0, 1\}^\lambda$,

$$\Pr_{\text{BiasYZ}}[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda, H}} : \mathbf{v} \leftarrow \text{BiasYZ}(|\text{adv}_H\rangle, x)] \geq 1 - 2^{-\lambda} \geq \frac{2}{3}.$$

The verifier V operates as follows: it samples a uniformly random $x \in \{0, 1\}^\lambda$ and runs $\mathbf{v} \leftarrow \text{BiasYZ}(|\text{adv}_H\rangle, x)$. Finally, V queries O at (x, \mathbf{v}) and returns the output of O . The efficiency and uniformity of V follows from the efficiency and uniformity of BiasYZ combined with the fact that V makes one oracle query.

We now argue completeness and soundness. If $E = \{0, 1\}^\lambda \times \Sigma^q$, then

$$\Pr[V^{O[H, E]}(|\text{adv}_H\rangle) = 1] = \Pr_{\text{BiasYZ}, x}[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda, H}} : \mathbf{v} \leftarrow \text{BiasYZ}(|\text{adv}_H\rangle, x)] \geq \frac{2}{3}.$$

On the other hand, note that V will always output 0 if there does not exist a \mathbf{v} such that $(x, \mathbf{v}) \in E$. Thus, if $|E| \leq 2^\lambda/3$, then for all H and quantum states $|\text{adv}_H^*\rangle$,

$$\Pr[V^{O[H, E]}(|\text{adv}_H^*\rangle) = 1] \leq \Pr_{x \leftarrow \{0, 1\}^\lambda}[\exists \mathbf{v} : (x, \mathbf{v}) \in E] \leq \frac{|E|}{2^\lambda} \leq \frac{1}{3}. \quad \square$$

6.2 Non-Existence of QCMA Proof Systems

We begin by showing that any QCMA verifier/algorithm can be turned into a very good hash value guesser.

Lemma 6.4 (Good guessers from QCMA algorithms). *Assume there exists a QCMA algorithm \mathcal{A} such that for instances of size λ , \mathcal{A} takes a $t(\lambda)$ -bit witness and makes $Q(\lambda)$ oracle queries. Then for all $\ell \leq 2^{\lambda/3}$, there exists a algorithm Guesser which makes no queries such that for any $H \in \text{Good}$,*

$$\Pr \left[\begin{array}{l} \forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) \\ \wedge (x_i\|0^{q-\lambda}, \mathbf{v}_i) \in R_{C_{\lambda, H}} \end{array} : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \text{Guesser}(1^\ell) \right] \geq 2^{-t(\lambda)} \cdot \left(\frac{1}{144Q(\lambda)^2} \right)^\ell.$$

Proof. The algorithm $\mathcal{A}^{O[H, E]}$ can be thought of starting from a state $|w, 0\rangle$ and applying a sequence of unitaries V_0, \dots, V_Q interlaced with queries to $O[H, E]$ before measuring the first qubit in the standard basis. The state of the algorithm right before its final measurement is then given by

$$V_Q \cdot O[H, E] \cdot V_{Q-1} \cdot O[H, E] \dots O[H, E] \cdot V_0 |w, 0\rangle.$$

Let Guesser be the algorithm described in Figure 2 which outputs ℓ tuples of codewords and hash values.

Let G be the event that the witness w sampled is a good witness for H , and let E_i be the event that the i 'th round of Guesser appends a tuple $(x\|0^{q-\lambda}, \mathbf{v}) \notin \Delta_{i-1}$ such that $(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda, H}}$.

Claim 6.5. $\forall i \in [\ell], \Pr[E_i | E_{i-1} \wedge \dots \wedge E_1 \wedge G] \geq \frac{1}{144Q(\lambda)^2}$.

Proof. Fix an index $1 \leq i \leq \ell$. Observe that conditioned on $E_1 \wedge \dots \wedge E_{i-1}$ occurring, this means that Δ_{i-1} consists of $i-1$ distinct tuples $\{(x_j\|0^{q-\lambda}, \mathbf{v}_j)\}_{j=1}^{i-1}$ such that $(x_j\|0^{q-\lambda}, \mathbf{v}_j) \in R_{C_{\lambda, H}}$ for all $1 \leq j \leq i-1$.

Since $O_{\Delta_{i-1}} = O[H, \Delta_{i-1}]$ corresponds to a NO instance (as $|\Delta_{i-1}| = i-1 < \ell \leq 2^\lambda/3$) while $O[H, \{0, 1\}^\lambda \times \Sigma^q]$ corresponds to a YES instance, by the completeness and soundness of the QCMA algorithm \mathcal{A} and the fact that the event G implies we have a good witness, Theorem 4.8 implies that the query mass on the inputs where the two oracles differ must be at least $((2/3 - 1/3)/4)^2/Q = 1/144Q$.

As $O_{\Delta_{i-1}}$ and $O[H, \{0, 1\}^\lambda \times \Sigma^q]$ differ precisely on inputs $(x\|0^{q-\lambda}, \mathbf{v}) \notin \Delta_{i-1}$ where $(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_{\lambda, H}}$, it follows that measuring a random query of \mathcal{A} produces a good tuple $(x\|0^{q-\lambda}, \mathbf{v})$ with probability at least $\frac{1}{Q} \cdot \frac{1}{144Q} = \frac{1}{144Q^2}$. \square

Observing that $\Pr[G] \geq 2^{-t(\lambda)}$ as there is always at least one good witness for any H , we conclude that

$$\Pr[E_1 \wedge \dots \wedge E_\ell] \geq \Pr[G] \cdot \prod_{i=1}^\ell \Pr[E_i | E_{i-1} \wedge \dots \wedge E_1 \wedge G] \geq 2^{-t(\lambda)} \cdot \left(\frac{1}{144Q(\lambda)^2} \right)^\ell. \quad \square$$

Guesser(1^ℓ):

1. Sample a random $w \in \{0, 1\}^t$ and initialize $\Delta_0 = \emptyset$.
2. For $i \in [\ell]$:
 - (a) Sample $j \leftarrow \{0, \dots, Q-1\}$ uniformly randomly.
 - (b) Compute the state $V_j \mathcal{O}_{\Delta_{i-1}} V_{j-1} \dots \mathcal{O}_{\Delta_{i-1}} V_0 |w, 0\rangle$, where $\mathcal{O}_{\Delta_{i-1}}$ is the oracle unitary defined by
$$|x, \mathbf{v}\rangle |y\rangle |z\rangle \mapsto |x, \mathbf{v}\rangle |y \oplus f_{\Delta_{i-1}}(x, \mathbf{v})\rangle |z\rangle, \text{ where } f_{\Delta_{i-1}}(x, \mathbf{v}) := \begin{cases} 1 & \text{if } (x, \mathbf{v}) \in \Delta_{i-1}, \\ 0 & \text{otherwise.} \end{cases}$$
 - (c) Measure the first register in the standard basis for output (x, \mathbf{v}) and update $\Delta_i := \Delta_{i-1} \cup \{(x, \mathbf{v})\}$.
3. Output $\{(x_{[1:\lambda]}, \mathbf{v}) : (x, \mathbf{v}) \in \Delta_\ell\}$.

Figure 2: The Hash Value Guesser, given a successful QCMA verifier for the CISS problem.

Separately, we can show the following *upper bound* on the success probability of Guesser. The bound follows from the fact that Guesser is not making any queries to the oracle, and thus knows nothing about H .

Lemma 6.6 (Guessing probability upper bound). *For sufficiently large λ , the following holds: fix any algorithm Guesser which makes no oracle queries. Then, for all $\ell \leq 2^\lambda$,*

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\begin{array}{l} \forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) \\ \wedge (x_i \| 0^{q-\lambda}, \mathbf{v}_i) \in R_{C_\lambda, H} \end{array} : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \text{Guesser}(1^\ell) \right] \leq \left(1 - \frac{1}{\lambda^4}\right)^{\lambda^5 \ell / 2}.$$

Proof. Consider any output $\{(x_1, \mathbf{v}_1), \dots, (x_\ell, \mathbf{v}_\ell)\}$ of Guesser. First, observe that if (x_i, \mathbf{v}_i) are distinct and $H(\mathbf{v}_i) = x_i \| 0^{q-\lambda}$, then we must have distinct \mathbf{v}_i or else Guesser will fail. Since we require $\mathbf{v}_i \in C_\lambda$ for all $i \in [\ell]$, we now consider the sets $S_j := \{\sigma \in \Sigma \mid \exists i \in [\ell] : (\mathbf{v}_i)_j = \sigma\}$ for all $j \in [q]$.

By definition, for all $i \in [\ell]$ and $j \in [q]$, $(\mathbf{v}_i)_j \in S_j$, so if we think of $\{S_j\}_{j=1}^q$ as input lists, the output list for the code C_λ must contain $\mathbf{v}_i \in [\ell]$ and thus $|\{\mathbf{v} \in C_\lambda : \forall j \in [q], \mathbf{v}_j \in S_j\}| \geq \ell$.

Corollary 4.18 thus implies that $\frac{1}{q} \sum_j |S_j| \geq \ell/2$ if λ is sufficiently large. In order for Guesser to succeed, it must correctly guess the output of H on all symbols in $\bigcup_{i=1}^q S_i$ which contain at least $\sum_j |S_j| \geq \frac{q\ell}{2} = \lambda^5 \ell / 2$ distinct points.

Since we sample $H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}$, this occurs with probability at most $(1 - \frac{1}{\lambda^4})^{\lambda^5 \ell / 2}$. \square

We can combine the upper bound and lower bound to conclude that any QCMA algorithm \mathcal{A} for the code intersection subset size (CISS) problem must misclassify some YES or NO instance.

Lemma 6.7. *For all constants $a > 0$ and functions $Q(\lambda), t(\lambda)$ that satisfy $Q(\lambda) \leq a\lambda^a, t(\lambda) \leq a\lambda^a$ for sufficiently large λ . Then for sufficiently large λ , for all quantum query algorithms \mathcal{A} which take a classical witness of length $t(\lambda)$ and make $Q(\lambda)$ queries to the oracle $O[H, E]$ of size λ , there exists an oracle $O[H^*, E^*]$ of size λ such that $H^* \in \text{Good}$ and*

1. either $E^* = \{0, 1\}^\lambda \times \Sigma^q$, but for all witnesses w of length $t(\lambda)$,

$$\Pr[\mathcal{A}^{O[H^*, E^*]}(w) = 1] < \frac{2}{3},$$

2. or $E^* \subseteq F^* \times \Sigma^q \subseteq \{0, 1\}^\lambda \times \Sigma^q$ where $|F^*| \leq 2^\lambda / 3$, but there exists a witness \tilde{w} of length $t(\lambda)$ such that

$$\Pr[\mathcal{A}^{O[H^*, E^*]}(\tilde{w}) = 1] > \frac{1}{3}.$$

Proof. Suppose for the sake of contradiction that \mathcal{A} properly classifies all YES and NO instances. Setting $\ell = t \ll 2^\lambda/3$, we can apply Corollary 5.24 and Lemma 6.4, yielding a guesser Guesser where

$$\begin{aligned} & \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) \wedge (x_i \| 0^{q-\lambda}, \mathbf{v}_i) \in R_{C_\lambda, H} : \{(x_i, \mathbf{v}_i)\}_{i=1}^t \leftarrow \text{Guesser}(1^t) \right] \\ \geq & \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) \wedge (x_i \| 0^{q-\lambda}, \mathbf{v}_i) \in R_{C_\lambda, H} : \{(x_i, \mathbf{v}_i)\}_{i=1}^t \leftarrow \text{Guesser}(1^t) \mid H \in \text{Good} \right] \cdot \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} [H \in \text{Good}] \\ \geq & 2^{-t(\lambda)} \cdot \left(\frac{1}{144Q(\lambda)^2} \right)^{t(\lambda)} \cdot (1 - 2^{-\lambda}) \geq \left(\frac{1}{576Q(\lambda)^2} \right)^{t(\lambda)}. \end{aligned}$$

But this is impossible as Lemma 6.6 implies that the success probability of Guesser is at most

$$\left(1 - \frac{1}{\lambda^4} \right)^{\lambda^5 t(\lambda)/2} \leq (e^{-\lambda/2})^{t(\lambda)} \ll \left(\frac{1}{576Q(\lambda)^2} \right)^{t(\lambda)}. \quad \square$$

A straightforward diagonalization argument thus gives us our desired separation (see Section §B for details).

Theorem 6.8. *There exists a classical oracle $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$ such that $\text{QMA}^\mathcal{O} \cap \text{AM}^\mathcal{O} \not\subseteq \text{QCMA}^\mathcal{O}$.⁸*

7 Separating BQP/qpoly from BQP/poly

We begin by proving the main technical result of this section, which is a search-like separation between BQP/qpoly and BQP/poly. In particular, we prove that given as input $x \in \{0, 1\}^\lambda$ and oracle access to $O[H, \{0, 1\}^\lambda \times \Sigma^q]$ for H sampled from $\text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}$, the problem of finding \mathbf{v} such that $(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H}$ is in FBQP/qpoly but not FBQP/poly (on average).

Lemma 7.1. *For all security parameters $\lambda \in \mathbb{N}$, let $\lambda^5 < q < 2\lambda^5$ be a prime, $k = \lambda^3$, and $s = \lambda$, and define the code $C_\lambda = \text{Mult}_{s, \mathbb{F}_q, k}$. In addition, for all λ , define the set $E_\lambda := \{0, 1\}^\lambda \times \Sigma^q$, where $\Sigma = \mathbb{F}_q^s$. Then the following hold:*

1. *There is an polynomial-time uniform quantum query algorithm \mathcal{A} , such that for all oracles O there exists a family of $\text{poly}(\lambda)$ -qubit quantum advice states (depending only on the oracle) $\{|z_O\rangle\}_O$ such that*

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} [\forall x \in \{0, 1\}^\lambda, \Pr[(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \mathcal{A}^{O[H, E_\lambda]}(x, |z_O\rangle)] \geq 1 - \text{negl}(\lambda)] \geq 1 - \text{negl}(\lambda).$$

2. *For all unbounded-time quantum algorithms \mathcal{B} that make $Q(\lambda) = \text{poly}(\lambda)$ oracle queries to O , and all families of $t(\lambda) = \text{poly}(\lambda)$ -bit classical advice strings (depending only on the oracle) $\{z_O\}_O$,*

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}, x \leftarrow \{0, 1\}^\lambda} [(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \mathcal{B}^{O[H, E_\lambda]}(x, z_O)] \leq \text{negl}(\lambda).$$

Proof. Throughout this proof we will assume that λ is sufficiently large and argue with respect to asymptotics. To prove Item 1, we simply set the advice as $|z_O\rangle = |\text{adv}_H\rangle$ from Corollary 5.24, and the algorithm just runs the algorithm in Corollary 5.24, which implies that

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} [\forall x \in \{0, 1\}^\lambda, \Pr[(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \mathcal{A}(|z_O\rangle, x)] \geq 1 - 2^{-\lambda}] \geq 1 - 2^{-\lambda}.$$

We now move to proving Item 2. Suppose for the sake of contradiction that there exists a polynomial $p(\lambda)$, an adversary \mathcal{B} which makes $Q(\lambda) = \text{poly}(\lambda)$ oracle queries, and a family of $t(\lambda) = \text{poly}(\lambda)$ -bit classical advice $\{z_O\}_O$ such that for infinitely many λ ,

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}, x \leftarrow \{0, 1\}^\lambda} [(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : \mathbf{v} \leftarrow \mathcal{B}^{O[H, E_\lambda]}(z_O, x)] \geq \frac{1}{p(\lambda)}.$$

⁸We note that this separation can be easily strengthened to $\text{QMA}^\mathcal{O} \cap \text{SBP}^\mathcal{O} \not\subseteq \text{QCMA}^\mathcal{O}$, as with most set-approximation-flavored oracles.

Our goal will be to arrive at a contradiction by showing that this algorithm \mathcal{B} implies a (too good) sampler for $R_{C_\lambda, H}$ that works for infinitely many choices of λ . Consider the Q -query algorithm $\mathcal{A}_1^{O[H, E_\lambda]}(z_O)$ which samples a random $x \leftarrow \{0, 1\}^\lambda$, runs $\mathcal{B}^{O[H, E_\lambda]}(z_O, x)$ to get \mathbf{v} , and outputs (x, \mathbf{v}) . By the definition of \mathcal{B} , we have

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : (x, \mathbf{v}) \leftarrow \mathcal{A}_1^{O[H, E_\lambda]}(z_O) \right] \geq \frac{1}{p(\lambda)}, \quad (1)$$

for infinitely many λ . We denote with Λ the set of all λ for which this bound holds. At a very high level, we will first show that for all but finitely many $\lambda \in \Lambda$, \mathcal{A}_1 's queries are concentrated on very few points. Once we know that queries to \mathcal{A}_1 are concentrated on a couple points, replacing the real oracle with an oracle that only contains those few points will give rise to a sampler for many more points than are contained within the oracle itself.

For each function H and set $S \subseteq \{0, 1\}^\lambda \times \Sigma^q$, we define $M_{H, \bar{S}}$ as the query mass that \mathcal{A}_1 places on points which differ between $O[H, S]$ and $O[H, E]$. Let SmallSet_H be the event that there exists a list $L_H \subseteq \{0, 1\}^\lambda \times \Sigma^q$ such that $|L_H| \leq 2^{\lambda/2}$ and $M_{H, \bar{L}_H} \leq \frac{1}{256p(\lambda)^2 Q(\lambda)}$.

Claim 7.2. *Whenever SmallSet_H does not occur, for all $\ell \leq 2^{\lambda/2}$, there is an algorithm, Guesser, which makes no queries to an oracle, and outputs a list of ℓ distinct points from $R_{C_\lambda, H}$ with probability at least $2^{-t} \cdot \left(\frac{1}{256p(\lambda)^2 Q(\lambda)^2} \right)^\ell$.*

Proof. The algorithm Guesser is identical to the algorithm in Figure 2, except starting from \mathcal{A}_1 instead of a QCMA verifier for the code intersection subset size problem.

The proof follows similarly as well. Let X_i be the event that the i 'th round of Guesser(1^ℓ) outputs a tuple $(x \| 0^{q-\lambda}, \mathbf{v}) \notin \Delta_{i-1}$ such that $(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H}$ and G be the event that the advice is guessed correctly. By assumption, and because $|\Delta_{i-1}| \leq 2^{\lambda/2}$, we have that \mathcal{A}_1 places at least $\frac{1}{256p(\lambda)^2 Q(\lambda)}$ query mass on points which differ between $O[H, \Delta_{i-1}]$ and $O[H, E_\lambda]$. Therefore, we have that $\Pr[X_i | X_1 \wedge \dots \wedge X_{i-1} \wedge G] \geq \frac{1}{256p(\lambda)^2 Q(\lambda)^2}$.

Applying the chain rule, together with the fact that $\Pr[G] \geq 2^{-t(\lambda)}$, we get that the probability of sampling ℓ distinct points from $R_{C_\lambda, H}$ is at least

$$\Pr[X_1 \wedge \dots \wedge X_\ell \wedge G] \geq 2^{-t(\lambda)} \cdot \left(\frac{1}{256p(\lambda)^2 Q(\lambda)^2} \right)^\ell. \quad \square$$

As a corollary, we have that for all but finitely many $\lambda \in \Lambda$, SmallSet_H must occur with high probability.

Claim 7.3. *There are only finitely many $\lambda \in \Lambda$ such that $\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}}[\text{SmallSet}_H] < 1 - \frac{1}{4p(\lambda)}$.*

Proof. Assume for the sake of contradiction that there are infinitely many $\lambda \in \Lambda$ such that $\Pr[\text{SmallSet}_H] < 1 - \frac{1}{4p(\lambda)}$. Whenever SmallSet_H does not occur, the previous claim gives us a sampler. Thus, when H is sampled from $\text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}$, the probability of the sampler outputting ℓ distinct points from $R_{C_\lambda, H}$ is at least

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \text{Guesser}(1^\ell) \right] \geq \frac{1}{4p} \cdot 2^{-t} \cdot \left(\frac{1}{256p(\lambda)^2 Q(\lambda)^2} \right)^\ell.$$

Taking $\ell = \max\{t, \lambda\} \leq 2^{\lambda/2} \leq 2^\lambda$ gives a sampling success probability of

$$\frac{1}{4p(\lambda)} \cdot \left(\frac{1}{512p(\lambda)^2 Q(\lambda)^2} \right)^\ell \geq \left(\frac{1}{2048p(\lambda)^3 Q(\lambda)^2} \right)^\ell,$$

which is a contradiction since Lemma 6.6 implies this success probability should be at most

$$(e^{-\lambda/2})^\ell \ll \left(\frac{1}{2048p(\lambda)^3 Q(\lambda)^2} \right)^\ell. \quad \square$$

Now we know that SmallSet_H occurs with high probability for infinitely many $\lambda \in \Lambda$. For all H for which SmallSet_H occurs, we denote L_H^* to refer to any arbitrary set of size at most $2^{\lambda/2}$ such that $M_{H, L_H^*} \leq \frac{1}{256p(\lambda)^2 Q(\lambda)}$; if SmallSet_H does not occur, then we define $L_H^* := \emptyset$. For each function H , define the punctured oracle

$$O_H^*(x, \mathbf{v}) = \begin{cases} 1 & \text{if } (x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} \wedge (x, \mathbf{v}) \in L_H^*, \\ 0 & \text{otherwise.} \end{cases}$$

Then we have that, conditioned on λ being such that $\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}}[\text{SmallSet}_H] \geq 1 - \frac{1}{4p(\lambda)}$,

$$\begin{aligned} & \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : (x, \mathbf{v}) \leftarrow \mathcal{A}_1^{O_H^*}(z_O) \mid \text{SmallSet}_H \right] \\ & \geq \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : (x, \mathbf{v}) \leftarrow \mathcal{A}_1^{O[H, E]}(z_O) \mid \text{SmallSet}_H \right] - \frac{1}{4p(\lambda)} \\ & \geq \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[(x\|0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H} : (x, \mathbf{v}) \leftarrow \mathcal{A}_1^{O[H, E]}(z_O) \wedge \text{SmallSet}_H \right] - \frac{1}{4p(\lambda)} \\ & \geq \frac{1}{p(\lambda)} - \frac{1}{2p(\lambda)} = \frac{1}{2p(\lambda)}. \end{aligned}$$

In the first line, we use the hybrid lemma (Theorem 4.8) combined with our bound on the query mass of \mathcal{A}_1 outside of L_H . In the second line, we use the definition of conditional probability, together with the fact that $\Pr[\text{SmallSet}_H] \leq 1$. We conclude by using a union bound, together with the fact that $\Pr[\text{SmallSet}_H] \geq 1 - \frac{1}{4p(\lambda)}$ and the fact that the probability of \mathcal{A}_1 sampling a point in $R_{C_\lambda, H}$ is at least $\frac{1}{p(\lambda)}$ by assumption.

To arrive at a contradiction we construct yet *another* sampler. By the definition of \mathcal{A}_1 , running it produces a uniformly random $x \in \{0, 1\}^\lambda$, so there exists an algorithm \mathcal{A}_2 which, on input 1^ℓ , runs \mathcal{A}_1 ℓ times and satisfies

$$\begin{aligned} & \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \mathcal{A}_2^{O_H^*}(z_O, 1^\ell) \right] \\ & \geq \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} [\text{SmallSet}_H] \cdot \prod_{i=1}^\ell \left(\frac{1}{2p(\lambda)} - \frac{i-1}{2^\lambda} \right). \end{aligned}$$

Here we applied the definition of conditional probability and used the fact that since the x_i 's are uniformly random, the probability that $x_i \in \bigcup_{j=1}^{i-1} \{x_j\}$ is at most $\frac{i-1}{2^\lambda}$.

As O_H^* has at most $2^{\lambda/2}$ nonzero points, we can hardwire $2^{\lambda/2} \cdot (\lambda + \log |\Sigma|^q) \leq 2^{2\lambda/3}$ bits of advice and simulate O_H^* . By guessing this extra advice along with z_O , we get an algorithm \mathcal{A}_3 such that for $\ell = 2^{3\lambda/4}$,

$$\begin{aligned} & \Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \mathcal{A}_3(1^\ell) \right] \\ & \geq 2^{-(t(\lambda) + 2^{2\lambda/3})} \cdot \left(1 - \frac{1}{4p(\lambda)} \right) \cdot \prod_{i=1}^\ell \left(\frac{1}{2p(\lambda)} - \frac{i-1}{2^\lambda} \right) \\ & \geq 2^{-\ell} \cdot \prod_{i=1}^\ell \left(\frac{1}{2p(\lambda)} - \frac{\ell}{2^\lambda} \right) \geq 2^{-\ell} \cdot \left(\frac{1}{3p(\lambda)} \right)^\ell = \left(\frac{1}{6p(\lambda)} \right)^\ell. \end{aligned}$$

Applying Lemma 6.6 for $\ell = 2^{3\lambda/4} \leq 2^\lambda$, we have a contradicting upper bound on the success probability of

$$\Pr_{H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}} \left[\forall i \neq j, (x_i, \mathbf{v}_i) \neq (x_j, \mathbf{v}_j) : \{(x_i, \mathbf{v}_i)\}_{i=1}^\ell \leftarrow \mathcal{A}_3(1^\ell) \right] \leq \left(1 - \frac{1}{\lambda^4} \right)^{\lambda^{5\ell/2}} \leq e^{-\lambda\ell/2} \ll \left(\frac{1}{6p(\lambda)} \right)^\ell.$$

□

Having shown that there is a search problem outside of (average-case) FBQP/poly, we now apply Lemma 4.21. This was essentially established in [LLPY23], but we make minor modifications to deal with quantum queries.

Lemma 7.4. *There is a family of distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$, where \mathcal{D}_λ is supported on tuples (G, \mathcal{O}') of functions $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ and $\mathcal{O}' : \{0, 1\}^{\text{poly}(\lambda)} \rightarrow \{0, 1\}^{\text{poly}(\lambda)}$, satisfying the following:*

1. *There is an polynomial-time uniform quantum algorithm \mathcal{A} which makes one query to \mathcal{O}' such that for all \mathcal{O}' there exists a family of $\text{poly}(\lambda)$ -qubit quantum advice $\{|z_{\mathcal{O}'}\rangle\}_{\mathcal{O}'}$ such that*

$$\Pr_{(G, \mathcal{O}') \leftarrow \mathcal{D}_\lambda} [\forall x \in \{0, 1\}^\lambda, \Pr[\mathcal{A}^{\mathcal{O}'}(|z_{\mathcal{O}'}\rangle, x) = G(x)] \geq 1 - \text{negl}(\lambda)] \geq 1 - \text{negl}(\lambda).$$

2. *For all quantum query algorithms algorithm \mathcal{B} that makes $\text{poly}(\lambda)$ queries to \mathcal{O}' and receiving a family of $\text{poly}(\lambda)$ -bit classical advice $\{z_{\mathcal{O}'}\}_{\mathcal{O}'}$ depending only on \mathcal{O}' ,*

$$\Pr_{(G, \mathcal{O}') \leftarrow \mathcal{D}_\lambda, x \leftarrow \{0, 1\}^\lambda} [\mathcal{B}^{\mathcal{O}'}(z_{\mathcal{O}'}, x) = G(x)] \leq \frac{3}{5},$$

for all sufficiently large λ .

Proof. Define \mathcal{D}_λ as follows: first, sample a random function $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ and $H \leftarrow \text{Bias}_{q, 1/\lambda^4, \mathbb{F}_q^s}$ and let

$$\mathcal{O}'(x, \mathbf{v}) := \begin{cases} G(x) & \text{if } (x \| 0^{q-\lambda}, \mathbf{v}) \in R_{G, H}, \\ \perp & \text{otherwise.} \end{cases}$$

We begin by showing easiness with quantum advice. Let $(\mathcal{A}', \{|z'_H\rangle\}_H)$ be the algorithm and advice family from Item 1 of Lemma 7.1. We now construct an algorithm \mathcal{A} and family of mixed state advice $\{\rho_{\mathcal{O}'}\}_{\mathcal{O}'}$.⁹ We describe a randomized procedure to set $\rho_{\mathcal{O}'}$ given an oracle \mathcal{O}' , but in reality, we will set $\rho_{\mathcal{O}'}$ to be the mixed state corresponding to the mixture over outputs of this procedure. Sample (G, H) from the distribution of \mathcal{D}_λ conditioned on \mathcal{O}' ; by construction, the joint distribution of (G, H, \mathcal{O}') sampled in this procedure is identical to \mathcal{D}_λ . We then set our advice to be $\rho_{\mathcal{O}'} = |z'_H\rangle$. The algorithm \mathcal{A} on input x will run $\mathbf{v} \leftarrow \mathcal{A}'(\rho_{\mathcal{O}'}, x)$, query (x, \mathbf{v}) to \mathcal{O}' , and output whatever \mathcal{O}' returns. Item 1 of Lemma 7.1 then implies that \mathcal{A} is efficient/uniform and that

$$\Pr_{(G, \mathcal{O}') \leftarrow \mathcal{D}_\lambda} [\forall x \in \{0, 1\}^\lambda, \Pr[\mathcal{A}^{\mathcal{O}'}(\rho_{\mathcal{O}'}, x) = G(x)] \geq 1 - \text{negl}(\lambda)] \geq 1 - \text{negl}(\lambda).$$

Now suppose for the sake of contradiction that there was some algorithm \mathcal{B} which made $Q(\lambda) = \text{poly}(\lambda)$ queries and had $t(\lambda) = \text{poly}(\lambda)$ -bit classical advice $\{z_{\mathcal{O}'}\}_{\mathcal{O}'}$ such that for infinitely many λ ,

$$\Pr_{\substack{(G, \mathcal{O}') \leftarrow \mathcal{D}_\lambda \\ x \leftarrow \{0, 1\}^\lambda}} [\mathcal{B}^{\mathcal{O}'}(z_{\mathcal{O}'}, x) = G(x)] > \frac{3}{5}.$$

We know that \mathcal{O}' returns $G(x)$ only if the query $(x, \mathbf{v}) \in R_{G, H}$. Thus, by a direct reduction to Lemma 4.21, for a $\frac{1}{4000Q(\lambda)^2}$ fraction of $x \in \{0, 1\}^\lambda$, measuring a random query of \mathcal{B} to a randomly sampled oracle $\mathcal{O}' \leftarrow \mathcal{D}_\lambda$ will produce $(x, \mathbf{v}) \in R_{G, H}$ with probability at least $\frac{1}{3200Q(\lambda)^2}$.

But now observe that \mathcal{O}' can be simulated by querying G and $\mathcal{O} := O[H, \{0, 1\}^\lambda \times \Sigma^q]$. Thus, for each function G , we define the following Q -query algorithm $\mathcal{B}'[G]$ and classical advice $\{z'_\mathcal{O}[G]\}_{\mathcal{O}}$. First, construct \mathcal{O}' from (G, \mathcal{O}) (since \mathcal{O} uniquely determines H) before setting $z'_\mathcal{O}[G] := z_{\mathcal{O}'}$. $\mathcal{B}'[G]^{\mathcal{O}}(z'_\mathcal{O}[G], x)$ will run $\mathcal{B}^{\mathcal{O}'}(z'_\mathcal{O}[G], x)$ where \mathcal{B}' will simulate the oracle \mathcal{O}' using its own oracle \mathcal{O} and the hardwired oracle G and measure a uniformly chosen query of \mathcal{B} . As noted earlier, this means that

$$\Pr_{G, H, x} [(x, \mathbf{v}) \in R_{G, H} : \mathbf{v} \leftarrow \mathcal{B}'[G]^{\mathcal{O}}(z'_\mathcal{O}[G], x)] \geq \frac{1}{3200 \cdot 4000 \cdot Q(\lambda)^4} = \frac{1}{\text{poly}(\lambda)}.$$

By taking G^* which maximizes the above probability,¹⁰ $(\mathcal{B}'[G^*], \{z'_\mathcal{O}[G^*]\}_{\mathcal{O}})$ breaks Item 2 of Lemma 7.1. \square

A simple diagonalization argument gives us our desired separation (see Section §B for details).

Theorem 7.5. *There is a classical oracle \mathcal{O} such that $\text{BQP}^{\mathcal{O}}/\text{qpoly} \cap \text{NP}^{\mathcal{O}} \cap \text{coNP}^{\mathcal{O}} \subsetneq \text{BQP}^{\mathcal{O}}/\text{poly}$.¹¹*

⁹This is without loss of generality as a mixed state is a distribution over pure states and so there is always a pure state advice that is at least as good as the mixed state advice.

¹⁰As noted in [LLPY23], finding G^* does not actually require access to the specific H since \mathcal{B}' can find G^* by itself by using its unbounded computational power to enumerate over all possible G and \mathcal{O}_H .

¹¹It is not hard to extend this separation to show that $\text{YQP}^{\mathcal{O}} \cap \text{NP}^{\mathcal{O}} \cap \text{coNP}^{\mathcal{O}} \subsetneq \text{BQP}^{\mathcal{O}}/\text{poly}$, where YQP is the class of problems that can be decided by a BQP machine with *untrusted* quantum advice [Aar07, AD14].

Acknowledgments. We thank Scott Aaronson, Joe Carolan, Ryan Williams, Rohan Goyal, Venkatesan Guruswami, Mary Wootters and Rachel Zhang for patiently answering our many questions. A special thanks to Anand Natarajan for very helpful discussions in the early stages of this project and Alexandru Gheorghiu and Aparna Gupte for pointing out the efficiency benefits of using biased oracles in the Yamakawa-Zhandry algorithm. JB is supported by Henry Yuen’s AFOSR award FA9550-23-1-0363. VV gratefully acknowledges support from a Simons Investigator Award and a Ford Foundation Chair.

References

- [Aar04] Scott Aaronson. Limitations of quantum advice and one-way communication. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 320–332. IEEE, 2004. 6
- [Aar07] Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, 2007. 26
- [Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 229–242. IEEE, 2009. 7
- [Aar10] Scott Aaronson. BQP and the polynomial hierarchy. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing, STOC ’10*, page 141–150, New York, NY, USA, 2010. Association for Computing Machinery. 1
- [AD14] Scott Aaronson and Andrew Drucker. A full characterization of quantum advice. *SIAM Journal on Computing*, 43(3):1131–1183, 2014. 26
- [AK07] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC’07)*, pages 115–128. IEEE, 2007. 1, 5, 6, 8
- [AN02] Dorit Aharonov and Tomer Naveh. Quantum NP—a survey. *arXiv preprint quant-ph/0210077*, 2002. 1
- [BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997. 4, 10
- [BDK24] Shalev Ben-David and Srijita Kundu. Oracle separation of QMA and QCMA with bounded adaptivity. *arXiv preprint arXiv:2402.00298*, 2024. 1, 2, 3, 4, 6, 7, 8, 13, 20
- [BHNZ25] John Bostanci, Jonas Haferkamp, Chinmay Nirkhe, and Mark Zhandry. Separating QMA from QCMA with a classical oracle. *arXiv preprint arXiv:2511.09551*, 2025. 1, 2, 4, 5, 6, 7, 9, 20, 31
- [Bla24] Ian F Blake. *Essays on Coding Theory*. Cambridge University Press, 2024. 29
- [BNZ25] John Bostanci, Barak Nehoran, and Mark Zhandry. A general quantum duality for representations of groups with applications to quantum money, lightning, and fire. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 201–212, 2025. 7
- [Bor09] Émile Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909. 10
- [BV93] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 11–20, 1993. 14
- [Can17] Francesco Cantelli. Sulla probabilita come limita della frequenza. *Rend. Accad. Lincei*, 26:39, 1917. 10
- [CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 673–684. IEEE, 2020. 6, 12

- [CW02] Cleve and Watrous. Sharp quantum versus classical query complexity separations. *Algorithmica*, 34(4):449–461, 2002. 13
- [DKSS13] Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to kakeya sets and mergers. *SIAM Journal on Computing*, 42(6):2305–2328, 2013. 29
- [FGH⁺12] Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, and Peter Shor. Quantum money from knots. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 276–289, 2012. 7
- [FK15] Bill Fefferman and Shelby Kimmel. Quantum vs classical proofs and subset verification. *arXiv preprint arXiv:1510.06750*, 2015. 1, 6, 8
- [GGJL25] Mika Göös, Tom Gur, Siddhartha Jain, and Jiawei Li. Quantum communication advantage in tfnp. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1465–1475, 2025. 5
- [GGKM] Alexandru Gheorghiu, Aparna Gupte, and Greg Kahanamoku-Meyer. Personal communication. 5, 7
- [GS86] S Goldwasser and M Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, STOC ’86, page 59–68, New York, NY, USA, 1986. Association for Computing Machinery. 7
- [GUV09] Venkatesan Guruswami, Christopher Umans, and Salil Vadhan. Unbalanced expanders and randomness extractors from parvaresh–vardy codes. *Journal of the ACM (JACM)*, 56(4):1–34, 2009. 8, 11
- [Kit97] A. Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, 1997. 1
- [Kop15] Swastik Kopparty. Some remarks on multiplicity codes. *arXiv preprint arXiv:1505.07547*, 2015. 19
- [KSY14] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. *Journal of the ACM (JACM)*, 61(5):1–20, 2014. 11, 19
- [KTS22] Itay Kalev and Amnon Ta-Shma. Unbalanced expanders from multiplicity codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, pages 12–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022. 4, 8, 11
- [Liu23] Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 117–143. Springer, 2023. 2, 3, 7
- [LLPY23] Xingjian Li, Qipeng Liu, Angelos Pelecanos, and Takashi Yamakawa. Classical vs quantum advice and proofs under classically-accessible oracle. *arXiv preprint arXiv:2303.04298*, 2023. 1, 2, 3, 4, 6, 7, 12, 13, 20, 25, 26, 31
- [LMY25] Jiahui Liu, Saachi Mutreja, and Henry Yuen. QMA vs. QCMA and pseudorandomness, 2025. 1, 2, 6
- [LS25] Ray Li and Nikhil Shagrithaya. Near-optimal list-recovery of linear code families. *arXiv preprint arXiv:2502.13877*, 2025. 8
- [Lut11] Andrew Lutomirski. Component mixers and a hardness result for counterfeiting quantum money. *arXiv preprint arXiv:1107.0321*, 2011. 2
- [Nie01] Rasmus Refslund Nielsen. *List decoding of linear block codes*. Department of Mathematics, Technical University of Denmark, 2001. 11, 19
- [NN24] Anand Natarajan and Chinmay Nirkhe. A distribution testing oracle separation between qma and qcma. *Quantum*, 8:1377, 2024. 1, 2, 6

- [NY04] Harumichi Nishimura and Tomoyuki Yamakami. Polynomial time quantum computation with advice. *Information Processing Letters*, 90(4):195–204, 2004. 1
- [NZ23] Barak Nehoran and Mark Zhandry. A computational separation between quantum no-cloning and no-telegraphing. *arXiv preprint arXiv:2302.01858*, 2023. 6
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009. 2, 5
- [RT97] M Yu Rosenbloom and Michael Anatol’evich Tsfasman. Codes for the m-metric. *Problemy Peredachi Informatsii*, 33(1):55–63, 1997. 11
- [RZVW24] Noga Ron-Zewi, S Venkitesh, and Mary Wootters. Efficient list-decoding of polynomial ideal codes with optimal list size. *arXiv preprint arXiv:2401.14517*, 2024. 4, 12
- [vDHI06] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *SIAM Journal on Computing*, 36(3):763–778, 2006. 13
- [WB86] Lloyd R Welch and Elwyn R Berlekamp. Error correction for algebraic block codes, December 30 1986. US Patent 4,633,470. 19
- [YZ24] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. *Journal of the ACM*, 71(3):1–50, 2024. 2, 4, 5, 13, 14
- [Zha24] Mark Zhandry. Toward separating QMA from QCMA with a classical oracle. *arXiv preprint arXiv:2411.01718*, 2024. 1, 2, 4, 6
- [Zha25] Mark Zhandry. Quantum money from abelian group actions. *TheoretCS*, 4, 2025. 7

A Duals of Multiplicity Codes

For a field \mathbb{F}_q , the multiplicity $\text{mult}(f, \alpha)$ of a polynomial $f \in \mathbb{F}_q[X]$ at a point $\alpha \in \mathbb{F}_q$ is the largest integer m so that $f^{(i)}(\alpha) = 0$ for any non-negative integer $i < m$. The multiplicity Schwartz-Zippel Lemma from [DKSS13] says that a nonzero degree k univariate polynomial can vanish on at most k points, counting multiplicities.

Lemma A.1 ([DKSS13]). *Let $f \in \mathbb{F}_q[X]$ be a nonzero polynomial of degree at most k . Then $\sum_{\alpha \in \mathbb{F}_q} \text{mult}(f, \alpha) \leq k$.*

Fact A.2 (Hasse derivatives, see [Bla24]). *The following properties hold for the Hasse derivative:*

1. *For any polynomial $f(X) \in \mathbb{F}_q[X]$, integer $i \geq 0$, and point $\alpha \in \mathbb{F}_q$, $f^{(i)}(\alpha)$ is the coefficient of X^i in $f(X + \alpha)$.*
2. *(Linearity) For any $f, g \in \mathbb{F}_q[x]$, $\lambda, \mu \in \mathbb{F}_q$, and $i \geq 0$, $(\lambda \cdot f + \mu \cdot g)^{(i)} = \lambda \cdot f^{(i)} + \mu \cdot g^{(i)}$.*
3. *(Product rule) For any $f, g \in \mathbb{F}_q[x]$ and $i \geq 0$, we have $(f \cdot g)^{(i)} = \sum_{k=0}^i f^{(k)} \cdot g^{(i-k)}$.*

We first derive a natural analogue of Lagrange interpolation for the setting of Hasse derivatives.

Lemma A.3 (Hermite interpolation). *Let \mathbb{F}_q be a field, $s \geq 1$ be a positive integer, and $\alpha_1, \dots, \alpha_n$ be distinct points in \mathbb{F}_q . For $i \in [n]$ and $0 \leq j \leq s - 1$, define $\mu_i(X) := \prod_{i' \neq i} (X - \alpha_{i'})^s$ and $\eta_i(X) := (\mu_i(X))^{-1} \bmod (X - \alpha_i)^s$.*

Then, for all $f(X) \in (\mathbb{F}_q)_{<sn}[X]$, we can write $f(X) = \sum_{i=1}^n \sum_{j=0}^{s-1} A_{i,j}(X) f^{(j)}(\alpha_i)$, where

$$A_{i,j}(X) = \mu_i(X)(X - \alpha_i)^j \sum_{t=0}^{s-1-j} \eta_i^{(t)}(\alpha_i)(X - \alpha_i)^t.$$

Proof. We begin by showing that for any $i, i' \in [n]$ and $0 \leq j, j' \leq s-1$, $A_{i,j}^{(j')}(\alpha_{i'}) = 1$ if $(i, j) = (i', j')$ and 0 otherwise. First, if $i' \neq i$, then $A_{i,j}(X + \alpha_{i'}) = X^s \cdot B_{i,j}(X)$ for some polynomial $B_{i,j}(X)$ so $A_{i,j}^{(j')}(\alpha_{i'}) = 0$ for all j, j' . Similarly, since $A_{i,j}(X + \alpha_i) = X^j \cdot C_{i,j}(X)$ for some polynomial $C_{i,j}(X)$, $A_{i,j}^{(j')}(\alpha_i) = 0$ whenever $i = i'$ and $j' < j$. It thus remains to consider $i = i'$ and $j \leq j'$. By the product rule, we know that

$$\begin{aligned} A_{i,j}^{(j')}(\alpha_i) &= \sum_{k=0}^{j'} \mu_i^{(j'-k)}(\alpha_i) \cdot \left[\sum_{t=0}^{s-1-j} \eta_i^{(t)}(\alpha_i) (X - \alpha_i)^{j+t} \right]^{(k)} (\alpha_i) = \sum_{k=j}^{j'} \mu_i^{(j'-k)}(\alpha_i) \cdot \eta_i^{(k-j)}(\alpha_i) \\ &= \sum_{\ell=0}^{j'-j} \mu_i^{(j'-j-\ell)}(\alpha_i) \cdot \eta_i^{(\ell)}(\alpha_i) = (\mu_i \cdot \eta_i)^{(j'-j)}(\alpha_i). \end{aligned}$$

By construction, $(\mu_i \cdot \eta_i)(X) = 1 + h(X) \cdot (X - \alpha_i)^s$ for some polynomial $h(X)$, so $\mu_i(X + \alpha_i) \cdot \eta_i(X + \alpha_i) = 1 + X^s \cdot h(X + \alpha_i)$. As $0 \leq j' - j \leq s-1$, we conclude that $(\mu_i \cdot \eta_i)^{(j'-j)}(\alpha_i)$ equals 1 if $j' = j$ and 0 otherwise.

Now, let $g(X) = \sum_{i=1}^n \sum_{j=0}^{s-1} A_{i,j}(X) f^{(j)}(\alpha_i)$. Note that since $\deg A_{i,j}, \deg f \leq sn-1$, we know that $\deg(g-f) \leq sn-1$. But for any $i' \in [n]$ and $0 \leq j' \leq s-1$, we have that

$$(g-f)^{(j')}(\alpha_{i'}) = g^{(j')}(\alpha_{i'}) - f^{(j')}(\alpha_{i'}) = \left[\sum_{i=1}^n \sum_{j=0}^{s-1} A_{i,j}^{(j')}(\alpha_{i'}) f^{(j)}(\alpha_i) \right] - f^{(j')}(\alpha_{i'}) = f^{(j')}(\alpha_{i'}) - f^{(j')}(\alpha_{i'}) = 0.$$

Thus, by the multiplicity Schwartz-Zippel Lemma (Lemma A.1), $(g-f)(X) = 0$ and so $f(X) = g(X)$. \square

Definition A.4 (Generalized multiplicity codes). *For invertible matrices $U_1, \dots, U_n \in \mathbb{F}_q^{s \times s}$, define the generalized multiplicity (GM) code $\text{GM}_{s, \mathbb{F}_q}(U_1, \dots, U_n; \alpha_1, \dots, \alpha_n; k) := \{(U_1 \cdot c_1, \dots, U_n \cdot c_n) : c \in \text{Mult}_{s, \mathbb{F}_q}(\alpha_1, \dots, \alpha_n; k)\}$.*

Note that GM codes have distance at least $n - \frac{k-1}{s}$ by Lemma A.1.

Theorem A.5 (Duality of GM codes). *Let \mathbb{F}_q be a field, and $s \geq 1$ be a positive integer, and $\alpha_1, \dots, \alpha_n$ be distinct points in \mathbb{F}_q . Then there exist invertible matrices $U_1, \dots, U_n \in \mathbb{F}_q^{s \times s}$, so that for any positive integer $k < sn$,*

$$\text{Mult}_{s, \mathbb{F}_q}(\alpha_1, \dots, \alpha_n; k) = \text{GM}_{s, \mathbb{F}_q}(U_1, \dots, U_n; \alpha_1, \dots, \alpha_n; sn-k)^\perp.$$

Proof. Consider any pair of polynomials $f(X) \in (\mathbb{F}_q)_{<k}[X]$ and $g(X) \in (\mathbb{F}_q)_{<sn-k}[X]$. Let $h(X) := f(X) \cdot g(X)$, and note that $h(X)$ has degree at most $sn-2$. By Lemma A.3, there exist polynomials $A_{i,j}(X)$ such that

$$h(X) = \sum_{i=1}^n \sum_{j=0}^{s-1} A_{i,j}(X) h^{(j)}(\alpha_i).$$

Letting $a_{i,j}$ denote the coefficient of X^{sn-1} in $A_{i,j}(X)$, we see that the coefficient of X^{sn-1} in $h(X)$ is

$$\sum_{i=1}^n \sum_{j=0}^{s-1} a_{i,j} h^{(j)}(\alpha_i) = \sum_{i=1}^n \sum_{j=0}^{s-1} a_{i,j} \sum_{\ell=0}^j f^{(\ell)}(\alpha_i) g^{(j-\ell)}(\alpha_i) = 0,$$

by the product rule and the fact that $\deg h \leq sn-2$. Consider the following anti-triangular matrices:

$$U_i = \begin{bmatrix} a_{i,0} & a_{i,1} & \cdots & a_{i,s-1} \\ a_{i,1} & \ddots & \ddots & 0 \\ \vdots & \ddots & 0 & \vdots \\ a_{i,s-1} & 0 & \cdots & 0 \end{bmatrix} \in \mathbb{F}_q^{s \times s}.$$

If we define $f_i := (f^{(0)}(\alpha_i), \dots, f^{(s-1)}(\alpha_i)) \in \mathbb{F}_q^s$ and $g_i := (g^{(0)}(\alpha_i), \dots, g^{(s-1)}(\alpha_i)) \in \mathbb{F}_q^s$, we see that

$$\sum_{i=1}^n \langle f_i, U_i \cdot g_i \rangle = \sum_{i=1}^n f_i^T \cdot U_i \cdot g_i = \sum_{i=1}^n \sum_{j=0}^{s-1} a_{i,j} \sum_{\ell=0}^j f^{(\ell)}(\alpha_i) g^{(j-\ell)}(\alpha_i) = 0.$$

We claim that U_i are invertible. To see this, note that by Lemma A.3, for any $i \in [n]$,

$$A_{i,s-1}(X) = \mu_i(X)(X - \alpha_i)^{s-1}\eta_i(\alpha_i) = \eta_i(\alpha_i)(X - \alpha_i)^{s-1} \prod_{i' \neq i} (X - \alpha_{i'})^s,$$

so $a_{i,s-1} = \eta_i(\alpha_i)$. As $\mu_i \cdot \eta_i \equiv 1 \pmod{(X - \alpha_i)^s}$, $\mu_i(\alpha_i) \cdot \eta_i(\alpha_i) = 1$ and thus $a_{i,s-1} \neq 0$. Consequently,

$$\det(U_i) = a_{i,s-1}^s \cdot \det(J_s) = a_{i,s-1}^s \cdot (-1)^{s(s-1)/2} \neq 0,$$

where J_s is the $s \times s$ reversal/exchange matrix. Thus, we have shown that $\sum_{i=1}^n \langle \text{Enc}_C(f)_i, \text{Enc}_{C'}(g)_i \rangle = 0$ for all f and g , where $C := \text{Mult}_{s, \mathbb{F}_q}(\alpha_1, \dots, \alpha_n; k)$ and $C' := \text{GM}_{s, \mathbb{F}_q}(U_1, \dots, U_n; \alpha_1, \dots, \alpha_n; sn - k)$. We conclude that $C = (C')^\perp$ as both C and $(C')^\perp$ are vector spaces of dimension k . \square

Corollary A.6 (Theorem 4.19). *For all parameters s, q , and $k < sq$, $(\text{Mult}_{s, \mathbb{F}_q, k})^\perp$ has distance at least $\frac{k+1}{s}$.*

Proof. By Theorem A.5, $\text{dist}((\text{Mult}_{s, \mathbb{F}_q, k})^\perp) = \text{dist}(\text{GM}_{s, \mathbb{F}_q}(U_1, \dots, U_q; 1, \dots, q; sq - k)) \geq q - \frac{sq-k-1}{s} = \frac{k+1}{s}$. \square

B Diagonalization Arguments

Proof of Theorem 6.8. The proof is nearly identical to that of [BHNZ25], but we include it for completeness. Let $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$ be an oracle and let \mathcal{O}_λ be the restriction to λ^7 -bit inputs, where the lower threshold is λ_0 (per Remark 6.2). We define the unary (promise) language $\mathcal{L}^\mathcal{O}$ so that $1^\lambda \in \mathcal{L}^\mathcal{O}$ precisely when \mathcal{O}_λ is a YES instance and $1^\lambda \notin \mathcal{L}^\mathcal{O}$ precisely when \mathcal{O}_λ is a NO instance.

We consider only oracles \mathcal{O} such that each restriction to size λ^7 -bit inputs encodes either a YES or NO instance, and so the containment $\mathcal{L}^\mathcal{O} \in \text{QMA}^\mathcal{O}$ follows from Lemma 6.3 (as we can hardcode the values of $\mathcal{L}^\mathcal{O}$ on all inputs of length at most λ_0^7). Showing that $\mathcal{L}^\mathcal{O} \in \text{AM}^\mathcal{O}$ for all \mathcal{O} is simple: Arthur samples $x \leftarrow \{0, 1\}^\lambda$ and Merlin responds with any $\mathbf{v} \in \Sigma^q$ such that $(x \| 0^{q-\lambda}, \mathbf{v}) \in R_{C_\lambda, H}$ (which always exists as $H \in \text{Good}$). Arthur accepts iff $\mathcal{O}(x, \mathbf{v}) = 1$. Completeness and soundness follow essentially immediately.

Now we prove the lower bound for QCMA machines. Let M_1, M_2, \dots be an enumeration of all possible Turing machines. Identify any surjective function $\iota : \mathbb{N} \twoheadrightarrow \mathbb{N}^2$ and define functions $j, a : \mathbb{N} \rightarrow \mathbb{N}$ by $(j(\kappa), a(\kappa)) = \iota(\kappa)$. Define $F : \mathbb{N} \rightarrow \mathbb{N}$ so that $F(a)$ is the minimum value such that for all $\lambda \geq F(a)$, any $Q(\lambda) \leq a\lambda^a$ query algorithm with $t(\lambda) \leq a\lambda^a$ -length classical witness must misclassify some \mathcal{O}_λ . By Lemma 6.7, for every integer a , $F(a)$ is well-defined. We identify integers n_1, n_2, \dots where the oracles will be defined to be nonzero. Define integers $n_1 := 1 + F(a(1))$, $n_\kappa := 1 + \max\{F(a(\kappa)), a(\kappa - 1)(n_{\kappa-1})^{a(\kappa-1)}\}$. For any $n \in \mathbb{N} \setminus \{n_1, n_2, \dots\}$, let \mathcal{O} equal 0 everywhere. For these input lengths, \mathcal{O}_n is trivially a NO instance.

For each $\kappa \in \mathbb{N}$, run $M_{j(\kappa)}$ on input 1^{n_κ} for $a(\kappa)n_\kappa^{a(\kappa)}$ steps and interpret its output as a quantum query circuit \mathcal{A}_{n_κ} which accepts a classical witness. For every query that \mathcal{A}_{n_κ} makes of length $< n_\kappa$, use the previously generated definitions of the oracle \mathcal{O} to hardcode these answers. For queries \mathcal{A}_{n_κ} makes of length $> n_\kappa$, replace the oracle gates with identity circuits. The resulting circuit will be \mathcal{B}_{n_κ} , which only makes queries of length n_κ . This new algorithm \mathcal{B}_{n_κ} can be used to derive an oracle \mathcal{O}_{n_κ} by applying Lemma 6.7.

It remains to prove that no QCMA $^\mathcal{O}$ algorithm exists. Assume, for contradiction, that there exists a P-uniform family of oracle circuits $\{\mathcal{A}_\lambda\}$ that solves the code intersection problem with witnesses of length $t(\lambda) = \text{poly}(\lambda)$ and $Q(\lambda) = \text{poly}(\lambda)$ queries. Then, $\{\mathcal{A}_\lambda\}$ appears in the Turing machine enumeration as some M_{j^*} and there exists some a^* such that $t(\lambda), Q(\lambda) \leq a^*\lambda^{a^*}$. As ι is a surjection, there exists a κ^* such that $\iota(\kappa^*) = (j^*, a^*)$. Let $\mathcal{A}_{n_{\kappa^*}}$ be the quantum circuit for inputs of length n_{κ^*} . Since the oracle is defined as being 0 for inputs $\notin \{n_1, n_2, \dots\}$ and $n_{\kappa^*+1} > a^*n_{\kappa^*}^{a^*}$, each query gate for inputs of length $> n_{\kappa^*}$ is an identity gate. Thus the circuit $\mathcal{B}_{n_{\kappa^*}}$ has the exact same output as $\mathcal{A}_{n_{\kappa^*}}$ on inputs of size n_{κ^*} . However, using Lemma 6.7, we constructed an oracle $\mathcal{O}_{n_{\kappa^*}}$ that $\mathcal{B}_{n_{\kappa^*}}$ will misclassify. Therefore, $\mathcal{A}_{n_{\kappa^*}}$ will answer incorrectly on input $1^{n_{\kappa^*}}$, completing the proof. \square

Proof of Theorem 7.5. Our proof will closely follow [LLPY23]. Suppose that for each λ we generate $(G_\lambda, \mathcal{O}'_\lambda) \leftarrow \mathcal{D}_\lambda$ and define a language $\mathcal{L}^{\mathcal{O}' := \bigsqcup_{\lambda \in \mathbb{N}} G_\lambda^{-1}(1)}$ and an oracle \mathcal{O}' that returns $\mathcal{O}'_{|x|}(x)$ on a query $x \in \{0, 1\}^*$. It suffices to show that $\mathcal{L}^{\mathcal{O}'} \in \text{BQP}^{\mathcal{O}'}/\text{qpoly} \cap \text{NP}^{\mathcal{O}'} \cap \text{coNP}^{\mathcal{O}'}$ and $\mathcal{L}^{\mathcal{O}'} \notin \text{BQP}^{\mathcal{O}'}/\text{poly}$ with probability 1.

To see that $\mathcal{L}^{\mathcal{O}'} \in \text{BQP}^{\mathcal{O}'}/\text{qpoly}$ with probability 1, observe that Item 1 of Lemma 7.4 implies that there is a BQP machine $\mathcal{A}^{\mathcal{O}'}$ with polynomial-size quantum advice that decides $\mathcal{L}^{\mathcal{O}'}$ on all x of length λ with probability at least $1 - \frac{1}{\lambda^2}$

for sufficiently large λ . As $\sum_{\lambda=1}^{\infty} \frac{1}{\lambda^2} = \frac{\pi^2}{6} < \infty$, the Borel–Cantelli lemma (Lemma 4.10) implies that $\mathcal{A}^{\mathcal{O}'}$ decides $\mathcal{L}^{\mathcal{O}'}$ for all but finitely many λ with probability 1. By hard-coding all x 's where \mathcal{A} and \mathcal{L} disagree, \mathcal{A} can be modified into a BQP/qpoly machine $\mathcal{A}'^{\mathcal{O}'}$ that decides $\mathcal{L}^{\mathcal{O}'}$ on all $x \in \{0, 1\}^*$ with probability 1.

In addition, for any $x \in \{0, 1\}^\lambda$, we can give any \mathbf{v} where $(x \| 0^{a-\lambda}, \mathbf{v}) \in R_{C_\lambda, H}$ to certify $G(x)$. Thus, by a Chernoff/union bound, we know that with probability $1 - \text{negl}(\lambda)$ over \mathcal{D}_λ , there exists an NP/coNP certificate for all $x \in \{0, 1\}^\lambda$. We conclude via an identical argument that $\mathcal{L}^{\mathcal{O}'} \in \text{NP}^{\mathcal{O}'} \cap \text{coNP}^{\mathcal{O}'}$ with probability 1.

For a BQP machine \mathcal{B} that takes $\text{poly}(\lambda)$ -bit classical advice, we define $S_{\mathcal{B}}(\lambda)$ to be the event over the choice of (G, \mathcal{O}') that there is a $\text{poly}(\lambda)$ -bit classical advice family $\{z_{\mathcal{O}'}\}_{\mathcal{O}'}$ such that

$$\Pr[\forall x \in \{0, 1\}^\lambda, \mathcal{B}^{\mathcal{O}'}(z_{\mathcal{O}'}, x) = G(x)] \geq \frac{2}{3}.$$

By Item 2 of Lemma 7.4, there exists $\lambda_0 \in \mathbb{N}$ such that for all BQP machines \mathcal{B} , $\Pr_{\mathcal{D}_\lambda}[S_{\mathcal{B}}(\lambda)] \leq \frac{9}{10}$ for all $\lambda \geq \lambda_0$.

We will consider a sequence of input lengths $\lambda_1, \lambda_2, \dots$ defined by $\lambda_i := T(\lambda_{i-1}) + 1$, where $T(\lambda)$ is the running time of \mathcal{B} on input of length λ . This means that when \mathcal{B} 's input length is λ_{i-1} , it cannot query the oracle on input lengths $\geq \lambda_i$, so it must be the case that

$$\begin{aligned} \Pr[S_{\mathcal{B}}(\lambda_i) | S_{\mathcal{B}}(\lambda_0) \wedge \dots \wedge S_{\mathcal{B}}(\lambda_{i-1})] &= \Pr[S_{\mathcal{B}}(\lambda_i)] \\ \implies \Pr[S_{\mathcal{B}}(1) \wedge S_{\mathcal{B}}(2) \wedge \dots] &\leq \Pr\left[\bigwedge_{i=0}^{\infty} S_{\mathcal{B}}(\lambda_i)\right] = \prod_{i=0}^{\infty} \Pr[S_{\mathcal{B}}(\lambda_i) | S_{\mathcal{B}}(\lambda_0) \wedge \dots \wedge S_{\mathcal{B}}(\lambda_{i-1})] \leq \prod_{i=0}^{\infty} \frac{9}{10} = 0. \end{aligned}$$

But there are countably many BQP machines, so $\Pr[\exists \mathcal{B} : S_{\mathcal{B}}(1) \wedge S_{\mathcal{B}}(2) \wedge \dots] = 0$. We conclude that $\mathcal{L}^{\mathcal{O}'} \notin \text{BQP}^{\mathcal{O}'}/\text{poly}$ with probability 1 over the choice of (G, \mathcal{O}') , as desired. \square