Quantum vs. Classical Proofs and Subset Verification

Bill Fefferman
Department of EECS, University of California at Berkeley, Berkeley, CA and
NIST, Gaithersburg, MD, USA
wjf@berkeley.edu

Shelby Kimmel
Department of Computer Science, Middlebury College, Middlebury, VT, USA
skimmel@middlebury.edu

Abstract

We study the ability of efficient quantum verifiers to decide properties of exponentially large subsets given either a classical or quantum witness. We develop a general framework that can be used to prove that QCMA machines, with only classical witnesses, cannot verify certain properties of subsets given implicitly via an oracle. We use this framework to prove an oracle separation between QCMA and QMA using an "in-place" permutation oracle, making the first progress on this question since Aaronson and Kuperberg in 2007 [3]. We also use the framework to prove a particularly simple standard oracle separation between QCMA and AM.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory, Theory of computation → Complexity classes

Keywords and phrases Quantum Complexity Theory, Quantum Proofs

Digital Object Identifier 10.4230/LIPIcs.MFCS.2018.22

Acknowledgements We are grateful for multiple discussions with Stephen Jordan regarding permutation oracle verification strategies. We appreciate the many people who discussed this project with us, including Scott Aaronson, David Gosset, Gus Gutoski, Yi-Kai Liu, Ronald de Wolf, Robin Kothari, Dvir Kafri, and Chris Umans. SK completed much of this work while at the Joint Center for Quantum Information and Computer Science (QuICS), University of Maryland.

1 Introduction

How much computational power does an efficient quantum verifier gain when given a polynomial sized quantum state to support the validity of a mathematical claim? In particular, is there a problem that can be solved in this model, that cannot be solved if the verifier were instead given a classical bitstring? This question, the so-called QMA vs. QCMA problem, is fundamental in quantum complexity theory. To complexity theorists, the question can be motivated simply by trying to understand the power of quantum nondeterminism, where both QMA and QCMA can be seen as "quantum analogues" of NP. More physically, QMA is characterized by the k-local Hamiltonian problem, in which we must decide if the ground state energy of a local Hamiltonian is above or below a specified threshold [16, 5]. In this setting, the QMA vs. QCMA question asks whether there exists a purely *classical* description of the ground state that allows us to make this decision. For instance, if the ground state of any local Hamiltonian can be prepared by an efficient quantum circuit, then QMA = QCMA, as the classical witness for the k-local Hamiltonian problem would be the

classical description of this quantum circuit. It was this intuition that caused Aharonov and Naveh to conjecture that these classes are equal, in the paper that first defined QCMA [5].

It was recently established [12] that the witness to a QMA machine may always be replaced by a subset state, where a subset state on n qubits has the form $|S\rangle = 1/\sqrt{|S|}\sum_{i\in S}|i\rangle$ for some subset $S \subset [2^n]$. However, it seems difficult to create a classical witness on n bits that captures the information in a subset state $|S\rangle$. Therefore, problems involving subsets seem like ripe ground for understanding the QMA vs. QCMA problem. We investigate the following question: under what circumstances is it possible for a quantum machine to verify properties of a subset? This question is not answered by [12]; they study general properties of languages that are in QMA and QCMA, while we attempt to prove specific languages of interest (that are related to subsets) are either in or not in QMA or QCMA.

In the hopes of further exploring these questions, we exhibit a general framework that can be used to obtain oracle separations against QCMA for subset-based problems. We use this framework to prove the existence of an "in-place" permutation oracle $\boldsymbol{\mathcal{P}}$ (a unitary that permutes standard basis states within a single register) [10, 2] for which $\mathsf{QMA}^{\boldsymbol{\mathcal{P}}} \not\subseteq \mathsf{QCMA}^{\boldsymbol{\mathcal{P}}}$, making the first progress on this problem since Aaronson and Kuperberg in 2007 [3], who attained a "quantum oracle" separation (i.e., a separation relative to an *arbitrary* black-box unitary transformation acting on a polynomial number of qubits). In this problem, for the case of QMA, the in-place permutation oracle allows us to verify that the given witness is indeed the correct subset state. On the other hand, our framework allows us to prove the language is not in QCMA. Our framework is quite general, and we are also able to use it to establish a particularly simple example of a (conventional) oracle $\boldsymbol{\mathcal{O}}$ so that $\mathsf{AM}^{\boldsymbol{\mathcal{O}}} \not\subseteq \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.[1]

1.1 Subset-Verifying Oracle Problems

We consider two oracle problems related to verifying properties of subsets. In *Subset Size Checking*, we are given a black box function $f : [N^2] \to \{0, 1\}$, that marks elements with either a 0 or 1. We are promised that the number of marked items is either N or $0.99N$, and we would like to decide which is the case. We want to verify the size of the subset marked by f.

In the other oracle problem, *Preimage Checking*, we are given a black box permutation on N^2 elements. We are promised that the preimage of the first N elements under the permutation is either mostly even or mostly odd, and we would like to decide which is the case. In this problem, we want to verify this parity property of a subset of the preimage of the function.

Subset Size Checking is in AM [11], and we give a procedure that proves *Preimage Checking* is in QMA when the permutation is given as an in-place quantum oracle. An in-place permutation unitary $\mathcal{P}_\sigma$ acts as $\mathcal{P}_\sigma|i\rangle = |\sigma(i)\rangle$ for a permutation σ. For *Preimage Checking*, we are interested in the set $S_{\mathrm{pre}}(\sigma) = \{i : \sigma(i) \in [N]\}$. Given the subset state $|S_{\mathrm{pre}}(\sigma)\rangle$, it is easy to verify that the correct state was sent, because $\mathcal{P}_\sigma|S_{\mathrm{pre}}(\sigma)\rangle = |[N]\rangle$, which is easy to verify using a measurement in the Hadamard basis.

However, we do not expect subset-based oracle languages like *Subset Size Checking* and *Preimage Checking* to be in QCMA because the classical witness does not have enough information to identify the relevant subset. We make this intuition more precise by providing a general recipe for proving that subset-verifying oracle languages are not in QCMA. We

[1] Note there was previously an example of an oracle separating AM from PP [20]. Since QMA $\subseteq$ PP [18], this is formally a stronger result. Nonetheless, our oracle is substantially different, and uses completely different ideas.

apply this procedure to show that both *Preimage Checking* (with a randomized in-place oracle - see Section 3 for more details) and *Subset Size Checking* are not in QCMA. The procedure involves familiar tools, like the adversary bound [6] (although adapted to our in-place oracle when necessary), as well as a new tool, the *Fixing Procedure*, which finds subsets with nice structure within a large arbitrary set. We now sketch the recipe:

1. We show that for every QCMA machine, there are more valid oracles than possible classical witnesses, so by a counting argument, there must be one classical witness w^* that corresponds to a large number of potential oracles. We then restrict ourselves to considering oracles that correspond to w^*.
2. Because we are considering subset-verifying problems, if we have a collection of black box functions that corresponds to w^*, we immediately have some set of subsets that corresponds to w^*. At this point, we know nothing about this set of subsets except its size, thanks to the counting argument. We next show (using the *Fixing Procedure*) that if we have a set of subsets of a certain size, we can always find a subset of the original set that has nice structure.
3. We apply the adversary bound to the subset with nice structure to show that the number of quantum queries needed to distinguish between YES and NO cases is exponential.
4. We finally put these pieces together in a standard diagonalization argument.

1.2 Technical Contributions

Our adversary bound for in-place permutation oracles provides a query lower-bounding technique for unitary oracles when access is *not* given to the inverse of the oracle. (While Belovs [9] created an adversary bound for arbitrary unitaries, his results assume access to an inverse.) While we typically assume quantum oracles include access to an inverse or are self-inverting, in open quantum systems it is natural to not have an inverse.

When proving that *Preimage Checking* is not in QCMA, we use an oracle that is not unitary. The oracle is a completely-positive trace-preserving (CPTP) map that at each application applies one unitary chosen uniformly at random from among a set of unitaries. Standard lower bounding techniques fail for such an oracle. The closest result is from Regev and Schiff [19] who give a lower bound on solving Grover's problem with an oracle that produces errors. Regev and Schiff deal with the non-unitarity of the map by modeling the state of the system using pure states. This strategy does not work in our case. Instead, we use the fact that every non-unitary CPTP map can be implemented as a unitary on a larger system. In our case, we simulate our random oracle using a unitary black box oracle acting on subsystem B, followed by a fixed unitary that entangles subsystems A and B. The entangling operation can not be efficiently implemented, but as we are bounding query complexity, this is acceptable. This technique may be of use for similar problems; for example, we do not know the query complexity of solving Grover's problem with an oracle that produces a depolarizing error with each application. A depolarizing map is similar to our CPTP map in that both maps can be thought of as applying a unitary at random from among a set of unitaries, and so perhaps this approach will stimulate new approaches for the Grover problem.

1.3 Impact and Directions for Future Research

While Aaronson and Kuperberg have previously proved an oracle separation between QMA and QCMA [3], their oracle seems to be especially quantum, as it is defined by a Haar random quantum state. Our in-place oracle has more of a classical feel, in that it encodes

a classical permutation function. However, it is still not a standard quantum oracle, as it is not self-inverting. Is there a standard (i.e. not in-place) oracle language that separates QMA and QCMA? Although we can only prove a separation when our in-place oracle also has randomness, we believe our techniques could be adapted to prove a similar result but without oracle randomness. While we give a recipe for showing certain subset-based problems are not in QCMA, we believe some of these problems are also not in QMA; for example, is it possible to prove *Subset Size Checking* is not in QMA?

Our contributions to techniques for lower bounding query complexity for non-standard oracles raise several questions. Is there a general adversary bound [14, 17] for in-place permutation oracles? There are examples of problems for which in-place permutation oracles require exponentially fewer queries that standard permutation oracles e.g., [7]. We conjecture that there are also examples of problems for which standard permutation oracles require exponentially fewer queries than in-place oracles. In fact, we do not believe it is possible to obtain a Grover-type speed-up with an in-place oracle; we believe the problem of determining the inverse of an element of an in-place permutation oracle requires N queries for a permutation on N elements. However, in order to prove these results, we suspect one needs a more powerful tool, like a general adversary bound for in-place oracles.

1.4 Organization

The rest of the paper is organized as follows, in Section 2, we introduce notation that will be used throughout the rest of the paper, and define QMA and QCMA. In Section 3, we define and discuss standard, in-place, and randomized in-place quantum permutations, as well as state an adversary lower bound for in-place permutation unitaries. In Section 4, we describe the *Preimage Checking* Problem and prove it is in QMA. In Section 5, we lay out the general recipe for proving subset-based languages are not in QCMA. In Section 6, we apply this procedure to the *Subset Size Checking* problem, and use it to prove an oracle separation between AM and QCMA. Finally, in Section 7, we apply the procedure to *Preimage Checking* and show this problem is not in QCMA.

2 Definitions and Notation

We use the notation $[M] = \{1, 2, \ldots, M\}$. σ will refer to a permutation. Unless otherwise specified, the sets we use, generally denoted S, will be a set of positive integers. Also, we will use bold type-face to denote a set of sets. For instance, $\mathbf{S}$ will refer to a set of sets of positive integers. To make our notation clearer, we will refer to such a set of sets as a *set family*. Likewise, we denote $\boldsymbol{\sigma}$ to be a set of permutations acting on the same set of elements.

For S a set of positive integers, a subset state $|S\rangle$ is $|S\rangle = \frac{1}{\sqrt{|S|}} \sum_{i \in S} |i\rangle$.

Throughout, we use $N = 2^n$. All logarithms are in base 2. We use $\boldsymbol{\sigma}^n$ to be the set of permutations acting on N^2 elements. That is, if $\sigma \in \boldsymbol{\sigma}^n$, $\sigma : [N^2] \to [N^2]$. For positive integers $i > j$, let $\mathbf{C}(i, j)$ be the set family containing j elements of $[i]$: $\mathbf{C}(i, j) = \{S \subset [i] : |S| = j\}$.

We use calligraphic font $\mathcal{P}, \mathcal{U}$ to denote unitary operations. We use elaborated calligraphic font $\mathscr{P}$, $\mathscr{U}$ to denote CPTP maps. For a unitary CPTP map $\mathscr{U}$ acting on a density matrix ρ, we have $\mathscr{U}(\rho) = \mathcal{U}\rho\mathcal{U}^\dagger$, where $(\cdot)^\dagger$ denotes conjugate transpose. We will use $\mathcal{O}$ to denote a unitary oracle, and $\mathscr{O}$ to denote a CPTP map oracle.

We include the following standard definition for completeness (e.g., see also [1, 3]).

▶ **Definition 1.** QMA is the set of promise problems $A = (A_{Yes}, A_{No})$ so that there exists an efficient quantum verifier V_A and a polynomial $p(\cdot)$:

1. *Completeness*: For all $x \in A_{Yes}$ there exists a $p(|x|)$-qubit pure quantum state $|\psi\rangle$ so that $\Pr[V_A(x, |\psi\rangle) = 1] \geq 2/3$
2. *Soundness*: For all $x \in A_{No}$ and any pure quantum state $|\psi\rangle$, $\Pr[V_A(x, |\psi\rangle) = 1] \leq 1/3$.

QCMA is the same class, with the witness $|\psi\rangle$ replaced by a polynomial length classical bitstring.

3 Permutation Maps

3.1 Permutations as Oracles: In-Place Permutation vs. Standard Permutation

Black box permutation unitaries have been considered previously, most notably in the collision and element distinctness problems [2, 4]. However, the unitaries considered in these works were standard oracles. A standard oracle implements the permutation $\sigma \in \boldsymbol{\sigma}^n$ as $\mathcal{P}_\sigma^{\text{stand}}|i\rangle|b\rangle = |i\rangle|b \oplus \sigma(i)\rangle$ for $i, b \in [N^2]$, where $|i\rangle$ for $i \in [N^2]$ are standard basis states and $\oplus$ denotes bitwise XOR. Note that $(\mathcal{P}_\sigma^{\text{stand}})^2 = \mathbb{I}_{N^4}$; that is, acting with the unitary twice produces the $N^4 \times N^4$ identity operation.

We consider in-place permutation unitaries, which implement the permutation $\sigma \in \boldsymbol{\sigma}^n$ as $\mathcal{P}_\sigma|i\rangle = |\sigma(i)\rangle$. In general $(\mathcal{P}_\sigma)^2 \neq \mathbb{I}_{N^2}$. Crucially, given black box access to $\mathcal{P}_\sigma$, *we do not give black box access to its inverse.* In fact, in Section 3.2, we show that given only $\mathcal{P}_\sigma$, it is hard to invert its action. Non-self-inverting permutation unitaries have been considered previously, in [10, 2].

We believe standard and in-place permutation unitaries are of incomparable computational power. That is, given one type that implements σ, you can not efficiently simulate the other type implementing the permutation σ. For example, if we have the state $\sum_{y \in S} |y\rangle|\sigma(y)\rangle$ (normalization omitted), we can create the state $\sum_{y \in S} |y\rangle|0\rangle$ with a single query to $\mathcal{P}_\sigma^{\text{stand}}$. However, if we only have access to the in-place permutation $\mathcal{P}_\sigma$ and not to $\mathcal{P}_{(\sigma)^{-1}} = (\mathcal{P}_\sigma)^{-1}$, it seems difficult to create this state.

On the other hand, suppose we want to prepare the state $\sum_{y \in [N]} |\sigma(y)\rangle$ (normalization omitted). We can create this state in one query to the in-place permutation oracle $\mathcal{P}_\sigma$ by applying the oracle to the uniform superposition $\sum_{y \in [N]} |y\rangle$. In the standard model, this problem is called "index erasure," and requires an exponential number of queries in n to $\mathcal{P}_\sigma^{\text{stand}}$ [7].

3.2 An Adversary Bound for In-Place Permutation Oracles

In Appendix A, we prove a non-weighted adversary bound for in-place permutations oracles that is identical to what Ambainis proves in Theorem 6 in [6] for standard permutation oracles.

▶ **Lemma 2.** *Let $\boldsymbol{\sigma} \subset [V] \to [V]$ be a subset of permutations acting on the elements $[V]$. Let $f : \boldsymbol{\sigma} \to \{0, 1\}$ be a function of permutations. Let $\boldsymbol{\sigma}_X \subset \boldsymbol{\sigma}$ be a set of permutations such that if $\sigma \in \boldsymbol{\sigma}_X$, then $f(\sigma) = 1$. Let $\boldsymbol{\sigma}_Y \subset \boldsymbol{\sigma}$ be a permutation family such that if $\sigma \in \boldsymbol{\sigma}_Y$ then $f(\sigma) = 0$. Let $R \subset \boldsymbol{\sigma}_X \times \boldsymbol{\sigma}_Y$ be such that*

- *For every $\sigma_x \in \boldsymbol{\sigma}_X$, there exists at least m different $\sigma_y \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$.*
- *For every $\sigma_y \in \boldsymbol{\sigma}_Y$, there exists at least m' different $\sigma_x \in \boldsymbol{\sigma}_X$ such that $(\sigma_x, \sigma_y) \in R$.*

- *Let $l_{x,i}$ be the number of $\sigma_y \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$ and $\sigma_x(i) \neq \sigma_y(i)$. Let $l_{y,i}$ be the number of $\sigma_x \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$ and $\sigma_x(i) \neq \sigma_y(i)$. Then let $l_{max} = \max_{(\sigma_x,\sigma_y)\in R,i} l_{x,i} l_{y,i}$.*

Then given an in-place permutation oracle $\mathcal{P}_\sigma$ for $\sigma \in \boldsymbol{\sigma}$ that acts as $\mathcal{P}_\sigma|i\rangle = |\sigma(i)\rangle$, any quantum algorithm that correctly evaluates $f(\sigma)$ with probability $1-\epsilon$ for every element of $\boldsymbol{\sigma}_X$ and $\boldsymbol{\sigma}_Y$ must use $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right)\sqrt{\frac{mm'}{l_{max}}}$ queries to the oracle.

As a corollary of Lemma 2, (using exactly the same technique as Theorem 7 in [6]), we have that the query complexity of inverting an in-place permutation oracle on V elements is $\Omega(V^{1/2})$.

3.3 Permutations with Randomness

Additionally, we consider in-place permutation oracles with internal randomness that are CPTP (completely-positive trace-preserving) maps, rather than unitaries. Oracles with internal randomness were shown to cause a complete loss of quantum speed-up in [19], while in [13], such oracles were shown to give an infinite quantum speed-up.

We consider oracles that apply an in-place permutation at random from among a family of possible permutations. Let $\boldsymbol{\sigma} \subseteq \boldsymbol{\sigma}^n$ be a set of $|\boldsymbol{\sigma}|$ permutations. Then the CPTP map $\mathscr{P}_{\boldsymbol{\sigma}}$ acts as follows:

$$\mathscr{P}_{\boldsymbol{\sigma}}(\rho) = \frac{1}{|\boldsymbol{\sigma}|}\sum_{\sigma\in\boldsymbol{\sigma}} \mathcal{P}_\sigma \rho \mathcal{P}_\sigma^\dagger. \tag{1}$$

4 Pre-Image Checking

In this section, we define a property of oracle families which we call *randomized-preimage-correct*, and construct a decision language based on such oracles that is in QMA. Essentially, the problem is to decide whether the preimage of the first N elements of a permutation is mostly even or odd.

Given a permutation $\sigma \in \boldsymbol{\sigma}^n$, we associate a preimage subset $S_{\text{pre}}(\sigma)$ to that permutation ("pre" is for "preimage"), where $S_{\text{pre}}(\sigma) = \{j : \sigma(j) \in [N]\}$. That is, $S_{\text{pre}}(\sigma)$ is the subset of elements in $[N^2]$ whose image under σ is in $[N]$. Additionally, to each subset $S \subseteq [N^2]$ with $|S| = N$, we associate a set of permutations $\boldsymbol{\sigma}_{\text{pre}}(S)$, where $\boldsymbol{\sigma}_{\text{pre}}(S) = \{\sigma : \sigma \in \boldsymbol{\sigma}^n, S_{\text{pre}}(\sigma) = S\}$. Let

$$\begin{aligned} \boldsymbol{S}^n_{\text{even}} &= \{S : S \subset [N^2], |S| = N, |S \cap \mathbb{Z}_{\text{even}}| = 2/3N\} \\ \boldsymbol{S}^n_{\text{odd}} &= \{S : S \subset [N^2], |S| = N, |S \cap \mathbb{Z}_{\text{odd}}| = 2/3N\}. \end{aligned} \tag{2}$$

▶ **Definition 3** (randomized-preimage-correct oracles)**.** Let $\boldsymbol{\mathcal{O}}$ be a countably infinite set of quantum operators (CPTP maps): $\boldsymbol{\mathcal{O}} = \{\mathscr{O}_1, \mathscr{O}_2, \dots\}$, where each $\mathscr{O}_n$ implements an operation on $(2n)$-qubits. We say that $\boldsymbol{\mathcal{O}}$ is randomized-preimage-correct if for every n, $\mathscr{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{pre}(S)}$, with $S \in \boldsymbol{S}^n_{\text{even}} \cup \boldsymbol{S}^n_{\text{odd}}$.

▶ **Theorem 4.** *For any randomized-preimage-correct $\boldsymbol{\mathcal{O}}$, the unary language $L_{\boldsymbol{\mathcal{O}}}$, which contains those unary strings 1^n such that $\mathscr{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{\text{pre}}(S)}$ with $S \in \boldsymbol{S}^n_{even}$, is in $\mathsf{QMA}^{\boldsymbol{\mathcal{O}}}$.*

Proof. We first prove completeness. We assume $1^n \in L_{\boldsymbol{\mathcal{O}}}$, so $\mathscr{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{\text{pre}}(S)}$ for some $S \in \boldsymbol{S}^n_{\text{even}}$. We consider using as a witness the subset state $|S\rangle$ on $2n$ qubits. We analyze the following verifier: with probability $1/2$, do either

Test (i) Apply $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ to $|S\rangle$, and measure whether the resultant state is $|[N]\rangle$. This measurement can be done by applying $H^{\otimes n}$ to the first n qubits, and then measuring all qubits in the standard basis. If the outcome is 0, output 1; otherwise, output 0.

Test (ii) Measure $|S\rangle$ in the standard basis. Let i^* be the resulting standard basis state. If i^* is odd, output 0. Otherwise, apply $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ to $|i^*\rangle$ and measure the resultant standard basis state. If the resultant state is not in $[N]$, output 0; otherwise, output 1.

If Test (i) is implemented, the verifier always outputs 1 because all the permutations that might be applied by $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ transform $|S\rangle$ into $|[N]\rangle$. If Test (ii) is implemented, the verifier outputs 1 with probability 2/3. Averaging over both Tests, the verifier outputs 1 with probability 5/6.

Now we show soundness. Let $1^n \notin L_{\boldsymbol{\mathcal{O}}}$, so $\mathscr{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ for some $S \in \mathbf{S}^n_{\mathrm{odd}}$. Without loss of generality, let the witness be the $2n$-qubit state $|\psi(S)\rangle = \sum_{i=1}^{N^2} \beta_i |i\rangle$. If $p_{(i)}$ (resp. $p_{(ii)}$) is the probability the verifier outputs 1 after performing Test (i) (resp. Test (ii)), then we have

$$p_{(i)} = \frac{1}{N}\left|\sum_{i\in S}\beta_i\right|^2, \qquad p_{(ii)} = \sum_{i\in\mathbb{Z}_{\mathrm{even}}\cap S}|\beta_i|^2. \tag{3}$$

regardless of which permutation the map $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ applies.

Using Cauchy-Schwarz and the triangle inequality, we have $1 \geq \left(\sqrt{3p_{(i)}} + (\sqrt{2}-1)p_{(ii)}\right)/\sqrt{2}$. Thus the total probability that the verifier outputs 1 is

$$\frac{1}{2}\left(p_{(i)} + p_{(ii)}\right) \leq \frac{1}{2}\left(\frac{2}{3}\left(1 - \frac{\sqrt{2}-1}{\sqrt{2}}p_{(ii)}\right)^2 + p_{(ii)}\right). \tag{4}$$

The derivative of the right hand side is positive for $0 \leq p_{(ii)} \leq 1$, so to maximize the right hand side we take $p_{(ii)} = 1$. Doing this, we find the probability that the verifier outputs 1 is at most 2/3. ◀

We will show that the Preimage Checking problem is not in QCMA in Section 7.

Our proof that the Preimage Checking problem is in QMA works equally well for an in-place oracle without randomness. We use the randomness in our oracle in the proof that randomized-preimage-correct languages can not be decided by QCMA. We believe the separation holds even without randomness in the oracle.

5 Strategy for Proving Subset-Based Oracle Languages are not in QCMA

In this section, we describe a general strategy for showing that certain oracle languages are not in QCMA. In particular, we consider the case when the oracles are related to sets of integers:

▶ **Definition 5** (Subset-Based Oracle). Let $\boldsymbol{\mathcal{O}} = \{\mathscr{O}_1, \mathscr{O}_2, \dots\}$ be an oracle such that each $\mathscr{O}_n$ implements a $p_1(n)$-qubit CPTP map from some set of maps $\boldsymbol{\mathcal{O}}^n$. Then we say $\boldsymbol{\mathcal{O}}$ is a subset-based oracle if there exists a set of bijective functions $\{g^1, g^2, \dots\}$ with $g^n : \boldsymbol{\mathcal{O}}^n \to \mathbf{S}^n$ where $\mathbf{S}^n$ is the union of disjoint subset families $\mathbf{S}^n_X$ and $\mathbf{S}^n_Y$.

We also use the following definition:

▶ **Definition 6.** Given a subset family $\mathbf{S}$ containing subsets of positive integers, and $\beta \in \mathbb{R}$ such that $\beta > 0$, we say $\mathbf{S}$ is β-distributed if:

(1) There exists a (possibly empty) set S_{fixed} such that $S_{fixed} \subset S$ for all $S \in \mathbf{S}$.

(2) For every element $i \in \left(\bigcup_{S\in\mathbf{S}} S\right) \setminus S_{fixed}$, i appears in at most a $2^{-\beta}$-fraction of $S \in \mathbf{S}$.

We call S_{fixed} the "fixed subset" of $\mathbf{S}$.

We use the following Recipe for proving a subset-based oracle language is not in QCMA:

▶ **Recipe 1.**

Set-up. *Fix some enumeration over all poly(n)-size quantum verifiers $M_1, M_2, ...$, which we can do because the number of such machines is countably infinite (by the Solovay-Kitaev theorem [15]). Some of these verifiers may try to decide a language by trivially "hardwiring" its outputs; for example, by returning 1 independent of the input. We start by fixing a unary language L such that no machine M_i hardwires the language. We can always do this because there are more languages than poly(n)-sized machines. Then our goal is to associate a subset-based oracle $\boldsymbol{\mathcal{O}} = \{\mathscr{O}_1, \mathscr{O}_2, \dots\}$ with L, such that $1^n \in L$ if and only if $g^n(\mathscr{O}_n) \in \mathbf{S}_X^n$, and to show that even with access to $\boldsymbol{\mathcal{O}}$, no M_i can efficiently decide L for all n.*

Consider the QCMA *machine M_i, and suppose it is given access to a subset-based oracle $\boldsymbol{\mathcal{O}}$, as well as a witness of $p_{M_i}(n)$ bits for each input 1^n. Then for each $\mathscr{O}_n \in \boldsymbol{\mathcal{O}}$ there is some subset of integers $S \in \mathbf{S}^n$ such that $g^n(\mathscr{O}_n) = S$. Since g^n is bijective, S uniquely defines $\mathscr{O}_n$, so the optimal witness that causes M_i to accept $\mathscr{O}_n$ can be thought of as a function of S. Let $w_i(S)$ be the witness that gives the highest probability of success in convincing M_i that $S \in \mathbf{S}_X^n$. Then we denote $\mathbf{S}_{i,wit}(w) = \{S : S \in \mathbf{S}_X^n, w = w_i(S)\}$.*

Using the pigeonhole principle, there exists some string $w_{i,n}$ of $p_{M_i}(n)$ bits such that

$$|\mathbf{S}_{i,wit}(w_{i,n})| \geq \frac{1}{2^{p_{M_i}(n)}} |\mathbf{S}_X^n| . \tag{5}$$

That is, there exists a witness such that a large number of subsets correspond to that witness.

1. *Prove that for $n \geq n_i^*$, there is a subset family $\mathbf{S}_X \subseteq \mathbf{S}_{i,wit}(w_{i,n})$ that is α-distributed with fixed subset S_{fixed}. Let $\mathbf{S}_Y = \{S : S \in \mathbf{S}_Y^n, S_{fixed} \subset S\}$. Show the cardinality of $\mathbf{S}_Y$ is large.*
2. *Create a relation $\mathbf{R} \subseteq \{\mathscr{O} : \mathscr{O} \in \boldsymbol{\mathcal{O}}^n, g(\mathscr{O}) \in \mathbf{S}_X\} \times \{\mathscr{O} : \mathscr{O} \in \boldsymbol{\mathcal{O}}^n, g(\mathscr{O}) \in \mathbf{S}_Y\}$ and use $\mathbf{R}$ to apply an adversary bound to prove a lower bound of $\Omega(N^{\alpha/2}) = \Omega(2^{n\alpha/2})$ on the number of queries M_i requires to distinguish some oracle $\mathscr{O}_{x,n,i} \in \boldsymbol{\mathcal{O}}^n$ such that $g^n(\mathscr{O}_{x,n,i}) \in \mathbf{S}_X^n$ from an oracle $\mathscr{O}_{y,n,i} \in \boldsymbol{\mathcal{O}}^n$ such that $g^n(\mathscr{O}_{y,n,i}) \in \mathbf{S}_Y^n$.*
3. *Apply a standard Baker-Gill-Solovay diagonalization argument [8] to complete the proof. That is, for each M_i, choose a unique $n_i \geq n_i^*$, and if $1^{n_i} \in L$, set $\mathscr{O}_{n_i} = \mathscr{O}_{x,n_i,i}$ and if $1^{n_i} \notin L$, set $\mathscr{O}_{n_i} = \mathscr{O}_{y,n_i,i}$. Then no* QCMA *machine can efficiently decide the language.*

6 Subset Size Checking

In this section, we create a subset-based oracle language $L_{\boldsymbol{\mathcal{O}}}$, such that $L_{\boldsymbol{\mathcal{O}}} \in \mathsf{AM}^{\boldsymbol{\mathcal{O}}}$, but $L_{\boldsymbol{\mathcal{O}}} \notin \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$. We use the strategy of Section 5 to prove $L_{\boldsymbol{\mathcal{O}}} \notin \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.

Let f_S be a function that marks a subset $S \subset [N^2]$. That is $f_S : \{0,1\}^{2n} \to \{0,1\}$, such that $f_S(i) = 1$ if $i \in S$ and 0 otherwise. Let $\mathcal{F}_S$ be the unitary such that $\mathcal{F}_S|i\rangle = (-1)^{f_S(i)}|i\rangle$.

▶ **Definition 7.** Let $\boldsymbol{\mathcal{O}}$ be a countably infinite set of unitaries (resp. boolean functions): $\boldsymbol{\mathcal{O}} = \{\mathcal{O}_1, \mathcal{O}_2, ...\}$, where $\mathcal{O}_n$ implements a $2n$-qubit (resp. bit) unitary (function). We say $\boldsymbol{\mathcal{O}}$ is subset-gapped if for every n, $\mathcal{O}_n = \mathcal{F}_S$ (resp. $\mathcal{O}_n = f_S$) for $|S| = N$ or $|S| = 0.99N$.

Clearly $\boldsymbol{\mathcal{O}}$ is a subset-based oracle (see Definition 5), with $g^n(\mathcal{O}_n) = g^n(\mathcal{F}_S) = S$.

Then the following two lemmas give the desired oracle separation between AM and QCMA:

▶ **Lemma 8.** *For any subset-gapped $\boldsymbol{\mathcal{O}}$, the language $L_{\boldsymbol{\mathcal{O}}}$ that contains those strings 1^n such that $\mathcal{O}_n = f_S$ with $|S| = N$, is in $AM^{\boldsymbol{\mathcal{O}}}$.*

Lemma 8 is proven by Goldwasser and Sipser in [11].

▶ **Lemma 9.** *For any subset-gapped $\boldsymbol{\mathcal{O}}$, the language $L_{\boldsymbol{\mathcal{O}}}$ that contains those strings 1^n such that $\mathcal{O}_n = \mathcal{F}_S$ with $|S| = N$, is not in $\mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.*

To prove this lemma, we follow Recipe 1. We address step 2 of the recipe in Lemma 10:

▶ **Lemma 10.** *Let $0 < \alpha < 1/2$ be a constant and $p(\cdot)$ be a polynomial function. Then there exists a positive integer $n^*(p, \alpha)$, such that for every positive integer $n > n^*(p, \alpha)$, and every subset family $\mathbf{S} \subseteq \mathbf{C}(N^2, N)$ such that $|\mathbf{S}| \geq |\mathbf{C}(N^2, N)|2^{-p(n)}$, there exists a subset family $\mathbf{S}_X \subseteq \mathbf{S}$ such that $\mathbf{S}_X$ is α-distributed with $|S_{fixed}| < .5N$.*

(Since $|S_{\text{fixed}}| < .5N$, this implies $|\{S : S \in \mathbf{S}_Y^n, S_{\text{fixed}} \subset S\}|$ is large, as desired.)

Proof Sketch. (Full proof in Appendix B.) We prove the existence of $\mathbf{S}_X$ by construction. Let $\mathbf{S}$ be any subset of $\mathbf{C}(N^2, N)$ with $|\mathbf{S}| \geq |\mathbf{C}(N^2, N)|2^{-p(n)}$. We construct $\mathbf{S}_X$ using the Fixing Procedure:

Fixing Procedure

1. Set $\mathbf{S}_X = \mathbf{S}$, and set $S_{\text{fixed}} = \emptyset$.
2. **a.** Let $\nu(i)$ be the number of subsets $S \in \mathbf{S}_X$ such that $i \in S$.
   b. If there exists some element i for which $|\mathbf{S}_X| > \nu(i) \geq |\mathbf{S}_X|N^{-\alpha}$, set $\mathbf{S}' \leftarrow \{S : S \in \mathbf{S}_X \text{ and } i \in S\}$, set $S_{\text{fixed}} \leftarrow S_{\text{fixed}} \cup i$, and return to step 2(a). Otherwise exit the Fixing Procedure.

By construction, the Fixing Procedure returns a set that is α-distributed (see Definition 6), so we only need to ensure that not too many elements are fixed. We obtain a lower bound on the final size of $\mathbf{S}_X$ because each time an element is fixed, the size of the set decreases by at most $N^{-\alpha}$. On the other hand, because $\mathbf{S}_X$ is contained in $\mathbf{C}(N^2, N)$, if a certain number of items are fixed, we have an upper bound on the size of $\mathbf{S}_X$ using the structure of $\mathbf{C}(N^2, N)$ and a combinatorial argument. We show that if more than $.5N$ items are fixed, these upper and lower bounds contradict each other, proving that less than $.5N$ items must be fixed before the Fixing Procedure terminates.

We address Step 3 of Recipe 1 with the following Lemma:

▶ **Lemma 11.** *Suppose $\mathbf{S}_X \subset \mathbf{C}(N^2, N)$ is the α-distributed subset created using the Fixing Procedure of Lemma 10, with fixed subset S_{fixed}. Let $\mathbf{S}_Y = \{S : S \in \mathbf{C}(N^2, 0.99N), S_{fixed} \subset S\}$. Then we can construct an adversary bound to prove that for every quantum algorithm G, there exists $S_x \in \mathbf{S}_X$, and $S_y \in \mathbf{S}_Y$, (that depend on G) such that, given oracle access to $\mathcal{F}_{S_x}$ or $\mathcal{F}_{S_y}$, G can not distinguish them with probability $\epsilon > .5$ without using $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right) N^{\alpha/2}$ queries.*

Proof Sketch. (Full proof in Appendix B.) We use Theorem 6 from [6]. This result is identical to our Lemma 2, except with standard oracles rather than permutation oracles.

We let $\mathbf{R} = \mathbf{S}_X \times \mathbf{S}_Y$. To apply Theorem 6, we need to show that for elements i such that $i \in S_x$ but $i \notin S_y$ for $(S_x, S_y) \in \mathbf{R}$ that either (1) S_x is not connected to many other

sets S_y where $i \notin S_y$ or (2) S_y is not connected to many other sets S_x where $i \in S_x$. We use the α-distributed property of $\mathbf{S}_X$ to show that property (2) holds. We show a similar result for the case $i \notin S_x$ but $i \in S_y$ for $(S_x, S_y) \in \mathbf{R}$.

7 Oracle Separation of QMA and QCMA

In this section, we prove an oracle separation between QMA and QCMA. In particular, we show:

▶ **Theorem 12.** *There exists a randomized-preimage-correct oracle $\boldsymbol{\mathcal{O}}$, and a language $L_{\boldsymbol{\mathcal{O}}}$ which contains those unary strings 1^n where $\mathscr{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S)}$ with $S \in \boldsymbol{S}^n_{even}$ such that $L_{\boldsymbol{\mathcal{O}}} \notin \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.*

Combined with Theorem 4, this gives the desired separation between QMA and QCMA.

Really, we would like to prove a different result, one that involves preimage-correct oracles:

▶ **Definition 13** (preimage-correct oracles)**.** Let $\boldsymbol{\mathcal{O}}$ be a countably infinite set of unitaries: $\boldsymbol{\mathcal{O}} = \{\mathcal{O}_1, \mathcal{O}_2, \dots\}$, where each $\mathcal{O}_n$ implements an $(2n)$-qubit unitary. We say that $\boldsymbol{\mathcal{O}}$ is preimage-correct if for every n, $\mathcal{O}_n = \mathcal{P}_\sigma$, for some σ such that $S_{\mathrm{pre}}(\sigma) \in \mathbf{S}^n_{\mathrm{even}} \cup \mathbf{S}^n_{\mathrm{odd}}$.

The definition of preimage-correct oracles is very similar to that of randomized-preimage-correct oracles in Definition 3, except there is no randomness in preimage-correct oracles – they are unitaries. In fact, we believe:

▶ **Conjecture 1.** *There exists a preimage-correct oracle $\boldsymbol{\mathcal{O}}$, and a language $L_{\boldsymbol{\mathcal{O}}}$ which contains those unary strings 1^n where $\mathscr{O}_n = P_\sigma$ with $S_{pre}(\sigma) \in \boldsymbol{S}^n_{even}$, such that $L_{\boldsymbol{\mathcal{O}}} \notin \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.*

Theorem 4 applies equally well whether the oracle is preimage-correct or randomized-preimage-correct. So why is it harder to prove Conjecture 1 than Theorem 12? The answer is that Recipe 1 is much easier to use if the optimal witness depends only on a subset of integers. Note randomized-preimage-correct oracles have a one-to-one relationship with a subset of integers, and so the optimal witness only depends on that subset. However for preimage-correct oracles, the optimal witness might depend on some details of the permutation, which is more challenging to handle.

For convenience, we define the complexity class $\mathsf{QCMA}_{\mathsf{exp,poly}}$ to be the analogue of QCMA, in which the quantum verifier is allowed exponential time and space, but receives a polynomial length classical witness. While trivially bounded-error quantum exponential time, $\mathsf{BQEXP} = \mathsf{QCMA}_{\mathsf{exp,poly}}$, in general the query complexity of a $\mathsf{QCMA}_{\mathsf{exp,poly}}$ machine is not the same as the query complexity of a BQEXP machine.

Our proof works as follows. We first show that if there is a QCMA machine that decides $L_{\boldsymbol{\mathcal{O}}}$, for all randomized-preimage-correct oracles $\boldsymbol{\mathcal{O}}$, then there will be a $\mathsf{QCMA}_{\mathsf{exp,poly}}$ machine that decides $L_{\widetilde{\boldsymbol{\mathcal{O}}}}$ for any preimage-correct oracle $\widetilde{\boldsymbol{\mathcal{O}}}$, where, *crucially*, the optimal witness only depends on the pre-image subset of the permutation implemented by the oracle. Then using Recipe 1, we show that there is a language $L_{\widetilde{\boldsymbol{\mathcal{O}}}}$ for a preimage-correct oracle $\widetilde{\boldsymbol{\mathcal{O}}}$ such that no $\mathsf{QCMA}_{\mathsf{exp,poly}}$ machine that can decide the language using an efficient number of queries to $\widetilde{\boldsymbol{\mathcal{O}}}$ (with a witness that only depends on the pre-image subset). This implies that there is no QCMA machine that solves the randomized-preimage-correct oracle problem.

We first prove the reduction from deciding languages on pre-image correct oracles to languages on randomized pre-image correct oracles.

▶ **Lemma 14.** *Given a randomized-preimage-correct oracle $\boldsymbol{\mathcal{O}}$, let $1^n \in L_{\boldsymbol{\mathcal{O}}}$ if $\mathcal{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{pre}(S)}$ with $S \in \boldsymbol{S}^n_{even}$. Given a preimage-correct oracle $\tilde{\boldsymbol{\mathcal{O}}}$ let $1^n \in L_{\widetilde{\boldsymbol{\mathcal{O}}}}$ if $\mathcal{O}_n = \mathcal{P}_\sigma$ with $S_{pre}(\sigma) \in \mathbf{S}^n_{even}$. Then if there is a* QCMA *machine M that decides $L_{\boldsymbol{\mathcal{O}}}$ for every randomized-preimage-correct $\boldsymbol{\mathcal{O}}$, then there is a* $\mathsf{QCMA}_{\mathsf{exp,poly}}$ *machine $\widetilde{M}$ that decides $L_{\widetilde{\boldsymbol{\mathcal{O}}}}$ for every preimage-correct $\tilde{\boldsymbol{\mathcal{O}}}$ such that $\widetilde{M}$ uses at most a polynomial number of queries to $\tilde{\boldsymbol{\mathcal{O}}}$, and on input 1^n takes as input a classical witness w that depends only on $S_{pre}(\sigma)$.*

Proof Sketch. (Full proof in Appendix C.) Given a permutation σ, we can obtain all permutations σ' such that $S_{\text{pre}}(\sigma') = S_{\text{pre}}(\sigma)$ by first applying σ, and then permuting the first N elements and the last $N^2 - N$ elements separately. Consider a controlled-unitary that, if system A is in state $|i\rangle$, implements the i^{th} in-place permutation of the first N and last $N^2 - N$ elements on system B. If we start with system A in an equal superposition, apply $\mathcal{P}_\sigma$ to B, apply the control to A and B, and then trace out system A, the result is $\mathscr{P}_{S_{\text{pre}}(\sigma)}$ on system A. Thus, given any preimage-correct oracle $\mathcal{P}_\sigma$, we can simulate the randomized-preimage-correct oracle $\mathscr{P}_{S_{\text{pre}}(\sigma)}$.

Using this simulation trick, we can create an algorithm $\tilde{M}$ using a preimage-correct oracle that has the same outcomes as any algorithm M using a randomized-preimage-correct oracle, which uses the oracle the same number of times, and has a witness that only depends on the preimage subset. However, we do not believe the control permutation can be implemented efficiently, and that is why we must consider the class $\mathsf{QCMA}_{\mathsf{exp,poly}}$.

▶ **Lemma 15.** *There exists a preimage-correct $\boldsymbol{\mathcal{O}}$ such that there is no $\mathsf{QCMA}_{\mathsf{exp,poly}}^{\boldsymbol{\mathcal{O}}}$ machine M that decides $L_{\boldsymbol{\mathcal{O}}}$ using a polynomial number of queries, where the classical witness on input 1^n depends only on $S_{pre}(\sigma)$, when $\mathcal{O}_n = \mathcal{P}_\sigma$.*

Note that Lemma 15, combined with the contrapositive of Lemma 14, proves Theorem 12.

To prove Lemma 15, we use Recipe 1. Even though we do not have a true subset-based oracle (the function $g(\mathcal{P}_\sigma) = S_{\text{pre}}(\sigma)$ is not injective), using the constraint that the classical witness depends only on $S_{\text{pre}}(\sigma)$, we can apply the recipe.

Additionally, while Recipe 1 refers to the class QCMA, because we are only making a statement about query complexity (and say nothing about space or time complexity), the approach also applies to the query complexity of $\mathsf{QCMA}_{\mathsf{exp,poly}}$.

We prove steps 2 and 3 of Recipe 1 in Lemmas 16 and 17. These proofs are quite similar to the proofs of Lemmas 10 and 11; the full proofs can be found in Appendix D.

▶ **Lemma 16.** *Let $0 < \alpha < 1/2$ be a constant and $p(\cdot)$ be a polynomial function. Then there exists a positive integer $n^*(p, \alpha)$, such that for every $n > n^*(p, \alpha)$, and every subset family $\mathbf{S} \subset \mathbf{S}^n_{\text{even}}$ such that $|\mathbf{S}| \geq |\mathbf{S}^n_{\text{even}}|2^{-p(n)}$, there exists a subset family $\mathbf{S}_X \subset \mathbf{S}$ such that $\mathbf{S}_X$ is α-distributed. Furthermore the fixed subset S_{fixed} of $\mathbf{S}_X$ contains at most $N/3$ even elements.*

▶ **Lemma 17.** *Let $\mathbf{S}_X$ be the α-distributed set created using the Fixing Procedure from Lemma 16, with fixed subset S_{fixed}. Let $\mathbf{S}_Y = \{S : S \in \mathbf{S}^n_{odd}, S_{fixed} \subset S\}$. Then we can construct an adversary bound to prove that for every quantum algorithm G, there exists permutations $\sigma_x, \sigma_y \in \boldsymbol{\sigma}^n$ with $S_{pre}(\sigma_x) \in \mathbf{S}_X$ and $S_{pre}(\sigma_y) \in \mathbf{S}_Y$, (that depend on G) such that, given oracle access to $\mathcal{P}_{\sigma_x}$ or $\mathcal{P}_{\sigma_y}$, G can not distinguish them with probability $\epsilon > .5$ without using $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right) N^{\alpha/2}$ queries.*

The proof strategy of Lemma 16 (like Lemma 10) involves a Fixing Procedure. However, the details are slightly more complex because we must deal with fixing even and odd elements.

The proof strategy of Lemma 17 is similar to Lemma 11, except we use a more complex relation $\mathbf{R}$ for the adversary bound. The challenge is that for two similar subsets S_x and S_y, there exist permutations σ_x and σ_y that are extremely dissimilar but for which $S_{\text{pre}}(\sigma_x) = S_x$ and $S_{\text{pre}}(\sigma_y) = S_y$. We want to create a relationship $\mathbf{R}$ that connects similar permutations, while only having information about the structure of the related subsets. To address this problem, we note that for any two subsets S_x and S_y, we can create a one-to-one matching between the elements of $\boldsymbol{\sigma}_{\text{pre}}(S_x)$ and the elements of $\boldsymbol{\sigma}_{\text{pre}}(S_y)$ such that each permutation is matched with a similar permutation. Using this one-to-one matching, we create a relationship $\mathbf{R}$ between permutations that inherits the properties of the related subsets.

As an immediate corollary of Theorem 4 and Theorem 12, there exists a randomized-preimage-correct oracle $\boldsymbol{\mathcal{O}}$ and language $L_{\boldsymbol{\mathcal{O}}}$ such that $L \notin \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$ but $L \in \mathsf{QMA}^{\boldsymbol{\mathcal{O}}}$, and so $\mathsf{QMA}^{\boldsymbol{\mathcal{O}}} \not\subseteq \mathsf{QCMA}^{\boldsymbol{\mathcal{O}}}$.

References

1 Complexity zoo. URL: https://complexityzoo.uwaterloo.ca/Complexity_Zoo.

2 Scott Aaronson. Quantum lower bound for the collision problem. In John H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 635–642. ACM, 2002. doi:10.1145/509907.509999.

3 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Computational Complexity, 2007. CCC'07. Twenty-Second Annual IEEE Conference on*, pages 115–128. IEEE, 2007.

4 Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004. doi:10.1145/1008731.1008735.

5 Dorit Aharonov and Tomer Naveh. Quantum NP – a survey. *arXiv preprint quant-ph/0210077*, 2002.

6 Andris Ambainis. Quantum lower bounds by quantum arguments. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 636–643. ACM, 2000.

7 Andris Ambainis, Loïck Magnin, Martin Roetteler, and Jérémie Roland. Symmetry-assisted adversaries for quantum state generation. In *Computational Complexity (CCC), 2011 IEEE 26th Annual Conference on*, pages 167–177. IEEE, 2011.

8 Theodore P. Baker, John Gill, and Robert Solovay. Relativizations of the P =? NP question. *SIAM J. Comput.*, 4(4):431–442, 1975. doi:10.1137/0204037.

9 Aleksandrs Belovs. Variations on quantum adversary. *arXiv preprint 1504.06943*, 2015.

10 J Niel De Beaudrap, Richard Cleve, and John Watrous. Sharp quantum versus classical query complexity separations. *Algorithmica*, 34(4):449–461, 2002.

11 Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In Juris Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 59–68. ACM, 1986. doi:10.1145/12130.12137.

12 Alex B Grilo, Iordanis Kerenidis, and Jamie Sikora. QMA with subset state witnesses. *arXiv preprint arXiv:1410.2882*, 2014.

13 Aram W Harrow and David J Rosenbaum. Uselessness for an oracle model with internal randomness. *Quantum Information & Computation*, 14(7&8):608–624, 2014.

14 Peter Hoyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 526–535. ACM, 2007.

15 Alexei Yu Kitaev. Quantum computation: Algorithms and error correction. *Russian Math. Surveys*, 52(6):1191–1249, 1997.

16 Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. *Classical and quantum computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Soc., 2002. doi:10.1090/gsm/047.

17 Troy Lee, Rajat Mittal, Ben W Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of state conversion. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 344–353. IEEE, 2011.

18 Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Comput. Complex.*, 14(2):122–152, 2005. doi:10.1007/s00037-005-0194-x.

19 Oded Regev and Liron Schiff. Impossibility of a quantum speed-up with a faulty oracle. In *Automata, Languages and Programming*, pages 773–781. Springer, 2008.

20 Nikolay K. Vereshchagin. Oracle separation of complexity classes and lower bounds for perceptrons solving separation problems. *Izvestiya: Mathematics*, 59:1103–1122, 1995. doi:10.1070/IM1995v059n06ABEH000050.

A An Adversary Bound for Permutation Oracles

We will prove Lemma 2:

▶ **Lemma 2.** *Let $\boldsymbol{\sigma} \subset [V] \to [V]$ be a subset of permutations acting on the elements $[V]$. Let $f : \boldsymbol{\sigma} \to \{0,1\}$ be a function of permutations. Let $\boldsymbol{\sigma}_X \subset \boldsymbol{\sigma}$ be a set of permutations such that if $\sigma \in \boldsymbol{\sigma}_X$, then $f(\sigma) = 1$. Let $\boldsymbol{\sigma}_Y \subset \boldsymbol{\sigma}$ be a permutation family such that if $\sigma \in \boldsymbol{\sigma}_Y$ then $f(\sigma) = 0$. Let $R \subset \boldsymbol{\sigma}_X \times \boldsymbol{\sigma}_Y$ be such that*

- *For every $\sigma_x \in \boldsymbol{\sigma}_X$, there exists at least m different $\sigma_y \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$.*
- *For every $\sigma_y \in \boldsymbol{\sigma}_Y$, there exists at least m' different $\sigma_x \in \boldsymbol{\sigma}_X$ such that $(\sigma_x, \sigma_y) \in R$.*
- *Let $l_{x,i}$ be the number of $\sigma_y \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$ and $\sigma_x(i) \neq \sigma_y(i)$. Let $l_{y,i}$ be the number of $\sigma_x \in \boldsymbol{\sigma}_Y$ such that $(\sigma_x, \sigma_y) \in R$ and $\sigma_x(i) \neq \sigma_y(i)$. Then let $l_{max} = \max_{(\sigma_x,\sigma_y)\in R,i} l_{x,i} l_{y,i}$.*

Then given an in-place permutation oracle $\mathcal{P}_\sigma$ for $\sigma \in \boldsymbol{\sigma}$ that acts as $\mathcal{P}_\sigma|i\rangle = |\sigma(i)\rangle$, any quantum algorithm that correctly evaluates $f(\sigma)$ with probability $1-\epsilon$ for every element of $\boldsymbol{\sigma}_X$ and $\boldsymbol{\sigma}_Y$ must use $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right)\sqrt{\frac{mm'}{l_{max}}}$ queries to the oracle.

We note that this is identical to Ambainis' adversary bound for permutations (see Theorem 6 in [6]).

Proof. We assume that we have a control permutation oracle, that acts as

$$\mathcal{P}|x\rangle_C|i\rangle_A|z\rangle_Q = |x\rangle|\sigma_x(i)\rangle|z\rangle \tag{6}$$

where the Hilbert space $\mathcal{H}_C$ has dimension $|\boldsymbol{\sigma}|$,the Hilbert space $\mathcal{H}_A$ has dimension V and is where the permutation is carried out, and $\mathcal{H}_Q$ is a set of ancilla qubits.

Let $|\psi^t\rangle$ be the state of the system immediately after t uses of the control oracle. Let $|\varphi^t\rangle$ be the state of the system immediately before the t^{th} use of the control oracle. Let ρ^t be the reduced state of the system immediately after t uses of the control oracle, where systems A and Q have been traced out. That is, $\rho^t = \mathrm{tr}_{AQ}(|\psi^t\rangle\langle\psi^t|)$. Let $(\rho^t)_{xy}$ be the $(x,y)^{\text{th}}$ element of the density matrix. Then we will track the progress of the following measure:

$$W^t = \sum_{(\sigma_x,\sigma_y)\in R} \left|(\rho^t)_{xy}\right|. \tag{7}$$

Notice that unitaries that only act on the subsystems Q and A do not affect W^t.

If the state before the first use of the oracle is

$$|\psi^0\rangle = \left(\frac{1}{\sqrt{2|\boldsymbol{\sigma}_X|}} \sum_{\sigma_x \in \boldsymbol{\sigma}_X} |x\rangle + \frac{1}{\sqrt{2|\boldsymbol{\sigma}_Y|}} \sum_{\sigma_y \in \boldsymbol{\sigma}_Y} |y\rangle \right) \otimes |\phi\rangle_{AQ}, \tag{8}$$

then following Ambainis (e.g. Theorem 2 [6]), we have that for an algorithm to succeed with probability at least $1-\epsilon$ after T uses of the oracle, we must have

$$W^0 - W^T > \left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right)\sqrt{mm'} \tag{9}$$

Now we calculate how much W^t can change between uses of the oracle. Suppose without loss of generality that

$$|\varphi^t\rangle = \frac{1}{\sqrt{2|\boldsymbol{\sigma}_X|}} \sum_{\sigma_x \in \boldsymbol{\sigma}_X} \sum_{i,z} \alpha_{x,i,z} |x,i,z\rangle_{CAQ} + \frac{1}{\sqrt{2|\boldsymbol{\sigma}_Y|}} \sum_{\sigma_y \in \boldsymbol{\sigma}_Y} \sum_{i,z} \alpha_{y,i,z} |y,i,z\rangle_{CAQ}. \tag{10}$$

Then we have

$$\begin{aligned} |\psi^t\rangle &= \frac{1}{\sqrt{2|\boldsymbol{\sigma}_X|}} \sum_{\sigma_x \in \boldsymbol{\sigma}_X} \sum_{i,z} \alpha_{x,i,z} |x,\sigma_x(i),z\rangle_{CAQ} + \frac{1}{\sqrt{2|\boldsymbol{\sigma}_Y|}} \sum_{\sigma_y \in \boldsymbol{\sigma}_Y} \sum_{i,z} \alpha_{y,i,z} |y,\sigma_y(i),z\rangle_{CAQ} \\ &= \frac{1}{\sqrt{2|\boldsymbol{\sigma}_X|}} \sum_{\sigma_x \in \boldsymbol{\sigma}_X} \sum_{i,z} \alpha_{x,\sigma_x^{-1}(i),z} |x,i,z\rangle_{CAQ} + \frac{1}{\sqrt{2|\boldsymbol{\sigma}_Y|}} \sum_{\sigma_y \in \boldsymbol{\sigma}_Y} \sum_{i,z} \alpha_{y,\sigma_y^{-1}(i),z} |y,i,z\rangle_{CAQ}. \end{aligned} \tag{11}$$

Hence for $(\sigma_x, \sigma_y) \in R$, we have

$$\begin{aligned} (\rho^t)_{xy} &= \frac{1}{2\sqrt{|\boldsymbol{\sigma}_X||\boldsymbol{\sigma}_Y|}} \sum_{i,z} \alpha_{x,\sigma_x^{-1}(i),z} \alpha^*_{y,\sigma_y^{-1}(i),z} \\ (\rho^{t-1})_{xy} &= \frac{1}{2\sqrt{|\boldsymbol{\sigma}_X||\boldsymbol{\sigma}_Y|}} \sum_{i,z} \alpha_{x,i,z} \alpha^*_{y,i,z}, \end{aligned} \tag{12}$$

where $(\cdot)^*$ signifies the complex conjugate. Now we can calculate $W^t - W^{t-1}$:

$$\begin{aligned} W^{t-1} - W^t &= \sum_{(\sigma_x,\sigma_y)\in R} |(\rho^{t-1})_{xy}| - |(\rho^t)_{xy}| \\ &\leq \sum_{(\sigma_x,\sigma_y)\in R} |(\rho^{t-1})_{xy} - (\rho^t)_{xy}|. \end{aligned} \tag{13}$$

From Eq. (12), we see that whenever $\sigma_x^{-1}(i) = \sigma_y^{-1}(i)$, we have a cancellation between the corresponding elements of $(\rho^t)_{xy}$ and $(\rho^{t-1})_{xy}$. However, when $\sigma_x^{-1}(i) \neq \sigma_y^{-1}(i)$, terms do not cancel. To see this more explicitly, we rewrite Eq. (13) as

$$\begin{aligned} W^{t-1} - W^t \leq \\ \frac{1}{2\sqrt{|\boldsymbol{\sigma}_X||\boldsymbol{\sigma}_Y|}} \sum_{(\sigma_x,\sigma_y)\in R} \Bigg| \sum_z \Bigg[\sum_{i:\sigma_x(i)=\sigma_y(i)} \alpha_{x,i,z}\alpha^*_{y,i,z} + \sum_{i:\sigma_x(i)\neq\sigma_y(i)} \alpha_{x,i,z}\alpha^*_{y,i,z} \\ - \sum_{i:\sigma_x^{-1}(i)=\sigma_y^{-1}(i)} \alpha_{x,\sigma_x^{-1}(i),z}\alpha^*_{y,\sigma_y^{-1}(i),z} - \sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} \alpha_{x,\sigma_x^{-1}(i),z}\alpha^*_{y,\sigma_y^{-1}(i),z} \Bigg] \Bigg|. \end{aligned} \tag{14}$$

Consider the sets $T_{x,y} = \{i : \sigma_x(i) = \sigma_y(i)\}$ and $U_{x,y} = \{\sigma_x^{-1}(i) : \sigma_x^{-1}(i) = \sigma_y^{-1}(i)\}$. We will show $U_{x,y} = T_{x,y}$. Suppose $i \in T_{x,y}$. Then $\sigma_x(i) = \sigma_y(i) = i'$, for some i'. But that implies

$\sigma_x^{-1}(i') = \sigma_y^{-1}(i') = i$, so $\sigma_x^{-1}(i') = i \in U_{x,y}$, and thus $T_{x,y} \subset U_{x,y}$. The opposite direction is shown similarly. Therefore, those two sums in Eq. (14) cancel, and, moving the summation over z and i outside the absolute values by the triangle inequality, we are left with

$$W^{t-1} - W^t \leq \frac{1}{2\sqrt{|\boldsymbol{\sigma}_X||\boldsymbol{\sigma}_Y|}} \sum_{z,(\sigma_x,\sigma_y)\in R} \left(\sum_{i:\sigma_x(i)\neq\sigma_y(i)} \left|\alpha_{x,i,z}\alpha^*_{y,i,z}\right| + \sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} \left|\alpha_{x,\sigma_x^{-1}(i),z}\alpha^*_{y,\sigma_y^{-1}(i),z}\right| \right). \tag{15}$$

Now we use the AM-GM to bound the terms in the absolute values:

$$\begin{aligned} W^{t-1} - W^t \leq & \frac{1}{2} \sum_{z,(\sigma_x,\sigma_y)\in R} \left(\sum_{i:\sigma_x(i)\neq\sigma_y(i)} \left(\sqrt{\frac{l_{y,i}}{l_{x,i}}} \frac{|\alpha_{x,i,z}|^2}{2|\boldsymbol{\sigma}_X|} + \sqrt{\frac{l_{x,i}}{l_{y,i}}} \frac{|\alpha^*_{y,i,z}|^2}{2|\boldsymbol{\sigma}_Y|} \right) \right) \\ & + \frac{1}{2} \sum_{z,(\sigma_x,\sigma_y)\in R} \left(\sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} \left(\sqrt{\frac{l_{y,i}}{l_{x,i}}} \frac{\left|\alpha_{x,\sigma_x^{-1}(i),z}\right|^2}{2|\boldsymbol{\sigma}_X|} + \sqrt{\frac{l_{x,i}}{l_{y,i}}} \frac{\left|\alpha^*_{y,\sigma_y^{-1}(i),z}\right|^2}{2|\boldsymbol{\sigma}_Y|} \right) \right) \end{aligned} \tag{16}$$

We now show that for $(\sigma_x, \sigma_y) \in R$,

$$\begin{aligned} \sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} |\alpha_{x,\sigma_x^{-1}(i),z}|^2 &= \sum_{i:\sigma_x(i)\neq\sigma_y(i)} |\alpha_{x,i,z}|^2, \\ \sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} |\alpha_{y,\sigma_y^{-1}(i),z}|^2 &= \sum_{i:\sigma_x(i)\neq\sigma_y(i)} |\alpha_{y,i,z}|^2. \end{aligned} \tag{17}$$

We prove the first equality, and the second is proven similarly. We define

$$\begin{aligned} T'_{x,y} &= [V] \setminus T_{x,y}, \\ U'_{x,y} &= [V] \setminus U_{x,y}. \end{aligned} \tag{18}$$

Looking at the definition of $T_{x,y}$ and $U_{x,y}$, we see that

$$\begin{aligned} T'_{x,y} &= \{i : \sigma_x(i) \neq \sigma_y(i)\} \\ U'_{x,y} &= \{\sigma_x^{-1}(i) : \sigma_x^{-1}(i) \neq \sigma_y^{-1}(i)\}. \end{aligned} \tag{19}$$

We previously showed $T_{x,y} = U_{x,y}$, so we have $T'_{x,y} = U'_{x,y}$. Therefore

$$\sum_{i:\sigma_x^{-1}(i)\neq\sigma_y^{-1}(i)} |\alpha_{x,\sigma_x^{-1}(i),z}|^2 = \sum_{i:U'_{x,y}} |\alpha_{x,j,z}|^2 = \sum_{i:T'_{x,y}} |\alpha_{x,i,z}|^2 \sum_{i:\sigma_x(i)\neq\sigma_y(i)} |\alpha_{x,i,z}|^2. \tag{20}$$

Thus, Eq. (16) becomes

$$W^{t-1} - W^t \leq \sum_{z,(\sigma_x,\sigma_y)\in R} \left(\sum_{i:\sigma_x(i)\neq\sigma_y(i)} \left(\sqrt{\frac{l_{y,i}}{l_{x,i}}} \frac{|\alpha_{x,i,z}|^2}{2|\boldsymbol{\sigma}_X|} + \sqrt{\frac{l_{x,i}}{l_{y,i}}} \frac{|\alpha^*_{y,i,z}|^2}{2|\boldsymbol{\sigma}_Y|} \right) \right). \tag{21}$$

Now we switch the order of summation and then use the definition of $l_{x,i}$ and $l_{y,i}$ to get

$$\begin{aligned} & W^{t-1} - W^t \\ \leq & \sum_{i\in[V],z} \left(\sum_{(\sigma_x,\sigma_y)\in R:\sigma_x(i)\neq\sigma_y(i)} \left(\sqrt{\frac{l_{y,i}}{l_{x,i}}} \frac{|\alpha_{x,i,z}|^2}{2|\boldsymbol{\sigma}_X|} + \sqrt{\frac{l_{x,i}}{l_{y,i}}} \frac{|\alpha^*_{y,i,z}|^2}{2|\boldsymbol{\sigma}_Y|} \right) \right) \\ \leq & \sum_{i\in[V],z} \left(\sum_{\sigma_x\in\boldsymbol{\sigma}_X} \sqrt{l_{x,i} \max_{\sigma_y:(\sigma_x,\sigma_y)\in R} l_{y,i}} \frac{|\alpha_{x,i,z}|^2}{2|\boldsymbol{\sigma}_X|} + \sum_{\sigma_y\in\boldsymbol{\sigma}_Y} \sqrt{l_{y,i} \max_{\sigma_x:(\sigma_x,\sigma_y)\in R} l_{x,i}} \frac{|\alpha^*_{y,i,z}|^2}{2|\boldsymbol{\sigma}_Y|} \right) \end{aligned} \tag{22}$$

Finally, using the definition of l_{max} we have

$$\begin{aligned} W^{t-1} - W^t \leq & \sqrt{l_{max}} \sum_{i\in[V],z} \left(\sum_{x\in\boldsymbol{\sigma}_X} \frac{|\alpha_{x,i,z}|^2}{2|\boldsymbol{\sigma}_X|} + \sum_{\sigma_y\in\boldsymbol{\sigma}_Y} \frac{|\alpha^*_{y,i,z}|^2}{2|\boldsymbol{\sigma}_Y|} \right) \\ \leq & \sqrt{l_{max}}, \end{aligned} \tag{23}$$

where we have used that Eq. (10) is a normalized state. ◀

B Proofs of Lemmas 10 and 11

▶ **Lemma 10.** *Let $0 < \alpha < 1/2$ be a constant and $p(\cdot)$ be a polynomial function. Then there exists a positive integer $n^*(p,\alpha)$, such that for every positive integer $n > n^*(p,\alpha)$, and every subset family $\mathbf{S} \subseteq \mathbf{C}(N^2, N)$ such that $|\mathbf{S}| \geq |\mathbf{C}(N^2,N)|2^{-p(n)}$, there exists a subset family $\mathbf{S}_X \subseteq \mathbf{S}$ such that $\mathbf{S}_X$ is α-distributed with $|S_{fixed}| < .5N$.*

Proof. We prove the existence of $\mathbf{S}'$ by construction. Let $\mathbf{S}$ be any subset family of $\mathbf{C}(N^2,N)$ such that $|\mathbf{S}| \geq |\mathbf{C}(N^2,N)|2^{-p(n)}$. We construct $\mathbf{S}'$ using the following procedure:

Fixing Procedure

1. Set $\mathbf{S}' = \mathbf{S}$, and set $S_{\text{fixed}} = \emptyset$.

2. a. Let $\nu(i)$ be the number of subsets $S \in \mathbf{S}_X$ such that $i \in S$.

b. If there exists some element i for which

$$|\mathbf{S}_X| > \nu(i) \geq |\mathbf{S}_X| N^{-\alpha} \tag{24}$$

set $\mathbf{S}' \leftarrow \{S : S \in \mathbf{S}_X \text{ and } i \in S\}$, set $S_{\text{fixed}} \leftarrow S_{\text{fixed}} \cup i$, and return to step 2(a). Otherwise exit the Fixing Procedure.

The Fixing Procedure by construction will always return a set that satisfies Definition 6. Now we just need to bound the size of S_{fixed}.

Let's suppose that at some point in the Fixing Procedure, for sets $\mathbf{S}'$ and S_{fixed}, we have $.5N$ items fixed. Suppose for contradiction there is some element $i^* \notin S_{\text{fixed}}$ that appears in greater than $N^{-\alpha}$ fraction of $S \in \mathbf{S}'$.

Let us look at the set family $\mathbf{S}'' = \{S : S \in \mathbf{S}', i^* \in S\}$. Because $(S_{\text{fixed}} \cup i^*) \subset S$ for all $S \in \mathbf{S}''$, there are $.5N - 1$ elements in each $S \in \mathbf{S}''$ that can be chosen freely from the remaining $N^2 - .5N - 1$ un-fixed elements. Thus, we have

$$|\mathbf{S}''| \leq \binom{N^2 - .5N - 1}{.5N - 1}. \tag{25}$$

By assumption $|\mathbf{S}''| \geq |\mathbf{S}'|N^{-\alpha}$, so

$$\begin{aligned}
|\mathbf{S}'| &\leq \binom{N^2 - .5N - 1}{.5N - 1} N^{\alpha} \\
&\leq \binom{N^2}{.5N} N^{\alpha} \\
&\leq (2Ne)^{N/2} N^{\alpha} \\
&= 2^{N/2(\log(2e)+\log N)+\log(N)\alpha} \\
&= 2^{O(N)+(N/2)\log N}. \qquad (26)
\end{aligned}$$

However, we can also bound the size of $\mathbf{S}'$ from the Fixing Procedure. Notice that at every step of the Fixing Procedure, the size of $\mathbf{S}'$ is reduced by at most a factor $N^{-\alpha}$. Since we are assuming $.5N$ elements are in S_{fixed}, the Fixing Procedure can reduce the original set $\mathbf{S}$ by at most a factor $N^{-\alpha N/2}$. Since $|\mathbf{S}| \geq |\mathbf{C}(N^2, N)|2^{-p(n)}$, we have that at this point in the Fixing Procedure

$$\begin{aligned}
|\mathbf{S}'| &\geq |\mathbf{C}(N^2, N)|2^{-p(n)} N^{-\alpha N/2} \\
&= \binom{N^2}{N} 2^{-p(n)} N^{-\alpha N/2} \\
&\geq N^N 2^{-p(n)} N^{-\alpha N/2} \\
&= 2^{N\log N - p(n) - \log(N)\alpha N/2} \\
&= 2^{-O(N)+N\log N(1-\alpha/2)}. \qquad (27)
\end{aligned}$$

Notice that as long as $\alpha < 1$, for large enough N (in particular, for $N > 2^{n^*}$ for some positive integer n^*, where n^* depends on α and $p(\cdot)$), the bound of Eq. (27) will be larger than the bound of Eq. (26), giving a contradiction. Therefore, our assumption must have been false, and more than $N/2$ elements can not have been fixed during the Fixing Procedure. Therefore, the final set produced by the Fixing Procedure will satisfy point (2) of Definition 6. ◀

▶ **Lemma 11.** *Suppose $\mathbf{S}_X \subset \mathbf{C}(N^2, N)$ is the α-distributed subset created using the Fixing Procedure of Lemma 10, with fixed subset S_{fixed}. Let $\mathbf{S}_Y = \{S : S \in \mathbf{C}(N^2, 0.99N), S_{fixed} \subset S\}$. Then we can construct an adversary bound to prove that for every quantum algorithm G, there exists $S_x \in \mathbf{S}_X$, and $S_y \in \mathbf{S}_Y$, (that depend on G) such that, given oracle access to $\mathcal{F}_{S_x}$ or $\mathcal{F}_{S_y}$, G can not distinguish them with probability $\epsilon > .5$ without using $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right) N^{\alpha/2}$ queries.*

Proof. Note $\mathbf{S}_Y$ is non-empty, since only $.5N$ elements are in S_{fixed}.

We will use Theorem 6 from [6]. This result is identical to our Lemma 2, except with standard oracles rather than permutation oracles. We define the relation $\mathbf{R}$ as:

$$\mathbf{R} = \{(S_x, S_y) : S_x \in \mathbf{S}_X, S_y \in \mathbf{S}_Y\}. \qquad (28)$$

Notice that each $S_x \in \boldsymbol{S}_X$ is paired with every element of $\mathbf{S}_Y$. Thus $m = |\mathbf{S}_Y|$. Likewise $m' = |\mathbf{S}_X|$.

Now consider $(S_x, S_y) \in \mathbf{R}$. We first consider the case of some element j such that $j \in S_x$ but $j \notin S_y$. By our construction of $\mathbf{S}_Y$, $j \notin S_{\text{fixed}}$. We upper bound $l_{x,j}$, the number of $S_{y'}$ such that $(S_x, S_{y'}) \in \mathbf{R}$ and $j \notin S_{y'}$. We use the trivial upper bound $l_{x,j} \leq |\mathbf{S}_Y|$, which is sufficient for our purposes. Next we need to upper bound $l_{y,j}$, the number of $S_{x'}$ such that

$(S_{x'}, S_y) \in \mathbf{R}$ and $j \in S_x$. Since S_y is paired with every element of $\mathbf{S}_X$ in $\mathbf{R}$, we just need to determine the number of sets in $\mathbf{S}_X$ that contain j. Because $\mathbf{S}_X$ is α-distributed, there can be at most $N^{-\alpha}|\mathbf{S}_X|$ elements of $\mathbf{S}_X$ that contain j. In this case we have

$$l_{x,j} l_{y,j} \leq |\mathbf{S}_X||\mathbf{S}_Y| N^{-\alpha}. \tag{29}$$

We now consider the case that $j \in S_y$ but $j \notin S_x$. (Note this case only occurs when S_{fixed} contains less than $0.99N$ elements.) We upper bound $l_{y,j}$, the number of $S_{x'}$ such that $(S_{x'}, S_y) \in \mathbf{R}$ and $j \notin S_{x'}$. Again, we use the trivial upper bound of $l_{y,j} \leq |\mathbf{S}_X|$, which is sufficient for our analysis. Next we upper bound $l_{x,j}$, the number of $S_{y'}$ such that $(S_x, S_{y'}) \in \mathbf{R}$ and $j \in S_{y'}$. In our choice of $\mathbf{R}$, S_x is paired with every $S_y \in \mathbf{S}_Y$, so we need to count the number of $S \in \mathbf{S}_Y$ that contain j. We have

$$\begin{aligned} l_{x,j} =& \binom{N^2 - S_{\text{fixed}} - 1}{0.99N - S_{\text{fixed}} - 1} \\ =& \frac{0.99N - S_{\text{fixed}}}{N^2 - S_{\text{fixed}}} |\mathbf{S}_Y| \\ \leq& \frac{1}{N} |\mathbf{S}_Y|. \end{aligned} \tag{30}$$

Therefore in this case, we have

$$l_{x,j} l_{y,j} \leq |\mathbf{S}_X||\mathbf{S}_Y| N^{-1}. \tag{31}$$

Looking at Eq. (29) and Eq. (31), we see that because $\alpha < 1$, the bound of Eq. (29) dominates, and so we have that

$$\sqrt{\frac{mm'}{l_{x,j} l_{y,j}}} \geq \sqrt{\frac{|\mathbf{S}_X||\mathbf{S}_Y|}{|\mathbf{S}_X||\mathbf{S}_Y| N^{-\alpha}}} = N^{\alpha/2}. \tag{32}$$

Using the contrapositive of Lemma 2, if an algorithm G makes less than q queries to an oracle $\mathcal{F}_S$ where S is promised to be in $\mathbf{S}_X$ or $\mathbf{S}_Y$, there exists at least one element of $\mathbf{S}_X$ and one element of $\mathbf{S}_Y$ such that the probability of distinguishing between the corresponding oracles less than is $1/2 + \epsilon$, where

$$\frac{1}{2} \sqrt{2 N^{-\alpha/2} q} > \epsilon. \tag{33}$$

Equivalently, there exists at least one element of $\boldsymbol{S}_X$ and one element of $\boldsymbol{S}_Y$ such that in order for $\mathcal{A}$ to distinguish them with constant bias, one requires $\Omega(N^{\alpha/2})$ queries. ◀

C Proof of Lemma 14

▶ **Lemma 14.** *Given a randomized-preimage-correct oracle $\boldsymbol{\mathcal{O}}$, let $1^n \in L_{\boldsymbol{\mathcal{O}}}$ if $\mathcal{O}_n = \mathscr{P}_{\boldsymbol{\sigma}_{pre}(S)}$ with $S \in \boldsymbol{S}^n_{even}$. Given a preimage-correct oracle $\tilde{\boldsymbol{\mathcal{O}}}$ let $1^n \in L_{\tilde{\boldsymbol{\mathcal{O}}}}$ if $\mathcal{O}_n = \mathcal{P}_\sigma$ with $S_{pre}(\sigma) \in \mathbf{S}^n_{even}$. Then if there is a* QCMA *machine M that decides $L_{\boldsymbol{\mathcal{O}}}$ for every randomized-preimage-correct $\boldsymbol{\mathcal{O}}$, then there is a* $\mathsf{QCMA}_{\mathsf{exp,poly}}$ *machine $\widetilde{M}$ that decides $L_{\tilde{\boldsymbol{\mathcal{O}}}}$ for every preimage-correct $\tilde{\boldsymbol{\mathcal{O}}}$ such that $\widetilde{M}$ uses at most a polynomial number of queries to $\tilde{\boldsymbol{\mathcal{O}}}$, and on input 1^n takes as input a classical witness w that depends only on $S_{pre}(\sigma)$.*

Proof. We denote the composition of two CPTP maps with $\circ$, so $\mathscr{E} \circ \mathscr{F}$ means apply $\mathscr{F}$ first, and then $\mathscr{E}$.

For each input 1^n, M applies an algorithm that takes as input a standard basis state. Because S completely characterizes $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}(S)}}$, the optimal witness will depend only on S.

Suppose on input 1^n to M, the algorithm is the following:

$$\mathscr{L}_{AB} \circ (\mathscr{O})_A \circ (\mathscr{U}_t)_{AB} \circ \cdots \circ (\mathscr{U}_2)_{AB} \circ (\mathscr{O})_A \circ (\mathscr{U}_1)_{AB}(|w\rangle\langle w| \otimes |\psi_0\rangle\langle\psi_0|)_{AB} \tag{34}$$

where $|w\rangle\langle w|$ is the witness state (that depends only on S) in the standard basis and $\mathscr{U}_i$ are fixed unitaries and $\mathscr{L}$. The two subspaces A and B refer to the subset where the oracle acts (A) and the rest of the workspace (B). The two subspaces do *not* refer to the tensor product structure of the initial state.

For $i \in [N!(N^2-N)!]$ let $\boldsymbol{\tau}^n = \{\tau_i\}$ be the set of permutations on the elements of $[N^2]$ that do not mix the first N elements with the last $N^2 - N$ elements. Then let $\mathcal{P}_n^{\mathrm{C}}$ be the following control-permutation:

$$\mathcal{P}_n^{\mathrm{C}}|i\rangle|j\rangle = \begin{cases} |i\rangle|\tau_i(j)\rangle \text{ for } i \in [N!(N^2-N)!] \\ |i\rangle|j\rangle \text{ otherwise.} \end{cases} \tag{35}$$

$\mathscr{P}_n^{\mathrm{C}}$ is the respective CPTP map.

$\mathcal{P}_n^{\mathrm{C}}$ is a completely known unitary that is independent of the oracle, however, we do not know how to implement this unitary in polynomial time. This unitary is the reason we consider the class $\mathsf{QCMA}_{\mathsf{exp,poly}}$ in this proof rather than the more standard QCMA. Ultimately, we care about query complexity - not the complexity of the unitaries that occur between the oracle applications.

Let

$$|\chi_n\rangle = \frac{1}{\sqrt{N!(N^2-N)!}} \sum_{i=1}^{N!(N^2-N)!} |i\rangle \tag{36}$$

Then on input 1^n we have $\tilde{M}$ implement the algorithm

$$\begin{aligned} &\mathscr{L}_{AB} \circ (\mathscr{P}_n^{\mathrm{C}})_{C_t A} \circ (\mathscr{O})_A \circ (\mathscr{U}_t)_{AB} \circ \cdots \\ &\quad \circ (\mathscr{U}_2)_{AB} \circ (\mathscr{P}_n^{\mathrm{C}})_{C_1 A} \circ (\mathscr{O})_A \circ (\mathscr{U}_1)_{AB} \left(|\chi_n\rangle\langle\chi_n|_C^t \otimes (|w\rangle\langle w| \otimes |\psi_0\rangle\langle\psi_0|)_{AB}\right) \end{aligned} \tag{37}$$

where $(\mathscr{P}_\tau^{\mathrm{C}})_{C_i A}$ means the $C_i{}^{\mathrm{th}}$ register controls the A^{th} register, and initially, the $C_i{}^{\mathrm{th}}$ register is the i^{th} copy of $|\chi_n\rangle$, and $\mathscr{O}$ is the CPTP version of the oracle $\mathcal{O}$.

Let $\rho_i(\mathscr{O})$ (resp. $\tilde{\rho}_i(\mathcal{O})$) be the state of the system during the algorithm M (resp. $\tilde{M}$) after the i^{th} use of the oracle. Let $\rho_0(\mathscr{O})$ (resp. $\tilde{\rho}_0(\mathcal{O})$) be the initial state of the respective algorithms. Then we will show that

$$\rho_i(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}(S_{\mathrm{pre}}(\sigma))}}) = \mathrm{tr}_{C_1,\ldots,C_t}\left(\tilde{\rho}_i(\mathcal{P}_\sigma)\right). \tag{38}$$

As a consequence of this, the probability distribution of measurement outcome of the two algorithms will be identical.

We prove this by induction. For the initial step, we have

$$\rho_0(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}(S_{\mathrm{pre}}(\sigma))}}) = |w\rangle\langle w| \otimes |\psi_0\rangle\langle\psi_0| \tag{39}$$

while

$$\begin{aligned} \mathrm{tr}_C(\tilde{\rho}_0(\mathcal{P}_\sigma)) &= \mathrm{tr}_C\left(|\chi_n\rangle\langle\chi_n|_C^t \otimes (|w\rangle\langle w| \otimes |\psi_0\rangle\langle\psi_0|)_{AB}\right) \\ &= |w\rangle\langle w| \otimes |\psi_0\rangle\langle\psi_0|. \end{aligned} \tag{40}$$

For the induction step, we need to show

$$\rho_k(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}) = \mathrm{tr}_{C_1,\ldots,C_t}\left(\tilde{\rho}_k(\mathcal{P}_\sigma)\right). \tag{41}$$

Because $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}$ has an equal probability of applying $\mathcal{P}_\sigma$ for each σ such that $S(\sigma) = S$, we have

$$\begin{aligned}\rho_k(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}) =& \mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}\left(\mathcal{U}_k\rho_{k-1}(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))})\mathcal{U}_k^\dagger\right)\\ =& \frac{1}{N!(N^2-N)!}\sum_{i=1}^{N!(N^2-N)!}\mathcal{P}_{\tau_i}\mathcal{P}_\sigma\mathcal{U}_k\rho_{k-1}(\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))})\mathcal{U}_k^\dagger\mathcal{P}_\sigma^\dagger\mathcal{P}_{\tau_i}^\dagger. \end{aligned}\tag{42}$$

On the other hand

$$\begin{aligned}\mathrm{tr}_C\left(\tilde{\rho}_k(\mathcal{P}_\sigma)\right) =& \mathrm{tr}_C\left((\mathcal{P}_\tau^{\mathrm{C}})_{C_kA}\mathcal{P}_\sigma\mathcal{U}_k(\tilde{\rho}_{k-1}(\mathcal{P}_\sigma))\mathcal{U}_k^\dagger\mathcal{P}_\sigma^\dagger(\mathcal{P}_\tau^{\mathrm{C}})_{C_kA}^\dagger\right)\\ =& \frac{1}{N!(N^2-N)!}\sum_{i=1}^{N!(N^2-N)!}\mathcal{P}_{\tau_i}\mathcal{P}_\sigma\mathcal{U}_k\,\mathrm{tr}_C\left(\tilde{\rho}_{k-1}(\mathcal{P}_\sigma)\right)\mathcal{U}_k^\dagger\mathcal{P}_\sigma^\dagger\mathcal{P}_{\tau_i}^\dagger \end{aligned}\tag{43}$$

Now we need to show $\tilde{M}$ decides $L_{\boldsymbol{\mathcal{O}}}$ for a preimage-correct oracle $\boldsymbol{\mathcal{O}}$. Let's consider an input 1^n. Suppose $\mathcal{O}_n = \mathcal{P}_\sigma$, where $S_{\mathrm{pre}}(\sigma) \in \mathbf{S}^n_{\mathrm{even}}$. Then because M decides $L_{\boldsymbol{\mathcal{O}}}$ for any randomized-preimage-correct, there exists a witness w that depends only on $S_{\mathrm{pre}}(\sigma)$ such that when the oracle is $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}$ the output of M is 1 with probability at least $2/3$. Using the same witness w, $\tilde{M}$ will therefore produce output 1 with probability at least $2/3$.

Now consider an input 1^n such that $\mathcal{O}_n = \mathcal{P}_\sigma$ where $S_{\mathrm{pre}}(\sigma) \in \mathbf{S}^n_{\mathrm{odd}}$. Because M decides $L_{\boldsymbol{\mathcal{O}}}$ for a randomized-preimage-correct oracle $\boldsymbol{\mathcal{O}}$, when M is run with the oracle $\mathscr{P}_{\boldsymbol{\sigma}_{\mathrm{pre}}(S_{\mathrm{pre}}(\sigma))}$, for any witness w, M will output 1 with probability at most $1/3$. But because $\tilde{M}$ will have the same probability distribution of outcomes, this means that for any witness w to $\tilde{M}$, with oracle $\mathcal{P}_\sigma$, $\tilde{M}$ will output 1 with probability at most $1/3$. ◀

D Proofs of Lemmas 16 and 17

▶ **Lemma 16.** *Let $0 < \alpha < 1/2$ be a constant and $p(\cdot)$ be a polynomial function. Then there exists a positive integer $n^*(p,\alpha)$, such that for every $n > n^*(p,\alpha)$, and every subset family $\mathbf{S} \subset \mathbf{S}^n_{\mathrm{even}}$ such that $|\mathbf{S}| \geq |\mathbf{S}^n_{\mathrm{even}}|2^{-p(n)}$, there exists a subset family $\mathbf{S}_X \subset \mathbf{S}$ such that $\mathbf{S}_X$ is α-distributed. Furthermore the fixed subset S_{fixed} of $\mathbf{S}_X$ contains at most $N/3$ even elements.*

Proof. We prove the existence of $\mathbf{S}'$ by construction. Let $\mathbf{S}$ be any subset of $\mathbf{S}^n_{\mathrm{even}}$ such that $|\mathbf{S}| \geq |\mathbf{S}^n_{\mathrm{even}}|2^{-p(n)}$. We construct $\mathbf{S}_X$ using the following procedure, which we call the fixing procedure.

Fixing Procedure

1. Set $\mathbf{S}_X = \mathbf{S}$, and set $S_{\mathrm{fixed}} = \emptyset$.
2. **a.** Let $\nu(i)$ be the number of subsets $S \in \mathbf{S}_X$ such that $i \in S$.
   b. If there exists some element i for which

$$|\mathbf{S}_X| > \nu(i) \geq |\mathbf{S}_X|N^{-\alpha} \tag{44}$$

set $\mathbf{S}_X \leftarrow \{S : S \in \mathbf{S}_X \text{ and } i \in S\}$, set $S_{\mathrm{fixed}} \leftarrow S_{\mathrm{fixed}} \cup i$, and return to step 2(a). Otherwise exit the Fixing Procedure.

By construction, $\mathbf{S}_X$ will satisfy Definition 6. Now we need to check that the Fixing Procedure stops before fixing more than $N/3$ even items.

Let's suppose that at some point in the Fixing Procedure, for sets $\mathbf{S}_X$ and S_{fixed}, we have $N/3$ even items fixed. Suppose for contradiction, that at this point, there is some even element i^* such that i^* appears in greater than $N^{-\alpha}$ fraction of $S \in \mathbf{S}_X$. Let's also assume without loss of generality that $|S_{\text{fixed}} \cap \mathbb{Z}_{\text{odd}}| = k_{\text{odd}} \leq N/3$.

Let us look at the set

$$\mathbf{S}'' = \{S : S \in \mathbf{S}_X, i^* \in S\}. \tag{45}$$

Because $(S_{\text{fixed}} \cup i^*) \subset S$ for all $S \in \mathbf{S}''$, there are $N/3 - 1$ even elements that can be freely chosen for $S \in \mathbf{S}''$ and $N/3 - k_{\text{odd}}$ odd elements that can be freely chosen. Thus, we have

$$|\mathbf{S}''| \leq \binom{N^2/2 - N/3 - 1}{N/3 - 1}\binom{N^2/2 - k_{\text{odd}}}{N/3 - k_{\text{odd}}}. \tag{46}$$

By assumption

$$|\mathbf{S}''| \geq |\mathbf{S}_X| N^{-\alpha}, \tag{47}$$

so

$$\begin{aligned}
|\mathbf{S}_X| &\leq \binom{N^2/2 - N/3 - 1}{N/3 - 1}\binom{N^2/2 - k_{\text{odd}}}{N/3 - k_{\text{odd}}} N^{\alpha} \\
&\leq \binom{N^2/2}{N/3}\binom{N^2/2}{N/3} N^{\alpha} \\
&\leq (3Ne/2)^{2N/3} N^{\alpha} \\
&= 2^{2N/3(\log(3e/2) + \log N) + \log(N)\alpha} \\
&= 2^{O(N) + (2N/3)\log N}.
\end{aligned} \tag{48}$$

However, we can also bound the size of $\mathbf{S}_X$ from the Fixing Procedure. Notice that at every step of the Fixing Procedure, the size of $\mathbf{S}_X$ is reduced by at most a factor $N^{-\alpha}$. Since we are assuming $N/3$ even elements are in S_{fixed} and $k_{\text{odd}} \leq N/3$ odd elements are in S_{fixed}, the Fixing Procedure can reduce the original set family $\mathbf{S}$ by at most a factor $N^{-\alpha(2N/3)}$. Since $|\mathbf{S}| \geq |\mathbf{S}^n_{\text{even}}| 2^{-p(n)}$, we have that at this point in the Fixing Procedure

$$\begin{aligned}
|\mathbf{S}_X| &\geq |\mathbf{S}^n_{\text{even}}| 2^{-p(n)} N^{-\alpha(2N/3)} \\
&= \binom{N^2/2}{2N/3}\binom{N^2/2}{N/3} 2^{-p(n)} N^{-\alpha(2N/3)} \\
&\leq (3N/4)^{2N/3} (3N/4)^{N/3} 2^{-p(n)} N^{-\alpha(2N/3)} \\
&= 2^{2N/3(\log(3/4) + \log N) + N/3(\log(3/4) + \log N) - p(n) - \log(N)\alpha 2N/3} \\
&= 2^{N \log N(1 - 2\alpha/3) + N\log(3/4) - p(\log(N))} \\
&= 2^{-O(N) + N\log N(1 - 2\alpha/3)}.
\end{aligned} \tag{49}$$

Notice that as long as $\alpha < 1/2$, for large enough N (in particular, for $N > 2^{n^*}$ for some positive integer n^*, where n^* depends on α and $p(\cdot)$), the bound of Eq. (49) will be larger than the bound of Eq. (48), giving a contradiction. Therefore, our assumption, must have been false, and at this point in the Fixing Procedure, all even elements $i \in [N^2]/S_{\text{fixed}}$ will appear in at most a fraction $N^{-\alpha}$ of $S \in \mathbf{S}_X$. Thus at the next step of the Fixing Procedure, an even element will not be added to S_{fixed}, and the number of even elements in S_{fixed} will stay bounded by $N/3$. The same logic can be reapplied at future steps of the Fixing Procedure, even if additional odd items are added. ◀

▶ **Lemma 17.** *Let $\mathbf{S}_X$ be the α-distributed set created using the Fixing Procedure from Lemma 16, with fixed subset S_{fixed}. Let $\mathbf{S}_Y = \{S : S \in \mathbf{S}^n_{odd}, S_{fixed} \subset S\}$. Then we can construct an adversary bound to prove that for every quantum algorithm G, there exists permutations $\sigma_x, \sigma_y \in \boldsymbol{\sigma}^n$ with $S_{pre}(\sigma_x) \in \mathbf{S}_X$ and $S_{pre}(\sigma_y) \in \mathbf{S}_Y$, (that depend on G) such that, given oracle access to $\mathcal{P}_{\sigma_x}$ or $\mathcal{P}_{\sigma_y}$, G can not distinguish them with probability $\epsilon > .5$ without using $\left(1 - 2\sqrt{\epsilon(1-\epsilon)}\right) N^{\alpha/2}$ queries.*

Proof. Since $\mathbf{S}_X$ is α-distributed, there exists a set S_{fixed} of elements such that $S_{\text{fixed}} \subset S$ for all $S \in \mathbf{S}_X$, where S_{fixed} contains at most $N/3$ odd elements and at most $N/3$ even elements. (Otherwise Condition (2) of Definition 6 will not be satisfied.) We choose

$$\begin{aligned}\boldsymbol{\sigma}_Y &= \{\sigma : S_{\text{pre}}(\sigma) \in \mathbf{S}_Y\},\\ \boldsymbol{\sigma}_X &= \{\sigma : S_{\text{pre}}(\sigma) \in \mathbf{S}_X\}.\end{aligned} \tag{50}$$

We now define the relation $\mathbf{R}$ needed to apply our adversary bound. For each $(S_x, S_y) \in \mathbf{S}_X \times \mathbf{S}_Y$, we will create a one-to-one matching in $\mathbf{R}$ between the elements of $\boldsymbol{\sigma}_{\text{pre}}(S_x)$ and $\boldsymbol{\sigma}_{\text{pre}}(S_y)$. We first choose any element $\sigma^*_x \in \boldsymbol{\sigma}_{\text{pre}}(S_x)$. Then we choose a permutation $\sigma^*_y \in \boldsymbol{\sigma}_{\text{pre}}(S_y)$ such that

- $\forall j \in (S_x \cap S_y), \sigma^*_x(j) = \sigma^*_y(j)$,
- $\forall j \in [N^2] \setminus (S_x \cup S_y), \sigma^*_x(j) = \sigma^*_y(j)$,
- $\forall j \in S_x \setminus (S_x \cap S_y), \exists i \in S_y \setminus (S_x \cap S_y)$ such that $\sigma^*_x(j) = \sigma^*_y(i)$ and $\sigma^*_x(i) = \sigma^*_y(j)$.

Since every permutation corresponding to S_y is in $\boldsymbol{\sigma}_{\text{pre}}(S_y)$, there will always be such a σ^*_y that satisfies the above criterion. We choose $(\sigma^*_x, \sigma^*_y) \in \mathbf{R}$.

For $i \in [N!(N^2 - N)!]$ let $\boldsymbol{\tau}^n = \{\tau_i\}$ be the set of permutations on the elements of $[N^2]$ that do not mix the first N elements with the last $N^2 - N$ elements. By $\sigma_a \circ \sigma_b$, we mean apply first permutation σ_b, and then permutation σ_a. Notice that

$$\begin{aligned}\boldsymbol{\sigma}_{\text{pre}}(S_x) =& \{\tau \circ \sigma^*_x : \tau \in \boldsymbol{\tau}^n\}\\ \boldsymbol{\sigma}_{\text{pre}}(S_y) =& \{\tau \circ \sigma^*_y : \tau \in \boldsymbol{\tau}^n\}.\end{aligned} \tag{51}$$

Furthermore given $\tau \in \boldsymbol{\tau}^n$, we have

- $\forall j \in (S_x \cap S_y), \tau \circ \sigma^*_x(j) = \tau \circ \sigma^*_y(j)$,
- $\forall j \in [N^2] \setminus (S_x \cup S_y), \tau \circ \sigma^*_x(j) = \tau \circ \sigma^*_y(j)$,
- $\forall j \in S_x \setminus (S_x \cap S_y), \exists i \in S_y \setminus (S_x \cap S_y)$ such that $\tau \circ \sigma^*_x(j) = \tau \circ \sigma^*_y(i)$ and $\sigma \circ \sigma^*_x(i) = \sigma \circ \sigma^*_y(j)$.

For every $\tau \in \boldsymbol{\tau}^n$, we set $(\tau \circ \sigma^*_x, \tau \circ \sigma^*_y) \in \mathbf{R}$. In doing so, we create a one-to-one correspondance in $\mathbf{R}$ between the elements of $\boldsymbol{\sigma}_{\text{pre}}(S_x)$ and $\boldsymbol{\sigma}_{\text{pre}}(S_y)$. We then repeat this process for all pairs $(S_x, S_y) \in \mathbf{S}_X \times \mathbf{S}_Y$. The end result is the $\mathbf{R}$ that we will use.

Now we need to analyze the properties of this $\mathbf{R}$. Notice that each $\sigma_x \in \boldsymbol{\sigma}_X$ is paired to exactly one element of $\boldsymbol{\sigma}_{\text{pre}}(S_y)$ for each $S_y \in \mathbf{S}_Y$. Thus $m = |\mathbf{S}_Y|$. Likewise $m' = |\mathbf{S}_X|$.

Now consider $(\sigma_x, \sigma_y) \in \mathbf{R}$. We consider some element j such that $\sigma_x(j) \neq \sigma_y(j)$. We first consider the case that $j \in S_x$. We upper bound $l_{x,j}$, the number of $\sigma_{y'}$ such that $(\sigma_x, \sigma_{y'}) \in \mathbf{R}$ and $\sigma_{y'}(j) \neq \sigma_x(j)$. To simplify analysis, we use the simple upper bound $l_{x,j} \leq |\mathbf{S}_Y|$, which is sufficient for our purposes. Next we need to upper bound $l_{y,j}$, the number of $\sigma_{x'}$ such that $(\sigma_{x'}, \sigma_y) \in \mathbf{R}$ and $\sigma_{x'}(j) \neq \sigma_y(j)$. By our construction of $\mathbf{R}$, we have $j \notin S_y$. Also, by construction, if $j \notin S_y$, $\sigma_{x'}(j) \neq \sigma_y(j)$ if and only if $j \in S_{x'}$. Since σ_y is paired to only one element σ_x for each set S_x, $l_{y,j}$ is bounded by the number of sets

$S_x \in \mathbf{S}_X$ such that $j \in S_x$. Because $\mathbf{S}_X$ is α-distributed, at most a fraction $N^{-\alpha}$ of the sets of $\mathbf{S}_X$ can contain j, so $l_{y,j} \leq |\mathbf{S}_X| N^{-\alpha}$. In this case we have

$$l_{x,j} l_{y,j} \leq |\mathbf{S}_X||\mathbf{S}_Y| N^{-\alpha}. \tag{52}$$

We now consider the case that $j \in S_y$. We upper bound $l_{y,j}$, the number of $\sigma_{x'}$ such that $(\sigma_{x'}, \sigma_y) \in \mathbf{R}$ and $\sigma_{x'}(j) \neq \sigma_y(j)$. To simplify analysis, we use the upper bound of $l_{y,j} \leq |\mathbf{S}_X|$, which is sufficient for our analysis. Next we need to upper bound $l_{x,j}$, the number of $\sigma_{y'}$ such that $(\sigma_x, \sigma_{y'}) \in \mathbf{R}$ and $\sigma_{y'}(j) \neq \sigma_x(j)$. By our construction of $\mathbf{R}$, we have $j \notin S_x$. Also, by construction, if $j \notin S_x$, $\sigma_{y'}(j) \neq \sigma_x(j)$ if and only if $j \in S_{y'}$. Since σ_x is paired to only one element σ_x for each set S_x, $l_{x,j}$ is bounded by the number of sets $S_y \in \mathbf{S}_Y$ such that $j \in S_y$. Suppose S_{fixed} contains k_{even} even elements and k_{odd} odd elements. If j is odd, we have

$$\begin{aligned} l_{x,j} &= \binom{N^2/2 - k_{\text{odd}} - 1}{2N/3 - k_{\text{odd}} - 1} \binom{N^2/2 - k_{\text{even}}}{N/3 - k_{\text{even}}} \\ &\leq \frac{2N/3}{N^2/2 - N/3} |\mathbf{S}_Y|, \end{aligned} \tag{53}$$

while if j is even (in that case, we must have $k_{\text{even}} < N/3$), we have

$$\begin{aligned} l_{x,j} &= \binom{N^2/2 - k_{\text{odd}}}{2N/3 - k_{\text{odd}}} \binom{N^2/2 - k_{\text{even}} - 1}{N/3 - k_{\text{even}} - 1} \\ &\leq \frac{N/3}{N^2/2 - N/3} |\mathbf{S}_Y|, \end{aligned} \tag{54}$$

where we've used that

$$|\mathbf{S}_Y| = \binom{N^2/2 - k_{\text{odd}}}{2N/3 - k_{\text{odd}}} \binom{N^2/2 - k_{\text{even}}}{N/3 - k_{\text{even}}}. \tag{55}$$

Therefore in this case, we have

$$l_{x,j} l_{y,j} = |\mathbf{S}_X||\mathbf{S}_Y| O(N^{-1}). \tag{56}$$

Looking at Eq. (52) and Eq. (56), we see that because $\alpha < 1$, the bound of Eq. (52) dominates, and so we have that

$$\sqrt{\frac{mm'}{l_{x,j} l_{y,j}}} \geq \sqrt{\frac{|\mathbf{S}_X||\mathbf{S}_Y|}{|\mathbf{S}_X||\mathbf{S}_Y| N^{-\alpha}}} = N^{\alpha/2}. \tag{57}$$

Using the contrapositive of Lemma 2, if an algorithm G makes less than q queries to an oracle $\mathcal{O}_{\sigma_x}$ where σ_x is promised to be in $\boldsymbol{\sigma}_X$ or $\boldsymbol{\sigma}_Y$, there exists at least one element of $\boldsymbol{\sigma}_X$ and one element of $\boldsymbol{\sigma}_Y$ such that the probability of distinguishing between the corresponding oracles less than is $1/2 + \epsilon$, where

$$\frac{1}{2} \sqrt{2 N^{-\alpha/2} q} > \epsilon. \tag{58}$$

Equivalently, there exists at least one element of $\boldsymbol{\sigma}_X$ and one element of $\boldsymbol{\sigma}_Y$ such that in order for $\mathcal{A}$ to distinguish them with constant bias, one requires $\Omega(N^{\alpha/2})$ queries. ◀