

Post-Quantum Security of the Even-Mansour Cipher

Gorjan Alagic¹, Chen Bai², Jonathan Katz³, and Christian Majenz⁴

¹ QuICS, University of Maryland, and NIST

² Dept. of Electrical and Computer Engineering, University of Maryland

³ Dept. of Computer Science, University of Maryland

⁴ Department of Applied Mathematics and Computer Science, Technical University of Denmark

Abstract. The Even-Mansour cipher is a simple method for constructing a (keyed) pseudorandom permutation E from a public random permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$. It is secure against classical attacks, with optimal attacks requiring q_E queries to E and q_P queries to P such that $q_E \cdot q_P \approx 2^n$. If the attacker is given *quantum* access to both E and P , however, the cipher is completely insecure, with attacks using $q_E, q_P = O(n)$ queries known.

In any plausible real-world setting, however, a quantum attacker would have only *classical* access to the keyed permutation E implemented by honest parties, even while retaining quantum access to P . Attacks in this setting with $q_E \cdot q_P^2 \approx 2^n$ are known, showing that security degrades as compared to the purely classical case, but leaving open the question as to whether the Even-Mansour cipher can still be proven secure in this natural, “post-quantum” setting.

We resolve this question, showing that any attack in that setting requires $q_E \cdot q_P^2 + q_P \cdot q_E^2 \approx 2^n$. Our results apply to both the two-key and single-key variants of Even-Mansour. Along the way, we establish several generalizations of results from prior work on quantum-query lower bounds that may be of independent interest.

1 Introduction

The Even-Mansour cipher [11] is a well-known approach for constructing a block cipher E from a public random permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$. The cipher $E : \{0, 1\}^{2n} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined as

$$E_{k_1, k_2}(x) = P(x \oplus k_1) \oplus k_2$$

where, at least in the original construction, k_1, k_2 are uniform and independent. Security in the standard (classical) setting is well understood [11, 9]: roughly, an unbounded attacker with access to P and P^{-1} cannot distinguish whether it is interacting with E_{k_1, k_2} and E_{k_1, k_2}^{-1} (for uniform k_1, k_2) or R and R^{-1} (for an independent, random permutation R) unless it makes $\approx 2^{n/2}$ queries to its oracles. A variant where k_1 is uniform and $k_2 = k_1$ has the same security [9]. These bounds are tight, and key-recovery attacks using $O(2^{n/2})$ queries are known [11, 9].

Unfortunately, the Even-Mansour construction is insecure against a fully quantum attack in which the attacker is given *quantum* access to its oracles [20, 17]. In such a setting, the adversary can evaluate the unitary operators

$$\begin{aligned} U_P &: |x\rangle|y\rangle \mapsto |x\rangle|y \oplus P(x)\rangle \\ U_{E_{k_1, k_2}} &: |x\rangle|y\rangle \mapsto |x\rangle|y \oplus E_{k_1, k_2}(x)\rangle \end{aligned}$$

(and the analogous unitaries for P^{-1} and E_{k_1, k_2}^{-1}) on any quantum state it prepares. Here, Simon’s algorithm [22] can be applied to $E_{k_1, k_2} \oplus P$ to give a key-recovery attack using only $O(n)$ queries.

To place this seemingly devastating attack in context, it is worth recalling that the original motivation for considering unitary oracles of the form above in quantum-query complexity was that one can always transform a classical circuit for a function f into a reversible (and hence unitary) quantum circuit for U_f .

In a cryptographic context, it is thus reasonable (indeed, necessary) to consider adversaries that use U_f whenever f is a function whose circuit they know. On the other hand, if the circuit for f is *not* known to the adversary, then there is no mechanism by which it can implement U_f on its own. In particular, if f involves a private key, then the only way an adversary could possibly obtain quantum access to f would be if there were an explicit interface granting such access. In most (if not all) real-world applications, however, the honest parties using the keyed function f would implement f using a classical computer; even if they implement f on a quantum computer, there is no reason for them to support anything but a classical interface to f . In such cases, an adversary would have no way to evaluate the unitary operator corresponding to f .

In most real-world applications of Even-Mansour, therefore, an attacker would have only *classical* access to the keyed permutation E_{k_1, k_2} and its inverse, while retaining quantum access to P and P^{-1} . In particular, this seems to be the “right” attack model for most applications of the resulting block cipher, e.g., constructing a secure encryption scheme from the cipher using some mode of operation. The setting in which the attacker is given classical oracle access to keyed primitives but quantum access to public primitives is sometimes called the “Q1 setting” [5]; we will refer to it simply as the *post-quantum* setting.

Security of the Even-Mansour cipher in this setting is currently unclear. Kuwakado and Morii [20] showed a key-recovery attack on Even-Mansour in this setting that requires only $\approx 2^{n/3}$ oracle queries, using the BHT collision-finding algorithm [7]. Their attack uses exponential memory but this was improved in subsequent work [14,5], culminating in an attack using the same number of queries but with polynomial memory complexity. While these results demonstrate that the Even-Mansour construction is *quantitatively* less secure in the post-quantum setting than in the classical setting, and show post-quantum security in certain restricted settings, they do not answer the more important *qualitative* question of whether the Even-Mansour construction remains secure as a block cipher in the post-quantum setting, or whether attacks using polynomially many queries might be possible.

Concurrent results of Jaeger et al. [16] imply security of a forward-only variant of the Even-Mansour construction, as well as for the full Even-Mansour cipher against *non-adaptive* adversaries who make all their classical queries before any quantum queries. They explicitly leave open the question of proving adaptive security in the latter setting.

1.1 Our Results

As our main result, we prove a lower bound showing that $\approx 2^{n/3}$ queries are *necessary* for attacking the Even-Mansour cipher in the post-quantum setting. In more detail, if q_P denotes the number of (quantum) queries to P, P^{-1} and q_E denotes the number of (classical) queries to $E_{k_1, k_2}, E_{k_1, k_2}^{-1}$, we show that any attack succeeding with constant probability requires either $q_P^2 \cdot q_E = \Omega(2^n)$ or $q_P \cdot q_E^2 = \Omega(2^n)$. (Equating q_P and q_E gives the claimed result.) Formally:

Theorem 1. *Let \mathcal{A} be a quantum algorithm making q_E classical queries to its first oracle (including forward and inverse queries) and q_P quantum queries to its second oracle (including forward and inverse queries.) Then*

$$\left| \Pr_{k_1, k_2, P} [\mathcal{A}^{E_{k_1, k_2}, P}(1^n) = 1] - \Pr_{R, P} [\mathcal{A}^{R, P}(1^n) = 1] \right| \leq 10 \cdot 2^{-n/2} \cdot (q_E \sqrt{q_P} + q_P \sqrt{q_E}),$$

where P, R are uniform n -bit permutations, and the marginal distributions of $k_1, k_2 \in \{0, 1\}^n$ are uniform.

The above applies, in particular, to the two-key and one-key variants of the cipher. A simplified version of the proof works also for the case where P is a random function, we consider the cipher $E_k(x) = P(x \oplus k)$ with k uniform, and \mathcal{A} is given forward-only access to both P and E .

Real-world attackers are usually assumed to make far fewer queries to keyed, “online” primitives than to public, “offline” primitives. (Indeed, while an offline query is just a local computation, an online query requires, e.g., causing an honest user to encrypt a certain message.) In such a regime, where $q_E \ll q_P$, the

bound on the adversary’s advantage in [Theorem 1](#) simplifies to $O(q_P\sqrt{q_E}/2^{n/2})$. In that case $q_P^2q_E = \Omega(2^n)$ is necessary for constant success probability, which matches the BHT and offline Simon algorithms [20,5].¹

Techniques and new technical results. Proving [Theorem 1](#) required us to develop new techniques that we believe are interesting beyond our immediate application. We describe the main challenge and its resolution in what follows.

As we have already discussed, in the setting of post-quantum security adversaries may have a combination of classical and quantum oracles. In addition to the Even-Mansour setting [16], this is the case, in particular, when a post-quantum security notion that involves keyed oracles is analyzed in the quantum random oracle model (QROM), such as when analyzing the Fujisaki-Okamoto transform [23,13,4,26,19,8] or the Fiat-Shamir transform [24,18,12]. In general, dealing with a mix of quantum and classical oracles presents a problem: quantum-query lower bounds typically begin by “purifying” the adversary and postponing all measurements to the end of its execution, but this does not work if the adversary may decide what query to make to a classical oracle (or even whether to query a oracle at all) *based on the outcome* of an intermediate measurement. The works cited above address this problem in various ways (e.g. by specializing to non-adaptive adversaries [16]), but often do so by relaxing the problem and allowing quantum access to *all* oracles. This is not an option for us if we wish to prove security, because the Even-Mansour cipher is known not to be secure when the adversary has quantum access to all its oracles! The only other work we are aware of that solves this problem is the concurrent work [16]. Here, the authors overcome the described barrier for the forward-only Even-Mansour (see appendix A) using Zhandry’s compressed oracle technique (which is not available for inverse-accessible permutations). Like previous works they delay all measurements, enforcing the classical-query nature of the adversary in a different way.

Instead, we deal with the problem by dividing the execution of an algorithm that has classical access to some oracle O_1 and quantum access to another oracle O_2 into *stages*, where a stage corresponds to a period between classical queries to O_1 . We then analyze the algorithm stage-by-stage. In doing so, however, we introduce another problem: the adversary may adaptively choose the number of queries to O_2 in each stage based on outcomes of intermediate measurements. While it is possible to upper bound the number of queries to O_2 in each stage by the number of queries made to O_2 overall, this will (in general) result in a very loose security bound. To avoid such a loss, we extend the “blinding lemma” of Alagic et al. [1] so that (in addition to some other generalizations) we obtain a bound in terms of the *expected* number of queries made by a distinguisher.

Lemma 1 (Arbitrary reprogramming, informal). *Consider the following game played by a distinguisher \mathcal{D} making at most q queries in expectation.*

- Phase 1:** \mathcal{D} outputs a function F and a randomized algorithm \mathcal{B} that specifies how to reprogram F .
- Phase 2:** Randomness r is sampled and $\mathcal{B}(r)$ is run to reprogram F , giving F' . A uniform $b \in \{0,1\}$ is chosen, and \mathcal{D} receives oracle access to either F (if $b = 0$) or F' (if $b = 1$).
- Phase 3:** \mathcal{D} loses access to its oracle and receives r ; \mathcal{D} outputs a bit b' .

Then $|\Pr[\mathcal{D} \text{ outputs } 1 \mid b = 0] - \Pr[\mathcal{D} \text{ outputs } 1 \mid b = 1]| \leq 2q\sqrt{\epsilon}$, where ϵ is an upper bound on the probability that any given input is reprogrammed.

The name “arbitrary reprogramming” is motivated by the facts that F is arbitrary (and known), and the adversary can reprogram F arbitrarily—so long as some bound on the probability of reprogramming each individual input exists.

We also extend the “adaptive reprogramming lemma” of Grilo et al. [12] to the case of two-way-accessible, random permutations:

Lemma 2 (Resampling for permutations, informal). *Consider the following game played by a distinguisher \mathcal{D} making at most q queries.*

¹ While our bound is tight with respect to the number of queries, it is loose with regard to the attacker’s advantage, as both the BHT and offline Simon algorithms achieve advantage $\Theta(q_P^2q_E/2^n)$. Reducing this gap is an interesting open question.

Phase 1: \mathcal{D} makes at most q (forward or inverse) quantum queries to a uniform permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$.

Phase 2: A uniform $b \in \{0, 1\}$ is chosen, and \mathcal{D} is allowed to make an arbitrary number of queries to an oracle that is either equal to P (if $b = 0$) or P' (if $b = 1$), where P' is obtained from P by swapping the output values at two uniform points (which are given to \mathcal{D} .) Then \mathcal{D} outputs a bit b' .

Then $|\Pr[\mathcal{D} \text{ outputs } 1 \mid b = 0] - \Pr[\mathcal{D} \text{ outputs } 1 \mid b = 1]| \leq 4\sqrt{q} \cdot 2^{-n/2}$.

This is tight up to a constant factor (cf. [12, Theorem 7]). The name “resampling lemma” is motivated by the fact that here reprogramming is restricted to resampling output values from the *same* distribution used to initially sample outputs of P . While Lemma 1 allows for more general resampling, Lemma 2 gives a bound that is independent of the number of queries \mathcal{D} makes after the reprogramming occurs.

Implications for a variant of the Hidden Shift problem. In the well-studied Hidden Shift problem [25], one is asked to find an unknown shift s by querying an oracle for a (typically injective) function f on a group G and an oracle for its shift $f_s(x) = f(x \cdot s)$. If both oracles are classical, this problem has query complexity superpolynomial in $\log |G|$. If both oracles are quantum, then the query complexity is polynomial [10] but the algorithmic difficulty appears to depend critically on the structure of G (e.g., while $G = \mathbb{Z}_2^n$ is easy [22], $G = S_n$ appears to be intractable [2]).

The obvious connection between the Hidden Shift problem and security of Even-Mansour in general groups has been considered before [2,15,6]. In our case, it leads us to define two natural variants of the Hidden Shift problem:

1. “post-quantum” Hidden Shift: the oracle for f is quantum while the oracle for f_s is classical;
2. “two-sided” Hidden Shift: in place of f_s , use $f_{s_1, s_2}(x) = f(x \cdot s_1) \cdot s_2$; if f is a permutation, grant access to f^{-1} and f_{s_1, s_2}^{-1} as well.

These two variants can be considered jointly or separately and, for either variant, one can consider worst-case or average-case settings [2]. Our main result implies:

Theorem 2 (informal). *Solving the post-quantum Hidden Shift problem on any group G requires a number of queries that is superpolynomial in $\log |G|$. This holds for both the one-sided and two-sided versions of the problem, and for both the worst-case and the average-case settings.*

Theorem 2 follows from the proof of Theorem 1 via a few straightforward observations. First, an inspection of the proof shows that the particular structure of the underlying group (i.e., the XOR operation on $\{0, 1\}^n$) is not relevant; the proof works identically for any group, simply replacing 2^n with $|G|$ in the bounds. The two-sided case of Theorem 2 then follows almost immediately: worst-case search is at least as hard as average-case search, and average-case search is at least as hard as average-case decision, which is precisely Theorem 1 (with the appropriate underlying group). Finally, as noted earlier, an appropriate analogue of Theorem 1 also holds in the “forward-only” case where $E(x) = P(x \oplus k)$ and P is a random function. This yields the one-sided case of Theorem 2.

1.2 Paper Organization

In Section 2 we state the technical lemmas needed for our main result. In Section 3 we prove Theorem 1, showing post-quantum security of the Even-Mansour cipher (both the two-key and one-key variants), based on the technical lemmas. In Section 4 we prove the technical lemmas themselves. Finally, in Appendix A, we give a proof of post-quantum security for the one-key, “forward-only” variant of Even-Mansour. While this is a relatively straightforward adaptation of the proof of our main result, it does not follow directly from it; moreover, it is substantially simpler and so may serve as a good warm-up for the reader before tackling our main result.

2 Reprogramming Lemmas

In this section we collect some technical lemmas that we will need for the proof of [Theorem 1](#). We first discuss a particular extension of the “blinding lemma” of Alagic et al. [[1](#), Theorem 11], which formalizes [Lemma 1](#). We then state a generalization of the “reprogramming lemma” of Grilo et al. [[12](#)], which formalizes [Lemma 2](#). The complete proofs of these technical results are given in [Section 4](#).

We frequently consider adversaries with quantum access to some function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. This means the adversary is given access to a black-box gate implementing the $(n + m)$ -qubit unitary operator $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus f(x)\rangle$.

2.1 Arbitrary Reprogramming

Consider a reprogramming experiment that proceeds as follows. First, a distinguisher \mathcal{D} specifies an arbitrary function F along with a probabilistic algorithm \mathcal{B} which describes how to reprogram F . Specifically, the output of \mathcal{B} is a set of points B_1 at which F may be reprogrammed, along with the values the function should take at those potentially reprogrammed points. Then \mathcal{D} is given quantum access to either F or the reprogrammed version of F , and its goal is to determine which is the case. When \mathcal{D} is done making its oracle queries, it is also given the randomness that was used to run \mathcal{B} . Intuitively, the only way \mathcal{D} can tell if its oracle has been reprogrammed is by querying with significant amplitude on some point in B_1 . We bound \mathcal{D} 's advantage in terms of the probability that any particular value lies in the set B_1 defined by \mathcal{B} 's output.

By suitably modifying the proof of Alagic et al. [[1](#), Theorem 11], one can show that the distinguishing probability of \mathcal{D} in the game described above is at most $2q \cdot \sqrt{\epsilon}$, where q is an upper bound on the number of oracle queries and ϵ is an upper bound on the probability that any given input x is reprogrammed (i.e., that $x \in B_1$). However, that result is only proved for distinguishers with a fixed upper bound on the number of queries they make. To obtain a tighter bound for our application, we need a version of the result for distinguishers that may *adaptively* choose how many queries they make based on outcomes of intermediate measurements. We recover the aforementioned bound if we let q denote the number of queries made by \mathcal{D} *in expectation*.

For a function $F : \{0, 1\}^m \rightarrow \{0, 1\}^n$ and a set $B \subset \{0, 1\}^m \times \{0, 1\}^n$ such that each $x \in \{0, 1\}^m$ is the first element of at most one tuple in B , define

$$F^{(B)}(x) := \begin{cases} y & \text{if } (x, y) \in B \\ F(x) & \text{otherwise.} \end{cases}$$

We prove the following in [Section 4.1](#):

Lemma 3 (Formal version of [Lemma 1](#)). *Let \mathcal{D} be a distinguisher in the following game:*

Phase 1: \mathcal{D} outputs descriptions of a function $F_0 = F : \{0, 1\}^m \rightarrow \{0, 1\}^n$ and a randomized algorithm \mathcal{B} whose output is a set $B \subset \{0, 1\}^m \times \{0, 1\}^n$ where each $x \in \{0, 1\}^m$ is the first element of at most one tuple in B . Let $B_1 = \{x \mid \exists y : (x, y) \in B\}$ and $\epsilon = \max_{x \in \{0, 1\}^m} \{\Pr_{B \leftarrow \mathcal{B}}[x \in B_1]\}$.

Phase 2: \mathcal{B} is run to obtain B . Let $F_1 = F^{(B)}$. A uniform bit b is chosen, and \mathcal{D} is given quantum access to F_b .

Phase 3: \mathcal{D} loses access to F_b , and receives the randomness r used to invoke \mathcal{B} in phase 2. Then \mathcal{D} outputs a guess b' .

For any \mathcal{D} making q queries in expectation when its oracle is F_0 , it holds that

$$|\Pr[\mathcal{D} \text{ outputs } 1 \mid b = 1] - \Pr[\mathcal{D} \text{ outputs } 1 \mid b = 0]| \leq 2q \cdot \sqrt{\epsilon}.$$

2.2 Resampling

Here, we consider the following experiment: first, a distinguisher \mathcal{D} is given quantum access to an oracle for a random function F ; then, in the second stage, F may be “reprogrammed” so its value on a single, uniform point s is changed to an independent, uniform value. Because the distribution of $F(s)$ is the same both before and after any reprogramming, we refer to this as “resampling.” The goal for \mathcal{D} is to determine whether or not its oracle was resampled. Intuitively, the only way \mathcal{D} can tell if this is the case—even if it is given s and unbounded access to the oracle in the second stage—is if \mathcal{D} happened to put a large amplitude on s in some query to the oracle in the first stage. The lemmas we state here formalize this intuition.

We begin by establishing notation and recalling a result of Grilo et al. [12]. Given a function $F : \{0, 1\}^m \rightarrow \{0, 1\}^n$ and $s \in \{0, 1\}^m$, $y \in \{0, 1\}^n$, define the “reprogrammed” function $F_{s \rightarrow y} : \{0, 1\}^m \rightarrow \{0, 1\}^n$ as

$$F_{s \rightarrow y}(w) = \begin{cases} y & \text{if } w = s \\ F(w) & \text{otherwise.} \end{cases}$$

The following is a special case of [12, Prop. 1]:

Lemma 4 (Resampling for random functions). *Let \mathcal{D} be a distinguisher in the following game:*

Phase 1: *A uniform $F : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is chosen, and \mathcal{D} is given quantum access to $F_0 = F$.*

Phase 2: *Uniform $s \in \{0, 1\}^m$, $y \in \{0, 1\}^n$ are chosen, and we let $F_1 = F_{s \rightarrow y}$. A uniform bit b is chosen, and \mathcal{D} is given s and quantum access to F_b . Then \mathcal{D} outputs a guess b' .*

For any \mathcal{D} making at most q queries to F_0 in phase 1, it holds that

$$|\Pr[\mathcal{D} \text{ outputs } 1 \mid b = 1] - \Pr[\mathcal{D} \text{ outputs } 1 \mid b = 0]| \leq 1.5\sqrt{q/2^m}.$$

We extend the above to the case of two-way accessible, random *permutations*. Now, a random permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is chosen in the first phase; in the second phase, P may be reprogrammed by swapping the outputs corresponding to two uniform inputs. For $a, b \in \{0, 1\}^n$, let $\text{swap}_{a,b} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be the permutation that maps $a \mapsto b$ and $b \mapsto a$ but is otherwise the identity. We prove the following in [Section 4.2](#):

Lemma 5 (Formal version of Lemma 2). *Let \mathcal{D} be a distinguisher in the following game:*

Phase 1: *A uniform permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is chosen, and \mathcal{D} is given quantum access to $P_0 = P$ and $P_0^{-1} = P^{-1}$.*

Phase 2: *Uniform $s_0, s_1 \in \{0, 1\}^n$ are chosen, and we let $P_1 = P \circ \text{swap}_{s_0, s_1}$. Uniform $b \in \{0, 1\}$ is chosen, and \mathcal{D} is given s_0, s_1 , and quantum access to P_b, P_b^{-1} . Then \mathcal{D} outputs a guess b' .*

For any \mathcal{D} making at most q queries (combined) to P_0, P_0^{-1} in the first phase, $|\Pr[\mathcal{D} \text{ outputs } 1 \mid b = 1] - \Pr[\mathcal{D} \text{ outputs } 1 \mid b = 0]| \leq 4\sqrt{q/2^n}$.

3 Post-Quantum Security of Even-Mansour

We now establish the post-quantum security of the Even-Mansour cipher based on the lemmas from the previous section. Recall that the Even-Mansour cipher is defined as $E_k(x) := P(x \oplus k_1) \oplus k_2$, where $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a public random permutation and $k = (k_1, k_2) \in \{0, 1\}^{2n}$ is a key. Our proof assumes only that the marginal distributions of k_1 and k_2 are each uniform. This covers the original Even-Mansour cipher [11] where k is uniform over $\{0, 1\}^{2n}$ as well as the one-key variant [9] where k_1 is uniform and then k_2 is set equal to k_1 .

For E_k to be efficiently invertible, the permutation P must itself support efficient inversion; that is, the oracle for P must be accessible in both the forward and inverse directions. We thus consider adversaries \mathcal{A} who can access both the cipher E_k and the permutation P in both the forward and inverse directions. The

goal of \mathcal{A} is to distinguish this world from the ideal world in which it interacts with independent random permutations R, P . In this section, it will be implicit in our notation that all oracles are two-way accessible.

In the following, we let \mathcal{P}_n be the set of all permutations of $\{0, 1\}^n$. We write $E_k[P]$ to denote the Even-Mansour cipher using permutation P and key k ; we do this both to emphasize the dependence on P , and to enable references to Even-Mansour with a permutation other than P . Our main result is as follows:

Theorem 3 (Theorem 1, restated). *Let D be a distribution over $k = (k_1, k_2)$ such that the marginal distributions of k_1 and k_2 are each uniform, and let \mathcal{A} be an adversary making q_E classical queries to its first oracle and q_P quantum queries to its second oracle. Then*

$$\left| \Pr_{\substack{k \leftarrow D \\ P \leftarrow \mathcal{P}_n}} [\mathcal{A}^{E_k[P], P}(1^n) = 1] - \Pr_{R, P \leftarrow \mathcal{P}_n} [\mathcal{A}^{R, P}(1^n) = 1] \right| \leq 10 \cdot 2^{-n/2} (q_E \sqrt{q_P} + q_P \sqrt{q_E}).$$

Proof. Without loss of generality, we assume \mathcal{A} never makes a redundant classical query; that is, once it learns an input/output pair (x, y) by making a query to its classical oracle, it never again submits the query x (respectively, y) to the forward (respectively, inverse) direction of that oracle.

We divide an execution of \mathcal{A} into $q_E + 1$ stages $0, \dots, q_E$, where the j th stage corresponds to the time between the j th and $(j + 1)$ st classical queries of \mathcal{A} . In particular, the 0th stage corresponds to the period of time before \mathcal{A} makes its first classical query, and the q_E th stage corresponds to the period of time after \mathcal{A} makes its last classical query. We allow \mathcal{A} to adaptively distribute its q_P quantum queries between these stages arbitrarily. We let $q_{P,j}$ denote the expected number of queries \mathcal{A} makes in the j th stage in the ideal world $\mathcal{A}^{R,P}$; note that $\sum_{j=0}^{q_E} q_{P,j} = q_P$.

Recall that $\text{swap}_{a,b}$ swaps a and b . Given a permutation P , an ordered list of pairs $T = ((x_1, y_1), \dots, (x_t, y_t))$, and a key $k = (k_1, k_2)$, define

$$P_{T,k} = \text{swap}_{P(x_1 \oplus k_1), y_1 \oplus k_2} \circ \dots \circ \text{swap}_{P(x_t \oplus k_1), y_t \oplus k_2} \circ P. \quad (1)$$

(If T is empty, then $P_{T,k} = P$.) Intuitively, assuming the $\{x_i\}$ are distinct and the $\{y_i\}$ are distinct, $P_{T,k}$ is a “small” modification of P for which $E_k[P_{T,k}](x_i) = y_i$ for all i . (Note, however, that this may fail to hold if there is an “internal collision,” i.e., $P(x_i \oplus k_1) = y_j \oplus k_2$ for some $i \neq j$. But such collisions occur with low probability over choice of k_1, k_2 .)

We now define a sequence of experiments \mathbf{H}_j , for $j = 0, \dots, q_E$.

Experiment \mathbf{H}_j . Sample $R, P \leftarrow \mathcal{P}_n$ and $k \leftarrow D$. Then:

1. Run \mathcal{A} , answering its classical queries using R and its quantum queries using P , stopping immediately *before* its $(j + 1)$ st classical query. Let $T_j = ((x_1, y_1), \dots, (x_j, y_j))$ be the ordered list of all input/output pairs that \mathcal{A} received from its classical oracle.
2. For the remainder of the execution of \mathcal{A} , answer its classical queries using $E_k[P]$ and its quantum queries using $P_{T_j,k}$.

We can compactly represent \mathbf{H}_j as the experiment in which \mathcal{A} 's queries are answered using the oracle sequence

$$\underbrace{P, R, P, \dots, R, P}_{j \text{ classical queries}}, \underbrace{E_k[P], P_{T_j,k}, \dots, E_k[P], P_{T_j,k}}_{q_E - j \text{ classical queries}}.$$

Each appearance of R or $E_k[P]$ indicates a single classical query. Each appearance of P or $P_{T_j,k}$ indicates a stage during which \mathcal{A} makes multiple (quantum) queries to that oracle but no queries to its classical oracle. Observe that \mathbf{H}_0 corresponds to the execution of \mathcal{A} in the real world, i.e., $\mathcal{A}^{E_k[P], P}$, and that \mathbf{H}_{q_E} is the execution of \mathcal{A} in the ideal world, i.e., $\mathcal{A}^{R,P}$.

For $j = 0, \dots, q_E - 1$, we introduce additional experiments \mathbf{H}'_j :

Experiment \mathbf{H}'_j . Sample $R, P \leftarrow \mathcal{P}_n$ and $k \leftarrow D$. Then:

1. Run \mathcal{A} , answering its classical queries using R and its quantum queries using P , stopping immediately *after* its $(j+1)$ st classical query. Let $T_{j+1} = ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$ be the ordered list of all input/output pairs that \mathcal{A} learned from its classical oracle.
2. For the remainder of the execution of \mathcal{A} , answer its classical queries using $E_k[P]$ and its quantum queries using $P_{T_{j+1},k}$.

Thus, \mathbf{H}'_j corresponds to running \mathcal{A} using the oracle sequence

$$\underbrace{P, R, P, \dots, R, P, R, P_{T_{j+1},k}}_{j \text{ classical queries}}, \underbrace{E_k[P], P_{T_{j+1},k}, \dots, E_k[P], P_{T_{j+1},k}}_{q_E - j - 1 \text{ classical queries}}.$$

In Lemmas 6 and 7, we establish bounds on the distinguishability of \mathbf{H}'_j and \mathbf{H}_{j+1} , as well as \mathbf{H}_j and \mathbf{H}'_j . For $0 \leq j < q_E$ these give:

$$\begin{aligned} |\Pr[\mathcal{A}(\mathbf{H}'_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{j+1}) = 1]| &\leq 2 \cdot q_{P,j+1} \cdot \sqrt{\frac{2 \cdot (j+1)}{2^n}}. \\ |\Pr[\mathcal{A}(\mathbf{H}_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}'_j) = 1]| &\leq 8 \cdot \sqrt{\frac{q_P}{2^n}} + 2q_E \cdot 2^{-n} \end{aligned}$$

Using the above, we have

$$\begin{aligned} &|\Pr[\mathcal{A}(\mathbf{H}_0) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{q_E}) = 1]| \\ &\leq \sum_{j=0}^{q_E-1} \left(8 \cdot \sqrt{\frac{q_P}{2^n}} + 2q_E \cdot 2^{-n} + 2 \cdot q_{P,j+1} \sqrt{\frac{2 \cdot (j+1)}{2^n}} \right) \\ &\leq 2q_E^2 \cdot 2^{-n} + \sum_{j=0}^{q_E-1} \left(8 \cdot \sqrt{\frac{q_P}{2^n}} + 2 \cdot q_{P,j+1} \sqrt{\frac{2q_E}{2^n}} \right) \\ &\leq 2q_E^2 \cdot 2^{-n} + 2^{-n/2} \cdot (8q_E \sqrt{q_P} + 2 \cdot q_P \sqrt{2q_E}). \end{aligned}$$

We now simplify the bound further. If $q_P = 0$, then E_k and R are perfectly indistinguishable and the theorem holds; thus, we may assume $q_P \geq 1$. We can also assume $q_E < 2^{n/2}$ since otherwise the bound is larger than 1. Under these assumptions, we have $q_E^2 \cdot 2^{-n} \leq q_E \cdot 2^{-n/2} \leq q_E \sqrt{q_P} \cdot 2^{-n/2}$ and so

$$\begin{aligned} &2q_E^2 \cdot 2^{-n} + 2^{-n/2} (8q_E \sqrt{q_P} + 2q_P \sqrt{2q_E}) \\ &\leq 2 \cdot q_E \sqrt{q_P} \cdot 2^{-n/2} + 2^{-n/2} (8q_E \sqrt{q_P} + 2q_P \sqrt{2q_E}) \\ &\leq 10 \cdot 2^{-n/2} (q_E \sqrt{q_P} + q_P \sqrt{q_E}), \end{aligned}$$

as claimed. □

To complete the proof of [Theorem 3](#), we now establish the two lemmas showing that \mathbf{H}'_j is close to \mathbf{H}_{j+1} and \mathbf{H}_j is close to \mathbf{H}'_j for $0 \leq j < q_E$.

Lemma 6. For $j = 0, \dots, q_E - 1$,

$$\Pr[\mathcal{A}(\mathbf{H}'_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{j+1}) = 1] \leq 2 \cdot q_{P,j+1} \sqrt{2 \cdot (j+1)/2^n},$$

where $q_{P,j+1}$ is the expected number of queries \mathcal{A} makes to P in the $(j+1)$ st stage in the ideal world (i.e., in \mathbf{H}_{q_E}).

Proof. Recall we can write the oracle sequences defined by \mathbf{H}'_j and \mathbf{H}_{j+1} as

$$\begin{aligned} \mathbf{H}'_j &: P, R, P, \dots, R, P, R, P_{T_{j+1},k}, E_k[P], P_{T_{j+1},k}, \dots, E_k[P], P_{T_{j+1},k} \\ \mathbf{H}_{j+1} &: \underbrace{P, R, P, \dots, R, P, R, P}_{j \text{ classical queries}}, \underbrace{E_k[P], P_{T_{j+1},k}, \dots, E_k[P], P_{T_{j+1},k}}_{q_E - j - 1 \text{ classical queries}}. \end{aligned}$$

Let \mathcal{A} be a distinguisher between \mathbf{H}'_j and \mathbf{H}_{j+1} . We construct from \mathcal{A} a distinguisher \mathcal{D} for the blinding experiment from [Lemma 3](#):

Phase 1: \mathcal{D} samples $P, R \leftarrow \mathcal{P}_n$. It then runs \mathcal{A} , answering its quantum queries using P and its classical queries using R , until after it responds to \mathcal{A} 's $(j+1)$ st classical query. Let $T_{j+1} = ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$ be the list of input/output pairs \mathcal{A} received from its classical oracle thus far. \mathcal{D} defines $F(t, x) := P^t(x)$ for $t \in \{1, -1\}$. It also defines the following randomized algorithm \mathcal{B} : sample $k \leftarrow D$ and then compute the set B of input/output pairs to be reprogrammed so that $F^{(B)}(t, x) = P_{T_{j+1},k}^t(x)$ for all t, x .

Phase 2: \mathcal{B} is run to generate B , and \mathcal{D} is given quantum access to an oracle F_b . \mathcal{D} resumes running \mathcal{A} , answering its quantum queries using $P^t = F_b(t, \cdot)$. Phase 2 ends when \mathcal{A} makes its next (i.e., $(j+2)$ nd) classical query.

Phase 3: \mathcal{D} is given the randomness used by \mathcal{B} to generate k . It resumes running \mathcal{A} , answering its classical queries using $E_k[P]$ and its quantum queries using $P_{T_{j+1},k}$. Finally, it outputs whatever \mathcal{A} outputs.

Observe that \mathcal{D} is a valid distinguisher for the reprogramming experiment of [Lemma 3](#). It is immediate that if $b = 0$ (i.e., \mathcal{D} 's oracle in phase 2 is $F_0 = F$), then \mathcal{A} 's output is identically distributed to its output in \mathbf{H}_{j+1} , whereas if $b = 1$ (i.e., \mathcal{D} 's oracle in phase 2 is $F_1 = F^{(B)}$), then \mathcal{A} 's output is identically distributed to its output in \mathbf{H}'_j . It follows that $|\Pr[\mathcal{A}(\mathbf{H}'_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{j+1}) = 1]|$ is equal to the distinguishing advantage of \mathcal{D} in the reprogramming experiment. To bound this quantity using [Lemma 3](#), we bound the reprogramming probability ϵ and the expected number of queries made by \mathcal{D} in phase 2 (when $F = F_0$.)

The reprogramming probability ϵ can be bounded using the definition of $P_{T_{j+1},k}$ and the fact that $F^{(B)}(t, x) = P_{T_{j+1},k}^t$. Fixing P and T_{j+1} , the probability that any given (t, x) is reprogrammed is at most the probability (over k) that it is in the set

$$\{(1, x_i \oplus k_1), (1, P^{-1}(y_i \oplus k_2)), (-1, P(x_i \oplus k_1)), (-1, y_i \oplus k_2)\}_{i=1}^{j+1}.$$

Taking a union bound and applying the fact that the marginal distributions of k_1 and k_2 are each uniform, we get $\epsilon \leq 2(j+1)/2^n$.

The expected number of queries made by \mathcal{D} in Phase 2 when $F = F_0$ is equal to the expected number of queries made by \mathcal{A} in its $(j+1)$ st stage in \mathbf{H}_{j+1} . Since \mathbf{H}_{j+1} and \mathbf{H}_{q_E} are identical until after the $(j+1)$ st stage is complete, this is precisely $q_{P,j+1}$. \square

Lemma 7. For $j = 0, \dots, q_E$,

$$|\Pr[\mathcal{A}(\mathbf{H}_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}'_j) = 1]| \leq 8 \cdot \sqrt{\frac{q_P}{2^n}} + 2q_E \cdot 2^{-n}.$$

Proof. Recall that we can write the oracle sequences defined by \mathbf{H}_j and \mathbf{H}'_j as

$$\begin{aligned} \mathbf{H}_j &: P, R, P, \dots, R, P, E_k[P], P_{T_j,k}, E_k[P], P_{T_j,k}, \dots, E_k[P], P_{T_j,k} \\ \mathbf{H}'_j &: \underbrace{P, R, P, \dots, R, P}_{j \text{ classical queries}}, R, P_{T_{j+1},k}, \underbrace{E_k[P], P_{T_{j+1},k}, \dots, E_k[P], P_{T_{j+1},k}}_{q_E - j - 1 \text{ classical queries}}. \end{aligned}$$

Let \mathcal{A} be a distinguisher between \mathbf{H}_j and \mathbf{H}'_j . We construct from \mathcal{A} a distinguisher \mathcal{D} for the reprogramming experiment of [Lemma 5](#):

Phase 1: \mathcal{D} is given quantum access to a permutation P . It samples $R \leftarrow \mathcal{P}_n$ and then runs \mathcal{A} , answering its quantum queries with P and its classical queries with R (in the appropriate directions), until² \mathcal{A} submits its $(j+1)$ st classical query x_{j+1} . At that point, \mathcal{D} has a list $T_j = ((x_1, y_1), \dots, (x_j, y_j))$ of the input/output pairs \mathcal{A} has received from its classical oracle thus far.

Phase 2: Now \mathcal{D} receives $s_0, s_1 \in \{0, 1\}^n$ and quantum oracle access to a permutation P_b . Then \mathcal{D} sets $k_1 := s_0 \oplus x_{j+1}$, chooses $k_2 \leftarrow D_{|k_1}$ (where this represents the conditional distribution on k_2 given k_1), and sets $k := (k_1, k_2)$. \mathcal{D} continues running \mathcal{A} , answering its remaining classical queries (including the $(j+1)$ st one) using $E_k[P_b]$, and its remaining quantum queries using

$$(P_b)_{T_j, k} = \text{swap}_{P_b(x_1 \oplus k_1), y_1 \oplus k_2} \circ \dots \circ \text{swap}_{P_b(x_j \oplus k_1), y_j \oplus k_2} \circ P_b.$$

Finally, \mathcal{D} outputs whatever \mathcal{A} outputs.

Note that although \mathcal{D} makes additional queries to P_b at the start of phase 2 (to determine $P_b(x_1 \oplus k_1), \dots, P_b(x_j \oplus k_1)$), the bound of Lemma 5 only depends on the number of quantum queries \mathcal{D} makes in phase 1, which is at most q_P .

We now analyze the execution of \mathcal{D} in the two cases of the game of Lemma 5: $b = 0$ (no reprogramming) and $b = 1$ (reprogramming). In both cases, P and R are independent, uniform permutations, and \mathcal{A} is run with quantum oracle P and classical oracle R until it makes its $(j+1)$ st classical query; thus, through the end of phase 1, the above execution of \mathcal{A} is consistent with both \mathbf{H}_j and \mathbf{H}'_j .

At the start of phase 2, uniform $s_0, s_1 \in \{0, 1\}^n$ are chosen. Since \mathcal{D} sets $k_1 := s_0 \oplus x_{j+1}$, the distribution of k_1 is uniform and hence k is distributed according to D . The two cases ($b = 0$ and $b = 1$) now begin to diverge.

Case $b = 0$ (no reprogramming). In this case, \mathcal{A} 's remaining classical queries (including its $(j+1)$ st classical query) are answered using $E_k[P_0] = E_k[P]$, and its remaining quantum queries are answered using $(P_0)_{T_j, k} = P_{T_j, k}$. The output of \mathcal{A} is thus distributed identically to its output in \mathbf{H}_j in this case.

Case $b = 1$ (reprogramming). In this case, we have

$$P_b = P_1 = P \circ \text{swap}_{s_0, s_1} = \text{swap}_{P(s_0), P(s_1)} \circ P = \text{swap}_{P(x_{j+1} \oplus k_1), P(s_1)} \circ P. \quad (2)$$

The response to \mathcal{A} 's $(j+1)$ st classical query is thus

$$y_{j+1} \stackrel{\text{def}}{=} E_k[P_1](x_{j+1}) = P_1(x_{j+1} \oplus k_1) \oplus k_2 = P_1(s_0) \oplus k_2 = P(s_1) \oplus k_2. \quad (3)$$

The remaining classical queries of \mathcal{A} are then answered using $E_k[P_1]$, while its remaining quantum queries are answered using $(P_1)_{T_j, k}$. If we let Expt_j refer to the experiment in which \mathcal{D} executes \mathcal{A} as a subroutine when $b = 1$, it follows from Lemma 5 that

$$|\Pr[\mathcal{A}(\mathbf{H}_j) = 1] - \Pr[\mathcal{A}(\text{Expt}_j) = 1]| \leq 4\sqrt{q_P/2^n}. \quad (4)$$

We now define three events:

1. bad_1 is the event that $y_{j+1} \in \{y_1, \dots, y_j\}$.
2. bad_2 is the event that $s_1 \oplus k_1 \in \{x_1, \dots, x_j\}$.
3. bad_3 is the event that, in phase 2, \mathcal{A} queries its classical oracle in the forward direction on $s_1 \oplus k_1$, or the inverse direction on $P(s_0) \oplus k_2$ (with result $s_1 \oplus k_1$).

Since $y_{j+1} = P(s_1) \oplus k_2$ is uniform (because k_2 is uniform and independent of P and s_1), it is immediate that $\Pr[\text{bad}_1] \leq j/2^n$. Similarly, $s_1 \oplus k_1 = s_1 \oplus s_0 \oplus x_{j+1}$ is uniform, and so $\Pr[\text{bad}_2] \leq j/2^n$. As for the last event, we have:

² We assume for simplicity that this query is in the forward direction, but the case where it is in the inverse direction can be handled entirely symmetrically (using the fact that the marginal distribution of k_2 is uniform). The strings s_0 and s_1 are in that case replaced by $P_b(s_0)$ and $P_b(s_1)$. See Appendix B.2 for details.

Claim. $\Pr[\text{bad}_3] \leq (q_E - j)/2^n + 4\sqrt{q_P/2^n}$.

Proof. Consider the algorithm \mathcal{D}' that behaves identically to \mathcal{D} in phases 1 and 2, but then when \mathcal{A} terminates outputs 1 iff event bad_3 occurred. When $b = 0$ (no reprogramming), the execution of \mathcal{A} is independent of s_1 , and so the probability that bad_3 occurs is at most $(q_E - j)/2^n$. Now observe that \mathcal{D}' is a distinguisher for the reprogramming game of Lemma 5, with advantage $|\Pr[\text{bad}_3|b = 1] - (q_E - j)/2^n|$. The claim then follows from Lemma 5. \square

<pre> 1 $P, R \leftarrow \mathcal{P}_n$ 2 Run \mathcal{A} with quantum access to P and classical access to R, until \mathcal{A} makes its $(j + 1)$st classical query x_{j+1}; let T_j be as in the text 3 $s_0, s_1 \leftarrow \{0, 1\}^n$, $P_1 := P \circ \text{swap}_{s_0, s_1}$ 4 $k_1 := s_0 \oplus x_{j+1}$, $k_2 \leftarrow D_{ k_1}$, $k := (k_1, k_2)$ 5 $y_{j+1} := E_k[P_1](x_{j+1})$ 6 $Q := (P_1)_{T_j, k}$ 7 if $y_{j+1} \in \{y_1, \dots, y_j\}$ then $\text{bad}_1 := \text{true}$, $y_{j+1} \leftarrow \{0, 1\}^n \setminus \{y_1, \dots, y_j\}$ 8 Give y_{j+1} to \mathcal{A} as the answer to its $(j + 1)$st classical query 9 $T_{j+1} := ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$ 10 if $s_1 \oplus k_1 \in \{x_1, \dots, x_j\}$ then $\text{bad}_2 := \text{true}$ 11 if $\text{bad}_1 = \text{true}$ <i>or</i> $\text{bad}_2 = \text{true}$ then $Q := P_{T_{j+1}, k}$ 12 Continue running \mathcal{A} with quantum access to Q and classical access to $\mathcal{O}/\mathcal{O}^{-1}$ </pre>	<pre> 13 $\mathcal{O}(x)$ 14 $y := E_k[P_1](x)$ 15 if $x = s_1 \oplus k_1$ then 16 $\text{bad}_3 := \text{true}$, $y := E_k[P](x)$ 17 return y </pre>	<pre> 18 $\mathcal{O}^{-1}(y)$ 19 $x := E_k^{-1}[P_1](y)$ 20 if $x = s_1 \oplus k_1$ then 21 $\text{bad}_3 := \text{true}$, $x := E_k^{-1}[P](y)$ 22 return x </pre>
---	--	---

Fig. 1. Expt'_j includes the boxed statements, whereas Expt_j does not.

In Figure 1, we show code for Expt_j and a related experiment Expt'_j . Note that Expt_j and Expt'_j are identical until either $\text{bad}_1, \text{bad}_2$, or bad_3 occur, and so by the fundamental lemma of game playing³ [3] we have

$$\begin{aligned} |\Pr[\mathcal{A}(\text{Expt}'_j) = 1] - \Pr[\mathcal{A}(\text{Expt}_j) = 1]| &\leq \Pr[\text{bad}_1 \vee \text{bad}_2 \vee \text{bad}_3] \\ &\leq 2q_E/2^n + 4\sqrt{q_P/2^n}. \end{aligned} \quad (5)$$

We complete the proof by arguing that Expt'_j is identical to \mathbf{H}'_j :

1. In Expt'_j , the oracle Q used in line 12 is always equal to $P_{T_{j+1}, k}$. When bad_1 or bad_2 occurs this is immediate (since then Q is set to $P_{T_{j+1}, k}$ in line 11). But if bad_1 does not occur then Equation (3) holds, and if bad_2 does not occur then for $i = 1, \dots, j$ we have $x_i \oplus k_1 \neq s_0$ and $x_i \oplus k_1 \neq s_1$ (where the former is because $x_{j+1} \oplus k_1 = s_0$ but $x_i \neq x_{j+1}$ by assumption, and the latter is by definition of bad_2). Thus

³ This lemma is an information-theoretic result, and can be applied in our setting since everything we say in what follows holds even if \mathcal{A} is given the entire function table for its quantum oracle Q in line 12.

$P_1(x_i \oplus k_1) = P(x_i \oplus k_1)$ for $i = 1, \dots, j$, and so

$$\begin{aligned} Q &= (P_1)_{T_j, k} = \text{swap}_{P_1(x_1 \oplus k_1), y_1 \oplus k_2} \circ \dots \circ \text{swap}_{P_1(x_j \oplus k_1), y_j \oplus k_2} \circ P_1 \\ &= \text{swap}_{P(x_1 \oplus k_1), y_1 \oplus k_2} \circ \dots \circ \text{swap}_{P(x_{j+1} \oplus k_1), y_{j+1} \oplus k_2} \circ P \\ &= P_{T_{j+1}, k}, \end{aligned}$$

using Equations (2) and (3).

2. In Expt'_j , the value y_{j+1} is uniformly distributed in $\{0, 1\}^n \setminus \{y_1, \dots, y_j\}$. Indeed, we have already argued above that the value y_{j+1} computed in line 14 is uniform in $\{0, 1\}^n$. But if that value lies in $\{y_1, \dots, y_j\}$ (and so bad_1 occurs) then y_{j+1} is re-sampled uniformly from $\{0, 1\}^n \setminus \{y_1, \dots, y_j\}$ in line 7.
3. In Expt'_j , the response from oracle $\mathcal{O}(x)$ is always equal to $E_k[P](x)$. When bad_3 occurs this is immediate. But if bad_3 does not occur then $x \neq s_1 \oplus k_1$; we also know that $x \neq s_0 \oplus k_1 = x_{j+1}$ by assumption. But then $P_1(x \oplus k_1) = P(x \oplus k_1)$ and so $E_k[P_1](x) = E_k[P](x)$. A similar argument shows that the response from $\mathcal{O}^{-1}(y)$ is always $E_k^{-1}[P](y)$.

Syntactically rewriting Expt'_j using the above observations yields an experiment that is identical to \mathbf{H}'_j . (See Appendix B.1 for further details.) Lemma 7 thus follows from Equations (4) and (5). \square

4 Proofs of the Technical Lemmas

In this section, we give the proofs of our technical lemmas: the ‘‘arbitrary reprogramming lemma’’ (Lemma 3) and the ‘‘resampling lemma’’ (Lemma 5).

4.1 Proof of the Arbitrary Reprogramming Lemma

Lemma 3 allows for distinguishers that choose the number of queries they make adaptively, e.g., depending on the oracle provided and the outcomes of any measurements, and the bound is in terms of the number of queries \mathcal{D} makes *in expectation*. As discussed in Section 1.1, the ability to directly handle such adaptive distinguishers is necessary for our proof, and to our knowledge has not been addressed before. To formally reason about adaptive distinguishers, we model the intermediate operations of the distinguisher and the measurements it makes as *quantum channels*. With this as our goal, we first recall some necessary background and establish some notation.

Recall that a density matrix ρ is a positive semidefinite matrix with unit trace. A quantum channel—the most general transformation between density matrices allowed by quantum theory—is a completely positive, trace-preserving, linear map. The quantum channel corresponding to the unitary operation U is the map $\rho \mapsto U\rho U^\dagger$. Another type of quantum channel is a *pinching*, which corresponds to the operation of making a measurement. Specializing to the only kind of pinching needed in our proof, consider the measurement of a single-qubit register C given by the projectors $\{\Pi_0, \Pi_1\}$ with $\Pi_b = |b\rangle\langle b|_C$. This corresponds to the pinching \mathcal{M}_C where

$$\mathcal{M}_C(\rho) = \Pi_0\rho\Pi_0 + \Pi_1\rho\Pi_1.$$

Observe that a pinching only produces the post-measurement state, and does not separately give the outcome (i.e., the result 0 or 1).

Consider a quantum algorithm \mathcal{D} with access to an oracle \mathcal{O} operating on registers X, Y (so $\mathcal{O}|x\rangle|y\rangle = |x\rangle|y \oplus \mathcal{O}(x)\rangle$). We define the unitary $c\mathcal{O}$ for the *controlled* version of \mathcal{O} , operating on registers C, X , and Y (with C a single-qubit register), as

$$c\mathcal{O}|c\rangle|x\rangle|y\rangle = |c\rangle|x\rangle|y \oplus c \cdot \mathcal{O}(x)\rangle.$$

With this in place, we may now view an execution of $\mathcal{D}^{\mathcal{O}}$ as follows. The algorithm uses registers C, X, Y , and E . Let q_{\max} be an upper bound on the number of queries \mathcal{D} ever makes. Then \mathcal{D} applies the quantum channel

$$(\Phi \circ c\mathcal{O} \circ \mathcal{M}_C)^{q_{\max}} \tag{6}$$

to some initial state $\rho = \rho_0^{(0)}$. That is, for each of q_{\max} iterations, \mathcal{D} applies to its current state the pinching \mathcal{M}_C followed by the controlled oracle $c\mathcal{O}$ and then an arbitrary quantum channel Φ (that we take to be the same in all iterations without loss of generality⁴) operating on all its registers. Finally, \mathcal{D} applies a measurement to produce its final output. If we let $\rho_{i-1}^{(0)}$ denote the intermediate state immediately before the pinching is applied in the i th iteration, then $p_{i-1} = \text{Tr} \left[|1\rangle\langle 1|_C \rho_{i-1}^{(0)} \right]$ represents the probability that the oracle is applied (or, equivalently, that a query is made) in the i th iteration, and so $q = \sum_{i=1}^{q_{\max}} p_{i-1}$ is the expected number of queries made by \mathcal{D} when interacting with oracle \mathcal{O} .

Proof of Lemma 3. An execution of \mathcal{D} takes the form of Equation (6) up to a final measurement. For some fixed value of the randomness r used to run \mathcal{B} , set $\Upsilon_b = \Phi \circ c\mathcal{O}_{F_b} \circ \mathcal{M}_C$, and define

$$\rho_k \stackrel{\text{def}}{=} \left(\Upsilon_1^{q_{\max}-k} \circ \Upsilon_0^k \right) (\rho),$$

so that ρ_k is the final state if the first k queries are answered using a (controlled) F_0 oracle and then the remaining $q_{\max} - k$ queries are answered using a (controlled) F_1 oracle. Furthermore, we define $\rho_i^{(0)} = \Upsilon_i^i(\rho)$. Note also that $\rho_{q_{\max}}$ (resp., ρ_0) is the final state of the algorithm when the F_0 oracle (resp., F_1 oracle) is used the entire time. We bound $\mathbb{E}_r \left[\delta(|r\rangle\langle r| \otimes \rho_{q_{\max}}, |r\rangle\langle r| \otimes \rho_0) \right]$, where $\delta(\cdot, \cdot)$ denotes the trace distance.

Define $\tilde{F}^{(B)}(x) = F(x) \oplus F^{(B)}(x)$, and note that $\tilde{F}^{(B)}(x) = 0^n$ for $x \notin B_1$. Since trace distance is non-increasing under quantum channels, for any r we have

$$\begin{aligned} \delta(|r\rangle\langle r| \otimes \rho_k, |r\rangle\langle r| \otimes \rho_{k-1}) &\leq \delta \left(c\mathcal{O}_{F_0} \circ \mathcal{M}_C \left(\rho_{k-1}^{(0)} \right), c\mathcal{O}_{F_1} \circ \mathcal{M}_C \left(\rho_{k-1}^{(0)} \right) \right) \\ &= \delta \left(\mathcal{M}_C \left(\rho_{k-1}^{(0)} \right), c\mathcal{O}_{\tilde{F}^{(B)}} \circ \mathcal{M}_C \left(\rho_{k-1}^{(0)} \right) \right). \end{aligned}$$

By definition of a controlled oracle,

$$\begin{aligned} c\mathcal{O}_{\tilde{F}^{(B)}} \circ \mathcal{M}_C \left(\rho_{k-1}^{(0)} \right) &= c\mathcal{O}_{\tilde{F}^{(B)}} \left(|1\rangle\langle 1|_C \rho_{k-1}^{(0)} |1\rangle\langle 1|_C + |0\rangle\langle 0|_C \rho_{k-1}^{(0)} |0\rangle\langle 0|_C \right) \\ &= \mathcal{O}_{\tilde{F}^{(B)}} \left(|1\rangle\langle 1|_C \rho_{k-1}^{(0)} |1\rangle\langle 1|_C \right) + |0\rangle\langle 0|_C \rho_{k-1}^{(0)} |0\rangle\langle 0|_C, \end{aligned}$$

and thus

$$\begin{aligned} &\delta \left(\mathcal{M}_C \left(\rho_{k-1}^{(0)} \right), c\mathcal{O}_{\tilde{F}^{(B)}} \circ \mathcal{M}_C \left(\rho_{k-1}^{(0)} \right) \right) \\ &= \delta \left(|1\rangle\langle 1|_C \rho_{k-1}^{(0)} |1\rangle\langle 1|_C, \mathcal{O}_{\tilde{F}^{(B)}} \left(|1\rangle\langle 1|_C \rho_{k-1}^{(0)} |1\rangle\langle 1|_C \right) \right) \\ &= p_{k-1} \cdot \delta \left(\sigma_{k-1}, \mathcal{O}_{\tilde{F}^{(B)}} \left(\sigma_{k-1} \right) \right) \end{aligned}$$

where, recall, $p_{k-1} = \text{Tr} \left[|1\rangle\langle 1|_C \rho_{k-1}^{(0)} \right]$ is the probability that a query is made in the k th iteration, and we define the normalized state $\sigma_{k-1} \stackrel{\text{def}}{=} \frac{|1\rangle\langle 1|_C \rho_{k-1}^{(0)} |1\rangle\langle 1|_C}{p_{k-1}}$. Therefore,

$$\begin{aligned} &\mathbb{E}_r \left[\delta(|r\rangle\langle r| \otimes \rho_{q_{\max}}, |r\rangle\langle r| \otimes \rho_0) \right] \\ &\leq \sum_{k=1}^{q_{\max}} \mathbb{E}_B \left[\delta(|r\rangle\langle r| \otimes \rho_k, |r\rangle\langle r| \otimes \rho_{k-1}) \right] \\ &\leq \sum_{k=1}^{q_{\max}} p_{k-1} \cdot \mathbb{E}_B \left[\delta \left(\sigma_{k-1}, \mathcal{O}_{\tilde{F}^{(B)}} \left(\sigma_{k-1} \right) \right) \right] \\ &\leq q \cdot \max_{\sigma} \mathbb{E}_B \left[\delta \left(\sigma, \mathcal{O}_{\tilde{F}^{(B)}} \left(\sigma \right) \right) \right], \end{aligned} \tag{7}$$

⁴ This can be done by having a register serve as a counter that is incremented with each application of Φ .

where we write \mathbb{E}_B for the expectation over the set B output by \mathcal{B} in place of \mathbb{E}_r .

Since σ can be purified to some state $|\psi\rangle$, and $\delta(|\psi\rangle, |\psi'\rangle) \leq \|\psi\rangle - |\psi'\rangle\|_2$ for pure states $|\psi\rangle, |\psi'\rangle$, we have

$$\begin{aligned} \max_{\sigma} \mathbb{E}_B [\delta(\sigma, \mathcal{O}_{\tilde{F}(B)}(\sigma))] &\leq \max_{|\psi\rangle} \mathbb{E}_B [\delta(|\psi\rangle, \mathcal{O}_{\tilde{F}(B)}|\psi\rangle)] \\ &\leq \max_{|\psi\rangle} \mathbb{E}_B [\|\psi\rangle - \mathcal{O}_{\tilde{F}(B)}|\psi\rangle\|_2]. \end{aligned}$$

Because $\mathcal{O}_{\tilde{F}(B)}$ acts as the identity on $(\mathbb{I} - \Pi_{B_1})|\psi\rangle$ for any $|\psi\rangle$, we have

$$\begin{aligned} &\mathbb{E}_B [\|\psi\rangle - \mathcal{O}_{\tilde{F}(B)}|\psi\rangle\|_2] \\ &= \mathbb{E}_B [\|\Pi_{B_1}|\psi\rangle - \mathcal{O}_{\tilde{F}(B)}\Pi_{B_1}|\psi\rangle + (\mathbb{I} - \mathcal{O}_{\tilde{F}(B)})(\mathbb{I} - \Pi_{B_1})|\psi\rangle\|_2] \\ &\leq \mathbb{E}_B [\|\Pi_{B_1}|\psi\rangle\|_2] + \mathbb{E}_B [\|\mathcal{O}_{\tilde{F}(B)}\Pi_{B_1}|\psi\rangle\|_2] \\ &= 2 \cdot \mathbb{E}_B [\|\Pi_{B_1}|\psi\rangle\|_2] \\ &\leq 2\sqrt{\mathbb{E}_B [\|\Pi_{B_1}|\psi\rangle\|_2^2]}, \end{aligned} \tag{8}$$

using Jensen's inequality in the last step. Let $|\psi\rangle = \sum_{x \in \{0,1\}^m, y \in \{0,1\}^n} \alpha_{x,y} |x\rangle|y\rangle$ where $\|\psi\rangle\|_2^2 = \sum_{x,y} \alpha_{x,y}^2 = 1$. Then

$$\begin{aligned} \mathbb{E}_B [\|\Pi_{B_1}|\psi\rangle\|_2^2] &= \mathbb{E}_B \left[\sum_{x,y: x \in B_1} \alpha_{x,y}^2 \right] \\ &= \sum_{x,y} \alpha_{x,y}^2 \cdot \Pr[x \in B_1] \leq \epsilon. \end{aligned}$$

Together with Equations (7) and (8), this gives the desired result. \square

4.2 Proof of the Resampling Lemma

We begin by introducing a superposition-oracle technique based on the one by Zhandry [26], but different in that our oracle represents a two-way accessible, uniform permutation (rather than a uniform function). We also do not need to ‘‘compress’’ the oracle, as an inefficient representation suffices for our purposes.

For an arbitrary function $f : \{0,1\}^n \rightarrow \{0,1\}^n$, define the state

$$|f\rangle_F = \bigotimes_{x \in \{0,1\}^n} |f(x)\rangle_{F_x},$$

where F is the collection of registers $\{F_x\}_{x \in \{0,1\}^n}$. We represent an evaluation of f via an operator O whose action on the computational basis is given by

$$O_{XYF} |x\rangle_X |y\rangle_Y |f\rangle_F = \text{CNOT}_{F_x:Y}^{\otimes n} |x\rangle_X |y\rangle_Y |f\rangle_F = |x\rangle_X |y \oplus f(x)\rangle_Y |f\rangle_F,$$

where X, Y are n -qubit registers. Handling inverse queries to f is more difficult. We want to define an inverse operator O^{inv} such that, for any permutation π ,

$$O_{XYF}^{\text{inv}} |\pi\rangle_F = \left(\sum_{x,y \in \{0,1\}^n} |y\rangle\langle y|_Y \otimes X_X^x \otimes |y\rangle\langle y|_{F_x} \right) |\pi\rangle_F \tag{9}$$

(where X is the Pauli- X operator, and for $x \in \{0,1\}^n$ we let $X^x := X^{x_1} \otimes X^{x_2} \otimes \dots \otimes X^{x_n}$ so that $X^x |\hat{x}\rangle = |\hat{x} \oplus x\rangle$); then,

$$O_{XYF}^{\text{inv}} |x\rangle_X |y\rangle_Y |\pi\rangle_F = |x \oplus \pi^{-1}(y)\rangle_X |y\rangle_Y |\pi\rangle_F.$$

In order for O^{inv} to be a well-defined unitary operator, however, we must extend its definition to the entire space of functions. A convenient extension is given by the following action on arbitrary computational basis states:

$$O_{XYF}^{\text{inv}} = \prod_{x' \in \{0,1\}^n} \left(X_X^{x'} \otimes |y\rangle\langle y|_{F_{x'}} + (\mathbb{1} - |y\rangle\langle y|)_{F_{x'}} \right),$$

so that

$$O_{XYF}^{\text{inv}} |x\rangle_X |y\rangle_Y |f\rangle_F = |x \oplus (\oplus_{x': f(x')=y} x')\rangle_X |y\rangle_Y |f\rangle_F.$$

In other words, the inverse operator XORs all preimages (under f) of the value in register Y into the contents of register X .

We may view a uniform permutation as a uniform superposition over all permutations in \mathcal{P}_n ; i.e., we model a uniform permutation as the state

$$|\phi_0\rangle_F = (2^n!)^{-\frac{1}{2}} \sum_{\pi \in \mathcal{P}_n} |\pi\rangle_F.$$

The final state of any oracle algorithm \mathcal{D} is identically distributed whether we (1) sample uniform $\pi \in \mathcal{P}_n$ and then run \mathcal{D} with access to π and π^{-1} , or (2) run \mathcal{D} with access to O and O^{inv} after initializing the F -registers to $|\phi_0\rangle_F$ (and, if desired, at the end of its execution, measure the F -registers to obtain π and the residual state of \mathcal{D}).

Our proof relies on the following lemma, which is a special case of the conclusion of implication (\diamond') in [21]. (Here and in the following, we denote the complementary projector of a projector P by $\bar{P} \stackrel{\text{def}}{=} \mathbb{1} - P$.)

Lemma 8 (Gentle measurement lemma). *Let $|\psi\rangle$ be a quantum state and let $\{P_i\}_{i=1}^q$ be a collection of projectors with $\|\bar{P}_i|\psi\rangle\|_2^2 \leq \epsilon_i$ for all i . Then*

$$1 - |\langle \psi | (P_q \cdots P_1) | \psi \rangle|^2 \leq \sum_{i=1}^q \epsilon_i.$$

Proof of Lemma 5. We split the distinguisher \mathcal{D} into two stages $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ corresponding to the first and second phases of the experiment in Lemma 5. As discussed above, we run the experiment using the superposition oracle $|\phi_0\rangle_F$ and then measure the F -registers at the end. Informally, our goal is to show that on average over the choice of reprogrammed positions s_0, s_1 , the adversary-oracle state after \mathcal{D}_0 finishes is almost invariant under the reprogramming operation (i.e., the swap of registers F_{s_0} and F_{s_1}) unless \mathcal{D}_0 makes a large number of oracle queries. This will follow from Lemma 8 because, on average over the choice of s_0, s_1 , any particular query of \mathcal{D}_0 (whether using O or O^{inv}) only involves F_{s_0} or F_{s_1} with negligible amplitude.

We begin by defining the projectors

$$(P_{s_0 s_1})_X = \begin{cases} \mathbb{1} & s_0 = s_1 \\ \mathbb{1} - |s_0\rangle\langle s_0| - |s_1\rangle\langle s_1| & s_0 \neq s_1 \end{cases}$$

$$(P_{s_0 s_1}^{\text{inv}})_{FY} = \begin{cases} \mathbb{1} & s_0 = s_1 \\ \sum_{y \in \{0,1\}^n} |y\rangle\langle y|_Y \otimes (\mathbb{1} - |y\rangle\langle y|)_{F_{s_0} F_{s_1}}^{\otimes 2} & s_0 \neq s_1. \end{cases}$$

It is straightforward to verify that for any s_0, s_1 :

$$\left[\text{Swap}_{F_{s_0} F_{s_1}}, O_{XYF} (P_{s_0 s_1})_X \right] = 0 \quad (10)$$

$$\left[\text{Swap}_{F_{s_0} F_{s_1}}, O_{XYF}^{\text{inv}} (P_{s_0 s_1}^{\text{inv}})_{FY} \right] = 0, \quad (11)$$

where $[\cdot, \cdot]$ denotes the commutator operation, and Swap_{AB} is the swap operator (i.e., $\text{Swap}_{A,B}|x\rangle_A |x'\rangle_B = |x'\rangle_A |x\rangle_B$ if the target registers A, B are distinct, and the identity if A and B refer to the same register). In

words, this means that if we project a forward query to inputs other than s_0, s_1 , then swapping the outputs of a function at s_0 and s_1 before evaluating that function has no effect; the same holds if we project an inverse query (for some associated function f) to the set of output values that are not equal to $f(s_0)$ or $f(s_1)$.

Since $\bar{P}_{s_0 s_1} \stackrel{\text{def}}{=} \mathbf{1} - P_{s_0 s_1} \leq |s_0\rangle\langle s_0| + |s_1\rangle\langle s_1|$ it follows that for any normalized state $|\psi\rangle_{XE}$ (where E is an arbitrary other register),

$$\begin{aligned} \mathbb{E}_{s_0, s_1} \left[\left\| (\bar{P}_{s_0 s_1})_X |\psi\rangle_{XE} \right\|_2^2 \right] &\leq \mathbb{E}_{s_0, s_1} [\langle \psi | (|s_0\rangle\langle s_0| + |s_1\rangle\langle s_1|) | \psi \rangle] \\ &= 2 \cdot 2^{-n}. \end{aligned} \quad (12)$$

We show a similar statement about $P_{s_0 s_1}^{\text{inv}}$. We can express a valid adversary/oracle state $|\psi\rangle_{YXEF}$ (that is thus only supported on the span of \mathcal{P}_n) as

$$|\psi\rangle_{YXEF} = \sum_{x, y \in \{0, 1\}^n} c_{xy} |y\rangle_Y |y\rangle_{E_x} |\psi_{xy}\rangle_{XEF_{x^c}}, \quad (13)$$

for some normalized quantum states $\{|\psi_{xy}\rangle\}_{x, y \in \{0, 1\}^n}$, with $\sum_{x, y \in \{0, 1\}^n} |c_{xy}|^2 = 1$ and $\langle y |_{E_x} |\psi_{xy}\rangle_{XEF_{x^c}} = 0$ for all $x' \neq x$. If $s_0 = s_1$, then $\left\| (\bar{P}_{s_0 s_1}^{\text{inv}})_{YF} |\psi\rangle_{YXEF} \right\|_2^2 = 0 \leq 2 \cdot 2^{-n}$. It is thus immediate from eq. (13) that

$$\mathbb{E}_{s_0, s_1} \left[\left\| (\bar{P}_{s_0 s_1}^{\text{inv}})_{YF} |\psi\rangle_{YXEF} \right\|_2^2 \right] \leq 2 \cdot 2^{-n} \quad (14)$$

Without loss of generality, we assume \mathcal{D}_0 starts with initial state $|\psi_0\rangle = |\psi'_0\rangle|\phi_0\rangle$ (which we take to include the superposition oracle's initial state $|\phi_0\rangle$), computes the state

$$|\psi\rangle = U_{\mathcal{D}_0} |\psi_0\rangle = U_q O_q U_{q-1} O_{q-1} \cdots U_1 O_1 |\psi_0\rangle,$$

and outputs all its registers as a state register E . Here, each $O_i \in \{O, O^{\text{inv}}\}$ acts on registers XYF , and each U_j acts on registers XYE . To each choice of s_0, s_1 we assign a decomposition $|\psi\rangle = |\psi_{\text{good}}(s_0, s_1)\rangle + |\psi_{\text{bad}}(s_0, s_1)\rangle$ by defining

$$|\psi_{\text{good}}(s_0, s_1)\rangle = z \cdot U_q O_q P_{s_0 s_1}^q U_{q-1} O_{q-1} P_{s_0 s_1}^{q-1} \cdots U_1 O_1 P_{s_0 s_1}^1 |\psi_0\rangle,$$

where $P_{s_0 s_1}^i = P_{s_0 s_1}$ if $O_i = O$, $P_{s_0 s_1}^i = P_{s_0 s_1}^{\text{inv}}$ if $O_i = O^{\text{inv}}$, and $z \in \mathbb{C}$ is such that $|z| = 1$ and $\langle \psi | \psi_{\text{good}}(s_0, s_1) \rangle \in \mathbb{R}_{\geq 0}$.

$$|\psi_{\text{good}}(s_0, s_1)\rangle = z \cdot U_{\mathcal{D}_0} Q_{s_0 s_1}^q \cdots Q_{s_0 s_1}^1 |\psi_0\rangle,$$

with $Q_{s_0 s_1}^i = \tilde{U}_i^\dagger P_{s_0 s_1}^i \tilde{U}_i$ for $\tilde{U}_i = U_{i-1} O_{i-1} \cdots U_1 O_1$. Let

$$\epsilon_i(s_0, s_1) = \left\| \bar{Q}_{s_0 s_1}^i |\psi_0\rangle \right\|_2^2 = \left\| \bar{P}_{s_0 s_1}^i \tilde{U}_i |\psi_0\rangle \right\|_2^2.$$

Applying Lemma 8 yields

$$1 - |\langle \psi | \psi_{\text{good}}(s_0, s_1) \rangle|^2 \leq \sum_{i=1}^q \epsilon_i(s_0, s_1). \quad (15)$$

We will now analyze the impact of reprogramming the superposition oracle after \mathcal{D}_0 has finished. Recall that reprogramming swaps the values of the permutation at points s_0 and s_1 , which is implemented in the superposition-oracle framework by applying $\text{Swap}_{F_{s_0} F_{s_1}}$. Note that $\text{Swap}_{F_{s_0} F_{s_1}} |\phi_0\rangle = |\phi_0\rangle$. As the adversary's internal unitaries U_i do not act on F , Equations (10) and (11) then imply that

$$\text{Swap}_{F_{s_0} F_{s_1}} |\psi_{\text{good}}(s_0, s_1)\rangle = |\psi_{\text{good}}(s_0, s_1)\rangle.$$

The standard formula for the trace distance of pure states thus yields

$$\frac{1}{2} \left\| |\psi\rangle\langle\psi| - \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle\langle\psi| \right\|_1 = \sqrt{1 - \left| \langle\psi| \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle \right|^2}. \quad (16)$$

We further have

$$\begin{aligned} \left| \langle\psi| \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle \right| &= \left| \langle\psi| \psi\rangle + \langle\psi_{\text{bad}}(s_0, s_1)| (\text{Swap}_{F_{s_0} F_{s_1}} - \mathbf{1}) |\psi_{\text{bad}}(s_0, s_1)\rangle \right| \\ &\geq 1 - 2 \|\psi_{\text{bad}}(s_0, s_1)\|_2^2 \end{aligned} \quad (17)$$

using the triangle and Cauchy-Schwarz inequalities. Combining Equations (16) and (17) we obtain

$$\frac{1}{2} \left\| |\psi\rangle\langle\psi| - \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle\langle\psi| \right\|_1 \leq 2 \cdot \|\psi_{\text{bad}}(s_0, s_1)\|_2.$$

But as $|\psi_{\text{bad}}(s_0, s_1)\rangle = |\psi\rangle - |\psi_{\text{good}}(s_0, s_1)\rangle$, we have

$$\begin{aligned} \|\psi_{\text{bad}}(s_0, s_1)\|_2^2 &= 2 - 2 \cdot \text{Re} \langle\psi| \psi_{\text{good}}(s_0, s_1)\rangle \\ &= 2 - 2 \cdot |\langle\psi| \psi_{\text{good}}(s_0, s_1)\rangle| \\ &\leq 2 \sum_{i=1}^q \epsilon_i(s_0, s_1). \end{aligned} \quad (18)$$

Combining the last two equations we obtain

$$\frac{1}{2} \left\| |\psi\rangle\langle\psi| - \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle\langle\psi| \right\|_1 \leq 2\sqrt{2} \sqrt{\sum_{i=1}^q \epsilon_i(s_0, s_1)}. \quad (19)$$

The remainder of the proof is the same as the analogous part of the proof of [12, Theorem 6]. \mathcal{D}_1 's task boils down to distinguishing the states $|\psi\rangle$ and $\text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle$, for uniform s_0, s_1 that \mathcal{D}_1 receives as input, using the limited set of instructions allowed by the superposition oracle. We can therefore bound \mathcal{D} 's advantage by the maximum distinguishing advantage for these two states when using arbitrary quantum computation, averaged over the choice of s_0, s_1 . Using the standard formula for this maximum distinguishing advantage we obtain

$$\begin{aligned} \Pr[\mathcal{D} \text{ outputs } b] - \frac{1}{2} &\leq \frac{1}{4} \mathbb{E}_{s_0, s_1} \left[\left\| |\psi\rangle\langle\psi| - \text{Swap}_{F_{s_0} F_{s_1}} |\psi\rangle\langle\psi| \right\|_1 \right] \\ &\leq \sqrt{2} \mathbb{E}_{s_0, s_1} \left[\sqrt{\sum_{i=1}^q \epsilon_i(s_0, s_1)} \right] \\ &\leq \sqrt{2} \sqrt{\mathbb{E}_{s_0, s_1} \left[\sum_{i=1}^q \epsilon_i(s_0, s_1) \right]} \leq 2\sqrt{\frac{q}{2^n}}, \end{aligned}$$

where the second inequality is Equation (19), the third is Jensen's inequality, and the last is from Equations (12)–(15). This implies the lemma. \square

Acknowledgments

The authors thank Andrew Childs and Bibhusa Rawal for useful discussions. Work of Jonathan Katz was supported in part by financial assistance award 70NANB19H126 from the U.S. Department of Commerce, National Institute of Standards and Technology. Work of Christian Majenz was funded by a NWO VENI grant (Project No. VI.Veni.192.159). Gorjan Alagic acknowledges support from the U.S. Army Research Office under Grant Number W911NF-20-1-0015, the U.S. Department of Energy under Award Number DE-SC0020312, and the AFOSR under Award Number FA9550-20-1-0108.

References

1. Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song. Quantum-access-secure message authentication via blind-unforgeability. In *Advances in Cryptology—Eurocrypt 2020, Part III*, volume 12107 of *LNCS*, pages 788–817. Springer, 2020.
2. Gorjan Alagic and Alexander Russell. Quantum-secure symmetric-key cryptography based on hidden shifts. In *Advances in Cryptology—Eurocrypt 2017, Part III*, volume 10212 of *LNCS*, pages 65–93. Springer, 2017.
3. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Advances in Cryptology—Eurocrypt 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, 2006. Full version available at <https://eprint.iacr.org/2004/331>.
4. Nina Bindel, Mike Hamburg, Kathrin Hövelmanns, Andreas Hülsing, and Edoardo Persichetti. Tighter proofs of CCA security in the quantum random oracle model. In *17th Theory of Cryptography Conference—TCC 2019, Part II*, volume 11892 of *LNCS*, pages 61–90. Springer, 2019.
5. Xavier Bonnetain, Akinori Hosoyamada, María Naya-Plasencia, Yu Sasaki, and André Schrottenloher. Quantum attacks without superposition queries: The offline Simon’s algorithm. In *Advances in Cryptology—Asiacrypt 2019, Part I*, volume 11921 of *LNCS*, pages 552–583. Springer, 2019.
6. Xavier Bonnetain and María Naya-Plasencia. Hidden shift quantum cryptanalysis and implications. In *Advances in Cryptology—Asiacrypt 2018, Part I*, volume 11272 of *LNCS*, pages 560–592. Springer, 2018.
7. Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum algorithm for the collision problem, 1997. Available at <https://arxiv.org/abs/quant-ph/9705002>.
8. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. Cryptology ePrint Archive, Report 2021/280, 2021. <https://eprint.iacr.org/2021/280>.
9. Orr Dunkelman, Nathan Keller, and Adi Shamir. Minimalism in cryptography: The Even-Mansour scheme revisited. In *Advances in Cryptology—Eurocrypt 2012*, volume 7237 of *LNCS*, pages 336–354. Springer, 2012.
10. Mark Ettinger, Peter Høyer, and Emanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, 91(1):43–48, 2004.
11. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–161, 1997.
12. Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the QROM. In *Advances in Cryptology—Asiacrypt 2021, Part I*, volume 13090 of *LNCS*, pages 637–667. Springer, 2021. Available at <https://eprint.iacr.org/2020/1361>.
13. Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In *15th Theory of Cryptography Conference—TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, 2017.
14. Akinori Hosoyamada and Yu Sasaki. Cryptanalysis against symmetric-key schemes with online classical queries and offline quantum computations. In *Topics in Cryptology—Cryptographers’ Track at the RSA Conference (CT-RSA) 2018*, volume 10808 of *LNCS*, pages 198–218. Springer, 2018.
15. Hector Bjølshavn Hougaard. How to generate pseudorandom permutations over other groups: Even-Mansour and Feistel revisited, 2017. Available at <https://arxiv.org/abs/1707.01699>.
16. Joseph Jaeger, Fang Song, and Stefano Tessaro. Quantum key-length extension. In *19th Theory of Cryptography Conference—TCC 2021, Part I*, volume 13042 of *LNCS*, pages 209–239. Springer, 2021.
17. Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In *Advances in Cryptology—Crypto 2016, Part II*, volume 9815 of *LNCS*, pages 207–237. Springer, 2016.
18. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In *Advances in Cryptology—Eurocrypt 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, 2018.
19. Veronika Kuchta, Amin Sakzad, Damien Stehlé, Ron Steinfeld, and Shifeng Sun. Measure-rewind-measure: Tighter quantum random oracle model proofs for one-way to hiding and CCA security. In *Advances in Cryptology—Eurocrypt 2020, Part III*, volume 12107 of *LNCS*, pages 703–728. Springer, 2020.
20. Hidenori Kuwakado and Masakatu Morii. Security on the quantum-type Even-Mansour cipher. In *Proc. International Symposium on Information Theory and its Applications*, pages 312–316. IEEE Computer Society, 2012.
21. Ryan O’Donnell and Ramgopal Venkateswaran. The quantum union bound made easy, 2021. Available at <https://arxiv.org/abs/2103.07827>.
22. Daniel R. Simon. On the power of quantum computation. *SIAM J. Computing*, 26(5):1474–1483, 1997.

23. Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In *14th Theory of Cryptography Conference—TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 192–216. Springer, 2016.
24. Dominique Unruh. Post-quantum security of Fiat-Shamir. In *Advances in Cryptology—Asiacrypt 2017, Part I*, volume 10624 of *LNCS*, pages 65–95. Springer, 2017.
25. Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *SIAM J. Computing*, 36(3):763–778, 2006.
26. Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In *Advances in Cryptology—Crypto 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, 2019.

A Security of Forward-Only Even-Mansour

In this section we consider a simpler case, where $E_k[F](x) := F(x \oplus k)$ for $F : \{0, 1\}^n \rightarrow \{0, 1\}^n$ a uniform function and k a uniform n -bit string. Here we restrict the adversary to forward queries only, i.e., the adversary has classical access to $E_k[F]$ and quantum access to F ; note that $E_k^{-1}[F]$ and F^{-1} may not even be well-defined. As mentioned in the main body, this setting was analyzed in the previously published work [16] as well, using different techniques.

We let \mathcal{F}_n denote the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}^n$.

Theorem 4. *Let \mathcal{A} be a quantum algorithm making q_E classical queries to its first oracle and q_F quantum queries to its second oracle. Then*

$$\left| \Pr_{\substack{k \leftarrow \{0, 1\}^n \\ F \leftarrow \mathcal{F}_n}} [\mathcal{A}^{E_k[F], F}(1^n) = 1] - \Pr_{R, F \leftarrow \mathcal{F}_n} [\mathcal{A}^{R, F}(1^n) = 1] \right| \leq 2^{-n/2} \cdot (2q_E \sqrt{q_F} + 2q_F \sqrt{q_E}).$$

Proof. We make the same assumptions about \mathcal{A} as in the initial paragraphs of the proof of [Theorem 3](#). We also adopt analogous notation for the stages of \mathcal{A} , now using q_E , q_F , and $q_{F,j}$ as appropriate.

Given a function $F : \{0, 1\}^n \rightarrow \{0, 1\}^n$, a set T of pairs where any $x \in \{0, 1\}^n$ is the first element of at most one pair in T , and a key $k \in \{0, 1\}^n$, we define the function $F_{T,k} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ as

$$F_{T,k}(x) := \begin{cases} y & \text{if } (x \oplus k, y) \in T \\ F(x) & \text{otherwise.} \end{cases}$$

Note that, in contrast to the analogous definition in [Theorem 3](#), here the order of the tuples in T does not matter and so we may take it to be a set. Note also that we are redefining the notation $F_{T,k}$ from how it was used in [Theorem 3](#); this new usage applies to this Appendix only.

We now define a sequence of experiments \mathbf{H}_j , for $j = 0, \dots, q_E$:

Experiment \mathbf{H}_j . Sample $R, F \leftarrow \mathcal{F}_n$ and $k \leftarrow \{0, 1\}^n$. Then:

1. Run \mathcal{A} , answering its classical queries using R and its quantum queries using F , stopping immediately before its $(j+1)$ st classical query. Let $T_j = \{(x_1, y_1), \dots, (x_j, y_j)\}$ be the set of all classical queries made by \mathcal{A} thus far and their corresponding responses.
2. For the remainder of the execution of \mathcal{A} , answer its classical queries using $E_k[F]$ and its quantum queries using $F_{T_j,k}$.

We can represent \mathbf{H}_j as the experiment in which \mathcal{A} 's queries are answered using the oracle sequence

$$\underbrace{F, R, F, \dots, R, F}_{j \text{ classical queries}}, \underbrace{E_k[F], F_{T_j,k}, \dots, E_k[F], F_{T_j,k}}_{q_E - j \text{ classical queries}}.$$

Note that \mathbf{H}_0 is exactly the real world (i.e., $\mathcal{A}^{E_k[F], F}$) and \mathbf{H}_{q_E} is exactly the ideal world (i.e., $\mathcal{A}^{R, F}$).

For $j = 0, \dots, q_E - 1$, we define an additional experiment \mathbf{H}'_j :

Experiment \mathbf{H}'_j . Sample $R, F \leftarrow \mathcal{F}_n$ and $k \leftarrow \{0, 1\}^n$. Then:

1. Run \mathcal{A} , answering its classical queries using R and its quantum queries using F , stopping immediately after its $(j+1)$ st classical query. Let $T_{j+1} = ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$ be the set of all classical queries made by \mathcal{A} thus far and their corresponding responses.
2. For the remainder of the execution of \mathcal{A} , answer its classical queries using $E_k[F]$ and its quantum queries using $F_{T_{j+1},k}$.

I.e., \mathbf{H}'_j corresponds to answering \mathcal{A} 's queries using the oracle sequence

$$\underbrace{F, R, F, \dots, R, F, R}_{j \text{ classical queries}}, F_{T_{j+1},k}, \underbrace{E_k[F], F_{T_{j+1},k}, \dots, E_k[F], F_{T_{j+1},k}}_{q_E - j - 1 \text{ classical queries}}.$$

We now show that \mathbf{H}'_j is close to \mathbf{H}_{j+1} and \mathbf{H}_j is close to \mathbf{H}'_j for $0 \leq j < q_E$.

Lemma 9. For $j = 0, \dots, q_E - 1$,

$$|\Pr[\mathcal{A}(\mathbf{H}'_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{j+1}) = 1]| \leq 2 \cdot q_{F,j+1} \sqrt{(j+1)/2^n}.$$

Proof. Given an adversary \mathcal{A} , we construct a distinguisher \mathcal{D} for the “blinding game” of [Lemma 3](#) that works as follows:

Phase 1: \mathcal{D} samples $F, R \leftarrow \mathcal{F}_n$. It then runs \mathcal{A} , answering its quantum queries with F and its classical queries with R , until it replies to \mathcal{A} 's $(j+1)$ st classical query. Let $T_{j+1} = \{(x_1, y_1), \dots, (x_{j+1}, y_{j+1})\}$ be the set of classical queries made by \mathcal{A} and their responses. \mathcal{D} defines algorithm \mathcal{B} as follows: on randomness $k \in \{0, 1\}^n$, output $B = \{(x_j \oplus k, y_j)\}_{j=1}^{j+1}$. Finally, \mathcal{D} outputs F and \mathcal{B} .

Phase 2: \mathcal{D} is given quantum access to a function F_b . It continues to run \mathcal{A} , answering its quantum queries with F_b until \mathcal{A} makes its next classical query.

Phase 3: \mathcal{D} is given the randomness k used to run \mathcal{B} . It continues running \mathcal{A} , answering its classical queries with $E_k[F]$ and its quantum queries with $F_{T_{j+1},k}$. Finally, \mathcal{D} outputs whatever \mathcal{A} outputs.

When $b = 0$ (so $F_b = F_0 = F$), then \mathcal{A} 's output is identically distributed to its output in \mathbf{H}_{j+1} . On the other hand, when $b = 1$ then $F_b = F_1 = F^{(B)} = F_{T_{j+1},k}$ and so \mathcal{A} 's output is identically distributed to its output in \mathbf{H}'_j . The expected number of queries made by \mathcal{D} in phase 2 when $F = F_0$ is the expected number of queries made by \mathcal{A} in stage $(j+1)$ in \mathbf{H}_{j+1} . Since \mathbf{H}_{j+1} and \mathbf{H}_{q_E} are identical until after the $(j+1)$ st stage, this is precisely $q_{F,j+1}$. Because k is uniform, we can apply [Lemma 3](#) with $\epsilon = (j+1)/2^n$. The lemma follows. \square

Lemma 10. For $j = 0, \dots, q_E$,

$$|\Pr[\mathcal{A}(\mathbf{H}_j) = 1] - \Pr[\mathcal{A}(\mathbf{H}'_j) = 1]| \leq 1.5 \cdot \sqrt{q_F/2^n}.$$

Proof. From any adversary \mathcal{A} , we construct a distinguisher \mathcal{D} for the game of [Lemma 4](#). \mathcal{D} works as follows:

Phase 1: \mathcal{D} is given quantum access to a (random) function F . It samples $R \leftarrow \mathcal{F}_n$ and then runs \mathcal{A} , answering its quantum queries using F and its classical queries using R , until \mathcal{A} submits its $(j+1)$ st classical query x_{j+1} . At that point, let $T_j = \{(x_1, y_1), \dots, (x_j, y_j)\}$ be the set of input/output pairs \mathcal{A} has received from its classical oracle thus far.

Phase 2: \mathcal{D} is given (uniform) $s \in \{0, 1\}^n$ and quantum oracle access to a function F_b . \mathcal{D} sets $k := s \oplus x_{j+1}$, and then continues running \mathcal{A} , answering its classical queries (including the $(j+1)$ st) using $E_k[F_b]$ and its quantum queries using the function $(F_b)_{T_j,k}$, i.e.,

$$x \mapsto \begin{cases} y & \text{if } (x \oplus k, y) \in T_j \\ F_b(x) & \text{otherwise.} \end{cases}$$

Finally, \mathcal{D} outputs whatever \mathcal{A} outputs.

We analyze the execution of \mathcal{D} in the two cases of the game of [Lemma 4](#). In either case, the quantum queries of \mathcal{A} in stages $0, \dots, j$ are answered using a random function F , and \mathcal{A} 's first j classical queries are answered using an independent random function R . Note further that since s is uniform, so is k .

Case 1: $b = 0$. In this case, all the remaining classical queries of \mathcal{A} (i.e., from the $(j+1)$ st on) are answered using $E_k[F]$, and the remaining quantum queries of \mathcal{A} are answered using $F_{T_j, k}$. The output of \mathcal{A} is thus distributed identically to its output in \mathbf{H}_j in this case.

Case 2: $b = 1$. Here, $F_b = F_1 = F_{s \rightarrow y}$ for a uniform y . Now, the response to the $(j+1)$ st classical query of \mathcal{A} is

$$E_k[F_b](x_{j+1}) = E_k[F_{s \rightarrow y}](x_{j+1}) = F_{s \rightarrow y}(k \oplus x_{j+1}) = F_{s \rightarrow y}(s) = y.$$

Since y is uniform and independent of anything else, and since \mathcal{A} has never previously queried x_{j+1} to its classical oracle, this is equivalent to answering the first $j+1$ classical queries of \mathcal{A} using a random function R . The remaining classical queries of \mathcal{A} are also answered using $E_k[F_{s \rightarrow y}]$. However, since $E_k[F_{s \rightarrow y}](x) = E_k[F](x)$ for all $x \neq x_{j+1}$ and \mathcal{A} never repeats the query x_{j+1} , this is equivalent to answering the remaining classical queries of \mathcal{A} using $E_k[F]$.

The remaining quantum queries of \mathcal{A} are answered with the function

$$x \mapsto \begin{cases} y' & \text{if } (x \oplus k, y') \in T_j \\ F_{s \rightarrow y}(x) & \text{otherwise.} \end{cases}$$

This, in turn, is precisely the function $F_{T_{j+1}, k}$, where T_{j+1} is obtained by adding (x_{j+1}, y) to T_j (and thus consists of the first $j+1$ classical queries made by \mathcal{A} and their corresponding responses). Thus, the output of \mathcal{A} in this case is distributed identically to its output in \mathbf{H}'_j .

The number of quantum queries made by \mathcal{D} in phase 1 is at most q_F . The claimed result thus follows from [Lemma 4](#). \square

Using [Lemmas 9](#) and [10](#), and the fact that $\sum_{j=1}^{q_E} q_{F,j} = q_F$, we have

$$\begin{aligned} |\Pr[\mathcal{A}(\mathbf{H}_0) = 1] - \Pr[\mathcal{A}(\mathbf{H}_{q_E}) = 1]| &\leq 1.5q_E \sqrt{q_F/2^n} + 2 \sum_{j=1}^{q_E} q_{F,j} \sqrt{j/2^n} \\ &\leq 1.5q_E \sqrt{q_F/2^n} + 2\sqrt{q_E/2^n} \sum_{j=1}^{q_E} q_{F,j} \\ &\leq 1.5q_E \sqrt{q_F/2^n} + 2q_F \sqrt{q_E/2^n}, \end{aligned}$$

as required. \square

B Further Details for the Proof of [Lemma 7](#)

B.1 Equivalence of Expt'_j and \mathbf{H}'_j

The code in the top portion of [Figure 2](#) is a syntactic rewriting of Expt'_j . (Flags that have no effect on the output of \mathcal{A} are omitted.) In [line 27](#), the computation of y_{j+1} has been expanded (note that $E_k[P_1](x_{j+1}) = P_1(s_0) \oplus k_2 = P(s_1) \oplus k_2$). In [line 31](#), Q has been replaced with $P_{T_{j+1}, k}$ and \mathcal{O} has been replaced with $E_k[P]$ as justified in the proof of [Lemma 7](#).

The code in the middle portion of [Figure 2](#) results from the following changes: first, rather than sampling uniform s_0 and then setting $k_1 := s_0 \oplus x_{j+1}$, the code now samples a uniform k_1 . Similarly, rather than choosing uniform s_1 and then setting $y_{j+1} := P(s_1) \oplus k_2$, the code now samples a uniform y_{j+1} (note that P is a permutation, so $P(s_1)$ is uniform). Since neither s_0 nor s_1 is used anywhere else, each can now be omitted.

The code in the bottom portion of [Figure 2](#) simply chooses $k = (k_1, k_2)$ according to distribution D , and chooses uniform $y_{j+1} \in \{0, 1\}^n \setminus \{y_1, \dots, y_j\}$. It can be verified by inspection that this final experiment is equivalent to \mathbf{H}'_j .

<p>23 $P, R \leftarrow \mathcal{P}_n$</p> <p>24 Run \mathcal{A} with quantum access to P and classical access to R, until \mathcal{A} makes its $(j+1)$st classical query x_{j+1}; let T_j be as in the text</p> <p>25 $s_0, s_1 \leftarrow \{0, 1\}^n$</p> <p>26 $k_1 := s_0 \oplus x_{j+1}, k_2 \leftarrow D_{ k_1}, k := (k_1, k_2)$</p> <p>27 $y_{j+1} := P(s_1) \oplus k_2$</p> <p>28 if $y_{j+1} \in \{y_1, \dots, y_j\}$ then $y_{j+1} \leftarrow \{0, 1\}^n \setminus \{y_1, \dots, y_j\}$</p> <p>29 Give y_{j+1} to \mathcal{A} as the answer to its $(j+1)$st classical query</p> <p>30 $T_{j+1} := ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$</p> <p>31 Continue running \mathcal{A} with quantum access to $P_{T_{j+1}, k}$ and classical access to $E_k[P]$</p>
<p>32 $P, R \leftarrow \mathcal{P}_n$</p> <p>33 Run \mathcal{A} with quantum access to P and classical access to R, until \mathcal{A} makes its $(j+1)$st classical query x_{j+1}; let T_j be as in the text</p> <p>34 $k_1 \leftarrow \{0, 1\}^n, k_2 \leftarrow D_{ k_1}, k := (k_1, k_2), y_{j+1} \leftarrow \{0, 1\}^n$</p> <p>35 if $y_{j+1} \in \{y_1, \dots, y_j\}$ then $y_{j+1} \leftarrow \{0, 1\}^n \setminus \{y_1, \dots, y_j\}$</p> <p>36 Give y_{j+1} to \mathcal{A} as the answer to its $(j+1)$st classical query</p> <p>37 $T_{j+1} := ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$</p> <p>38 Continue running \mathcal{A} with quantum access to $P_{T_{j+1}, k}$ and classical access to $E_k[P]$</p>
<p>39 $P, R \leftarrow \mathcal{P}_n$</p> <p>40 Run \mathcal{A} with quantum access to P and classical access to R, until \mathcal{A} makes its $(j+1)$st classical query x_{j+1}; let T_j be as in the text</p> <p>41 $k \leftarrow D, y_{j+1} \leftarrow \{0, 1\}^n \setminus \{y_1, \dots, y_j\}$</p> <p>42 Give y_{j+1} to \mathcal{A} as the answer to its $(j+1)$st classical query</p> <p>43 $T_{j+1} := ((x_1, y_1), \dots, (x_{j+1}, y_{j+1}))$</p> <p>44 Continue running \mathcal{A} with quantum access to $P_{T_{j+1}, k}$ and classical access to $E_k[P]$</p>

Fig. 2. Syntactic rewritings of Expt'_j .

B.2 Handling an Inverse Query

In this section we discuss the case where the $(j+1)$ st classical query of \mathcal{A} is an inverse query in the proof of [Lemma 7](#). Phase 1 is exactly as described in the proof of [Lemma 7](#), though we now let y_{j+1} denote the $(j+1)$ st classical query made by \mathcal{A} (assumed to be in the inverse direction).

Phase 2: \mathcal{D} receives $s_0, s_1 \in \{0, 1\}^n$ and quantum oracle access to a permutation P_b . First \mathcal{D} sets $t_0 := P_b(s_0)$ and $t_1 := P_b(s_1)$. It then sets $k_2 := t_0 \oplus y_{j+1}$, chooses $k_1 \leftarrow D_{|k_2}$ (where this represents the conditional distribution on k_1 given k_2), and sets $k := (k_1, k_2)$. \mathcal{D} continues running \mathcal{A} , answering its remaining classical queries (including the $(j+1)$ st one) using $E_k[P_b]$, and its remaining quantum queries using

$$(P_b)_{T_j, k} = \text{swap}_{P_b(x_1 \oplus k_1), y_1 \oplus k_2} \circ \dots \circ \text{swap}_{P_b(x_j \oplus k_1), y_j \oplus k_2} \circ P_b.$$

Finally, \mathcal{D} outputs whatever \mathcal{A} outputs.

Note that t_0, t_1 are uniform, and so k is distributed according to D . Then:

Case $b = 0$ (no reprogramming). In this case, \mathcal{A} 's remaining classical queries (including its $(j+1)$ st classical query) are answered using $E_k[P_0] = E_k[P]$, and its remaining quantum queries are answered using $(P_0)_{T_j, k} = P_{T_j, k}$. The output of \mathcal{A} is thus distributed identically to its output in \mathbf{H}_j in this case.

Case $b = 1$ (reprogramming). In this case, $k_2 = P_1(s_0) \oplus y_{j+1} = P(s_1) \oplus y_{j+1}$ and so

$$\begin{aligned} P_b^{-1} &= P_1^{-1} = (P \circ \text{swap}_{s_0, s_1})^{-1} = (\text{swap}_{P(s_0), P(s_1)} \circ P)^{-1} \\ &= P^{-1} \circ \text{swap}_{P(s_0), P(s_1)} \\ &= P^{-1} \circ \text{swap}_{P(s_0), y_{j+1} \oplus k_2}. \end{aligned}$$

The response to \mathcal{A} 's $(j + 1)$ st classical query is thus

$$x_{j+1} \stackrel{\text{def}}{=} E_k^{-1}[P_1](y_{j+1}) = P_1^{-1}(y_{j+1} \oplus k_2) \oplus k_1 = P_1^{-1}(P(s_1)) \oplus k_1 = s_0 \oplus k_1.$$

The remaining classical queries of \mathcal{A} are then answered using $E_k[P_1]$, while its remaining quantum queries are answered using $(P_1)_{T_j, k}$.

Now we define the following three events:

1. bad_1 is the event that $x_{j+1} \in \{x_1, \dots, x_j\}$.
2. bad_2 is the event that $P(s_0) \oplus k_2 \in \{y_1, \dots, y_j\}$.
3. bad_3 is the event that, in phase 2, \mathcal{A} queries its classical oracle in the forward direction on $s_1 \oplus k_1$, or the inverse direction on $P(s_0) \oplus k_2$.

Comparing the above to the proof of Lemma 7, we see (because P is a permutation) that the situation is entirely symmetric, and the analysis is therefore the same.