

Multi-Reprogramming with Quantum Advice in the Quantum Random Oracle Model

Yu-Hsuan Huang¹  and Heming Liao^{2,3} 

¹ Max Planck Institute for Security and Privacy (MPI-SP), Bochum, Germany.

² State Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, Chinese Academy of Sciences, Beijing 100085, China.

³ School of Cyber Security, University of Chinese Academy of Sciences, Beijing 100049, China.

yu-hsuan.huang@mpi-sp.org, liaoheming@iie.ac.cn

Abstract. Adaptive reprogramming is a central tool for proving security in the quantum random oracle model (QROM). Recently, it has been extended to the setting when an oracle-dependent advice is in presence. However, the existing result is restricted to classical advice, and addresses only a single reprogramming. Extending it to n reprogrammings through naive hybridization yields a loss that is linear in the number of reprogrammings.

In this paper, we establish an adaptive reprogramming theorem that accommodates quantum advice, and has a tighter overall security loss that only scales in the cubic-root of n , with an attack that matches the number of queries required in order to reach constant advantage.

Using our reprogramming theorem, we provide improved time-space trade-off bounds for PRG security of random oracles, and an IND-CPA security proof of salted pseudo one-time pad, in the quantum advice setting. Along the way, we also identify and prove a folklore lemma about indistinguishability of two oracles, which may be of independent interest.

As another application, we provide a generic security analysis of Fiat-Shamir and hash-and-sign signature schemes, in QROM with quantum advice. We do so via a two-step modular analysis, where in the first step we perform a generic UF-CMA to UF-NMA reduction using our reprogramming theorem, and in the second step we provide a characterization of when such a scheme satisfies UF-NMA security, which we then use to discuss security of several signature schemes in this model.

1 Introduction

Post-quantum cryptography. Since the proposal of Shor’s algorithm in 1994 [Sho94], the design and analysis of cryptographic schemes that resist quantum attacks — also known as *post-quantum cryptography* — have become an explicit goal. In December 2016, the US National Institute of Standards and Technology (NIST) launched a competition [NIST17], in which experts around the world join forces in order to select the future post-quantum cryptographic standards. As an important milestone, four winner schemes have been selected [SAB⁺22, LDK⁺22,

PFH⁺22,HBD⁺22], and an additional call for digital signature schemes [NIST26] has now advanced to the third round.

Following these efforts, industry has also begun transitioning toward post-quantum cryptography, with major technology companies deploying post-quantum primitives across communication systems: Apple has introduced PQ3 for iMessage, while Google and Microsoft have announced post-quantum deployments and migration roadmaps for cloud and security services⁴, which further motivates the development of efficient and (provably) secure post-quantum schemes. Moreover, the use of hybrid cryptographic schemes that use classical and post-quantum schemes simultaneously [GHP18,BCD⁺24], has also been recently popularized, and can be used in standardized protocols such as IKEv2 [KHN⁺14] and TLS 1.3 [Res26].

Random oracle model. In cryptography, it is common to perform a security proof in an idealized model, since otherwise such a proof is often out of reach (or in some cases, impossible [GO94, CGH98, Nie02, GW11, MW20]) in the plain, non-idealized model, or resulting in a much more inefficient scheme. Among others, the random oracle model (ROM) [BR93] has become one of the standard approaches in analyzing practical cryptographic schemes that use a hash function. In this model, the hash function H is idealized, treated as a random function, and correspondingly the attackers are then allowed to make oracle queries in order to learn its function values. The main benefit of this model is that it facilitates certain useful proof strategies. For example, a reduction may observe the queries made by the attacker; less obviously, a reduction also gains some power in modifying responses of the oracle, which we will discuss in later paragraphs.

Toward achieving security against quantum attacks, it is then necessary to consider a stronger form of queries that can be made in “quantum superposition.” This formalization is known as the *quantum random oracle model* (QROM) [BDF⁺11]. Indeed, almost all schemes in the NIST competition, including winners and candidates in each round, rely on security analyses in the ROM, with many of them also treated in the QROM [BBD⁺23, FFH26, KLS18, FGdG⁺26, BDK⁺18].

Advice in the random oracle model. From a theoretical standpoint, though, the ROM is still rather heuristic. Among refinements that are aimed at more realism, the incorporation of *oracle-dependent advice* has proved particularly fruitful. In practice, the design of a hash function is public and fixed long before (the attack of) a higher-level scheme could take place. Reasonably then, such an attack may receive preprocessed data about the hash function, to help it break the actual scheme at hand. This motivates the finer-grained security model, in which an attacker is given an auxiliary (quantum or classical) input called *advice*,

⁴ See Apple PQ3; Google Cloud KMS; Google Cloud PQC roadmap; and Microsoft quantum-safe security.

that is bounded in length, and may be arbitrarily dependent on the randomness of the oracle.

The ROM in the presence of advice, is analogous to the standard non-uniform model of computation (like circuit family), and is often considered more realistic in the cryptographic setting. For example, without advice, the random oracle itself is collision resistant, while in reality no fixed hash function is collision resistant, in the existential sense, because an attacker could have simply hardcoded a collision pair. It has also been substantially studied in the literature [DTT10, DGK17, CDGS18, ADGL23, AGL24], dating back to Hellman’s time-space tradeoff attack against function inversion [Hel80]. On the quantum provable security side, there are a number of different techniques, such as ones based on multi-instance game [CGLQ20, DFH⁺24], bit-fixing model [Unr07a, GLLZ21, Liu23], and indistinguishability with precomputation [CPZ24], and more. In other idealized models, such as the random permutation model, analogous notions have also been considered [CDG18, FGK22, ADGL23].

Reprogramming in the QROM. In the ROM, a reduction may be able to tweak certain oracle responses $H(x)$ dealt to the attacker in a query, without being noticed by the attacker. This is known as *reprogramming* of the random oracle — a methodology that is fundamental in many security proofs in the ROM. For example, a reduction algorithm may pick an input x which has not been queried before, or will not be queried afterward, by the attacker, and then rewrite the oracle output $H(x) := y$ to a different output of a suitable distribution. Since the attacker cannot compare the original, unprogrammed value of $H(x)$ with y , such reprogramming remains unnoticeable.

In the quantum setting (QROM), however, situation becomes much trickier, as an attacker may query the oracle in superposition over the entire domain, and consequently introduces subtle inconsistencies within such a query. Over the past decade and more, a substantial amount of research effort has been spent, in order to realize (various flavours of) reprogramming also in the quantum setting. This includes the family of one-way-to-hiding (O2H) lemmata [Unr14b, AHU19, BHH⁺19, KSS⁺20, GLX24], measure-and-reprogram techniques [DFMS19, DFM20], and the main topic of this work: the adaptive reprogramming theorems [Unr14a, Unr15, GHM21, PZ24, Jae25, LGXZ26].

The aforementioned reprogramming methods have also been applied, so as to provide QROM security proofs for many cryptographic applications. On the signature side, the adaptive reprogramming and measure-and-reprogram methods can be used to provide quantum security proofs for Fiat-Shamir transformation, or hash-and-sign design principles [KLS18, DFMS19, DFM20]. Popular variants such as Fiat-Shamir with aborts (FSwA) and hash-and-sign with aborts/retry (HSwA) have also been treated using more sophisticated reprogramming techniques [BBD⁺23, FFH26, Hua26]. On the encryption and KEM side, these techniques are used in analyses of the Fujisaki-Okamoto transformation [TU16, HHK17, JZC⁺18, BHH⁺19, KSS⁺20, HHM22, DFMS22, SGX23, GSX23, GLX24] and OAEP [TU16, Ebr22, LGCX25, LLHC26]. Recently, reprogramming methods in QROM also found its applications in permutation-based

cryptography [HHM⁺26]. These results demonstrate the broad applicability of reprogramming in the standard QROM.

Reprogramming with advice. A natural next step is to understand to what extent the aforementioned reprogramming methods extend to the advice setting. As a first attempt, the recent work [HHM⁺26] has taken the lead toward generalizing the adaptive reprogramming theorem into the advice setting. Specifically, they show that for an input X sampled from an (adversarially chosen) distribution of high min-entropy, if a reduction reprograms the random oracle as $H(X) := Y$ for a freshly chosen uniform randomness Y , an attacker given advice is unlikely to notice whether the reprogramming occurred.

However, the above work has two limitations:

First, it only applies when the attacker’s advice is restricted to a classical bit string. Such a restriction is rather undesirable, since in general an attacker can be given or work out an advice as a piece of quantum information. Moreover, as has been noted in [Liu23], in some (contrived) security game, obtaining quantum advice would indeed result in an exponentially more powerful attacker than in the classical advice case. Naturally, one would ask:

Can adaptive reprogramming theorems in QROM (with classical advice) be extended to the quantum advice setting?

Second, on a more technical front, the obtained security guarantee has only taken into account *one single reprogramming point*, with a concrete bound on the order of $O(\sqrt[3]{q_H L \epsilon})$. Here, $\epsilon > 0$ is a statistical parameter close to zero, L is the length of the advice, and q_H is the query complexity of the attacker. Though, in prominent use cases of adaptive reprogramming, such as Fiat-Shamir signature schemes [BBD⁺23, FFH26], one would often need to consider the case where reprogramming is done in multiple, say q_S , points. If one repeatedly invokes the reprogramming theorem in the course of a naive hybridization, the above would have resulted in a bound that scales on the order of $O(q_S \sqrt[3]{q_H L \epsilon})$. This seems to be rather non-tight, since in the no-advice setting, a much tighter bound accompanied with a matching attack of order $\Theta(\sqrt{q_S q_H \epsilon})$ can be shown. Hence, one may also ask:

Can one achieve a tight(er) adaptive reprogramming theorem in QROM with quantum advice, when considering multiple reprogramming points?

Considering the broad applicability of the adaptive reprogramming method in the no-advice setting, if the above questions can indeed be addressed, then one would ideally hope to use it to provide security analyses for practical cryptographic schemes that were only treated in the no-advice setting. It is therefore natural to wonder:

Using adaptive reprogramming methods, can one prove post-quantum security of useful cryptographic schemes, or improve state-of-the-art query complexity bound, in QROM with quantum advice?

1.1 Our Contributions

In this paper, we answer all of the above questions in the affirmative. Our contributions are three-fold, which we address in different paragraphs below.

Adaptive reprogramming with quantum advice. Our main technical contribution is that we provide an adaptive reprogramming theorem, in the style of [GHHM21, HHM⁺26], that works in QROM with quantum advice, and achieves an (almost) tight distinguishing bound. (See Table 1 for summary.)

Informally, what we show is the following: suppose an oracle algorithm $\mathcal{A}$ makes q_H quantum queries to a random oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$, is given a piece of H -dependent quantum advice of L qubits, and can sample q_S reprogramming points $X_1, \dots, X_{q_S}$, each with at least $\log(1/\epsilon)$ bits of min-entropy. Then, $\mathcal{A}$ cannot distinguish whether, throughout the game, all or none of the points chosen are reprogrammed $H(X_i) := Y_i \leftarrow \mathcal{Y}$ with a fresh randomness, except with a small distinguishing advantage of order $O(\sqrt[3]{q_S q_H L \epsilon})$.⁵

As already discussed, the prior work on reprogramming with advice [HHM⁺26] is restricted to classical advice (whereas ours holds for quantum advice), and applies to a single reprogramming point only. A naive hybridization over q_S reprogrammings would incur an additive loss summing up to $O(q_S \sqrt[3]{q_H L \epsilon})$. Therefore, our bound improves this loss by moving the factor q_S inside the cubic-root.

We also give a classical attack achieving advantage $\Omega(\sqrt{q_S q_H L \epsilon})$, which matches our bound in the resources (in q_S , q_H , and L) taken to reach constant distinguishing advantage. Though, there is still a gap between the cubic-root bound and the square-root attack.

	Bound	Advice	Attack
[HHM ⁺ 26]	$O(q_S \sqrt[3]{q_H L \epsilon})$ ⁶	classical	✗
Our work Theorem 1	$O(\sqrt[3]{q_S q_H L \epsilon})$	quantum	$\Omega(\sqrt{q_S q_H L \epsilon})$

Table 1. Reprogramming bounds in QROM when given L (qu)bits of advice, where q_H and q_S are the number of random oracle queries and reprogramming points respectively (assuming $q_H \geq q_S$), and each reprogramming point has $\log(1/\epsilon)$ bits of min-entropy.

Time-space tradeoff bounds for symmetric primitives. As an application of our reprogramming theorem with quantum advice, we improve state-of-the-art time-space tradeoff bounds, for PRG security of random oracles with quantum advice [Liu23]. Specifically, we show that, for n independent PRG challenge keys $K_1, \dots, K_n \leftarrow \mathcal{D}$ sampled according to some distribution $\mathcal{D}$ of high min-entropy, the random oracle outputs $(H(K_1), \dots, H(K_n))$ are indistinguishable to n independent uniform randomness, with advantage bounded by

⁵ Here we assume, as is typical, that $q_H \geq q_S$. See Theorem 1 for the concrete bound.

⁶ [HHM⁺26] only provides a bound for single reprogramming; the bound stated here is obtained via a hybrid argument.

$O(\sqrt[3]{nq_H L\epsilon} + nq_H^2\epsilon)$ (assuming $q_H \geq n$). Already in the special case where $n = 1$ our result yields a bound $O(\sqrt[3]{q_H L\epsilon} + q_H^2\epsilon)$, improving over Liu’s bound of order $O(\sqrt[3]{q_H L\epsilon} + q_H\sqrt{\epsilon})$. For general n , our bound also constitutes an improvement over the naive hybridization of Liu’s result, which would in total sum up to $O(n\sqrt[3]{q_H L\epsilon} + nq_H\sqrt{\epsilon})$, with n outside of the cubic/square roots. Using a similar approach, we also prove the IND-CPA security of salted pseudo one-time pad in QROM, when there is quantum advice. (See Table 2 for summary.)

	Single-challenge	n -challenge
[Liu23]	$O(\sqrt[3]{q_H L\epsilon} + q_H\sqrt{\epsilon})$	$O(n\sqrt[3]{q_H L\epsilon} + nq_H\sqrt{\epsilon})^7$
Our work Theorem 3	$O(\sqrt[3]{q_H L\epsilon} + \textcolor{yellow}{q_H^2\epsilon})$	$O(\sqrt[3]{nq_H L\epsilon} + n\textcolor{yellow}{q_H^2\epsilon})$

Table 2. PRG security bounds for random oracle with L qubits of quantum advice, in the single-challenge and n -challenge settings, where q_H is the number of (quantum) random oracle queries (assuming $q_H \geq n$), and ϵ denotes $\log(1/\epsilon)$ bits of min-entropy.

Fiat-Shamir and hash-and-sign signatures. Our last major contribution is a thorough analysis of Fiat-Shamir and hash-and-sign signature schemes in QROM with quantum advice. Like in [FFH26], we consider a generalization of Fiat-Shamir transformation $\text{FS}[\Sigma]$ (and the underlying Sigma protocol Σ). We then formally analyze the security of this generalized scheme $\text{FS}[\Sigma]$, which covers both Fiat-Shamir and hash-and-sign signature schemes simultaneously.

Our analysis is performed in two modular steps:

In the first step, we use our reprogramming theorem to give a generic security reduction from UF-CMA to UF-NMA in QROM (with a quantum advice ρ of L qubits), of the following form:

$$\text{Adv}_{\text{FS}[\Sigma]}^{\text{UF-CMA}}(\mathcal{A}, \rho) \leq \text{Adv}_{\text{FS}[\Sigma]}^{\text{UF-NMA}}(\mathcal{B}, \rho) + O(\sqrt[3]{q_S q_H L\epsilon}).$$

Here, we denote the UF-CMA and UF-NMA advantages using the Adv notation, $\epsilon > 0$ is a scheme-dependent parameter that is close to zero, and q_S, q_H are the number of queries that attacker is allowed to make to the signing oracle and the random oracle. Using prior state-of-the-art reprogramming bound ([HHM⁺26, Theorem 5]) would give such a reduction when the advice is restricted to a *classical string*. Moreover, the obtained bound would be of order $O(q_S \sqrt[3]{q_H L\epsilon})$, with a q_S factor loss *outside* of the cubic-root. Concretely, for NIST security requirement, q_S can be all the way up to 2^{64} . Thus, suppressing less relevant terms, our bound constitutes an improvement in tightness by a factor of $q_S^{2/3} \approx 2^{42.6}$.

Second, we give a characterization of UF-NMA security of $\text{FS}[\Sigma]$, in the presence of (quantum or classical) advice. Specifically, under certain nontriviality

⁷ [Liu23] only provides a bound for single challenge; the bound stated here is obtained via a naive hybrid argument.

condition $\neg \text{Trvl}$, we show that $\text{FS}[\Sigma]$ satisfies UF-NMA as long as Σ satisfies the standard identification security (see Section 5.3). Conversely, if there is an attack breaking either of the aforementioned security properties of Σ , then there is also an attack breaking the UF-NMA security of $\text{FS}[\Sigma]$ with advice. This gives us a useful way to predict exactly when the considered $\text{FS}[\Sigma]$ scheme is secure.

In Section 5.4 we discuss UF-NMA security of several concrete instantiations in QROM with (classical or quantum) advice. Specifically, we show that (non-public-key-hashing) Falcon is insecure when there’s a classical advice, while some other schemes such as a GPV-based signature scheme or (a simplified version of) Wave achieve provable security (UF-NMA) in QROM with quantum advice.

1.2 Technical Overview

In this technical overview, we briefly discuss the challenges and ideas of our main results. This is not an exhaustive overview of all our technical results.

Adaptive Reprogramming with quantum advice. At a high level, our reprogramming theorem says that a bounded-query attacker cannot distinguish whether a reprogramming $H(X) := Y \leftarrow \mathcal{Y}$ on a high min-entropy input is done, where we additionally allow $\mathcal{A}$ to obtain a H -dependent quantum advice ρ . The proof of this theorem consists of two major steps.

First, like in the case of [HHM⁺26], we transform this game into a two-stage game, in which the attacker can decide at the first stage whether to abort or not, and conditioned on the event that it doesn’t, distinguish whether the aforementioned reprogramming $H(X)$ is done in the second stage. This two-stage security game is closely related to the bit-fixing model, that was first put forward by [Unr07b], and later also discussed in [Liu23].⁸

Second, we show that any such successful two-stage attacker would only have a small distinguishing advantage. This step can be seen as a generalization of prior analyses performed in [HHM⁺26, Lemma 6] (which follows a similar strategy as in [FH23, Appendix A]), where we allow multiple reprogrammings to take place.

Recall that prior analyses in this setting use Zhandry’s compressed oracle technique (see [Zha19, CFHL21]). There, one exploits the fact that, when performing a measurement whether an input X with high min-entropy is recorded in the “database,” then most likely the outcome is “no,” in which case the reprogramming $H(X) := Y$ with a fresh randomness Y would not have caused any difference. The gentle-measurement lemma would control the disturbance caused by such a measurement, in trace distance $O(\sqrt{q\epsilon})$, which sums up to $O(q_S \sqrt{q\epsilon})$, with a q_S factor outside of the square-root.

⁸ Even though there is a generic lifting theorem that provides qAI-QROM security given its counterpart in (a quantum variant of) the bit-fixing model, it does not directly apply to our setting. This is because we need to take into account the reprogramming oracles that were not captured in the generic theorem.

However, as has been observed by [LGXZ26] (in the no-advice setting), this q_S -linear factor caused by performing gentle measurement lemma query-by-query is not necessary. Instead, with a hybrid argument tailored to our reprogramming game, we are able to benefit from the existing gentle sequential measurement lemma [OV22] (a decision version of quantum union bound [Gao15]), thereby obtaining a better distinguishing bound of order $O(q_S^2\epsilon + \sqrt{q_S q\epsilon})$. This would eventually allow us to obtain the main reprogramming bound of order $O(\sqrt[3]{q_H q_S L\epsilon})$ with the factor q_S inside of the cubic root.

Immediate applications. As an application of our reprogramming theorem, we provide security bounds of several symmetric primitives, such as random-oracle based PRG, or salted pseudo one-time pad, again when (quantum) advice is present.

The analyses of these primitives are similar, in that we proceed through some hybrid sequences where we first reprogram a random oracle $H(X) := Y \leftarrow \mathcal{Y}$ for a high min-entropy X , and then undo this reprogramming to only sampling $Y \leftarrow \mathcal{Y}$ via the fact that X is hidden (or partially hidden) away from the attacker. Thanks to our reprogramming theorem, the first step immediately follows, and the second step can be controlled via the standard one-way to hiding lemma (see Theorem 3 in the full version of [AHU19]).

Furthermore, we can achieve a tight(er) bound using a folklore indistinguishability bound (see Section 4.1), which bounds the advantage of a quantum oracle algorithm in distinguishing the all-zero oracle $\text{MEMBER}_\emptyset$ from $\text{MEMBER}_{\{X\}}$, which returns 1 only on input X .

To the best of our knowledge, this is believed true among the community, but there is no reference that addresses the proof explicitly. Therefore, as a side contribution, we provide a self-contained proof via polynomial method, which may be of independent interest.

Fiat-Shamir and hash-and-sign signature. We analyze Fiat-Shamir (and hash-and-sign) signature schemes $\text{FS}[\Sigma]$ in the presence of quantum advice. As in the standard no-advice setting, our analysis proceeds in two modular steps.

In the first step, we show that any UF-CMA attacker (with advice) against $\text{FS}[\Sigma]$ can be transformed into a UF-NMA attacker. This is via a standard argument, where we simulate each signing query $\text{Sign}(sk, m)$ via simulating the underlying Sigma-protocol transcript using an honest-verifier zero-knowledge simulator, and then reprogramming the transcript consistently to the random oracle. Indeed, the main technical challenge is to show that the simulated oracle cannot be distinguished from the real signing oracle, which can be shown directly, once our main reprogramming theorem is in place.

In the second step, we provide a characterization of the UF-NMA security (with advice). Surprisingly, providing such a security proof is trickier than one would initially expect. Recall that the standard analyses (when there is no advice) of this step are typically done via the *measure-and-reprogram* (M&R) technique. It transforms a UF-NMA attacker, to a similarly efficient attacker that

wins the identification security game of the underlying Sigma protocol Σ . However, not only does M&R not work in the advice setting, the intended statement is actually false! We discover and formalize another property of Σ , which we call *nontriviality* ($\neg \text{Trvl}$), without which the signature scheme $\text{FS}[\Sigma]$ admits a trivial UF-NMA attack (with advice). Indeed, a Sigma protocol as such may satisfy the standard identification security of Σ (so that standard M&R technique applies when there is no advice), and meanwhile still being insecure in the nontriviality property (so that $\text{FS}[\Sigma]$ is insecure in UF-NMA with advice).

Perhaps more surprisingly, these two properties of Σ are also sufficient to prove UF-NMA security of $\text{FS}[\Sigma]$ with advice. Toward this end, we first invoke the generic result in [Liu23, Theorem 5], which transforms a UF-NMA attacker into an attacker in the aforementioned quantum bit-fixing model P -Bit-Fixing-QROM. We then go through a hybrid argument, which effectively distinguishes the case whether the forgery produced by the attacker has been “queried” in the first stage or not. Breaking the bit-fixing game in the former case would require the attacker breaking the nontriviality condition of Σ , while in the latter case, we land at a hybrid game in which the M&R technique can be recycled. Formalizing this argument is more involved, since with quantum queries there is no obvious way to determine (e.g., by inspecting the compressed-oracle database) whether a query belongs to the first or second stage of the attacker.

2 Preliminaries

2.1 Basic Notations

For every positive integer n , we denote by $[n]$ the set $\{1, 2, \dots, n\}$, and for every predicate P , we denote by $\mathbb{1}_P$ the indicator function that outputs 1 whenever P is satisfied, and 0 otherwise. Likewise, in a probability space, where P is an event, $\mathbb{1}_P$ would then be the binary random variable that is set to one whenever the event is satisfied.

2.2 Arrow Notations

For any probability distribution D , we use the arrow notation $x \leftarrow D$ to denote that x is freshly sampled from D . When D is a finite set, we consider the uniform distribution over the set.

Likewise, when D is a procedure or algorithm, we consider the corresponding distribution defined by the output after running D once, i.e. $x \leftarrow D$ for x produced by running D once. Further abusing the notation, if the D has an internal state, we denote by $(y; \sigma_y) \leftarrow D$ for y produced by running D once, and the internal working state of D collapsing to σ_y conditioned on the output y .

Abusing the notation even further, for any fixed value y° of y , we write $y^\circ \leftarrow D$ for the event that $y = y^\circ$, and $(y^\circ; \sigma_{y^\circ}) \leftarrow D$ as a short-hand notation that also specifies the post-measurement state of D conditioned on $y = y^\circ$.

2.3 Probability Theory

Support. Let X be any random variable. If there is a countable set S such that $\Pr[X \in S] = 1$ and $\Pr[X = s] > 0$ for every $s \in S$, then this set is unique, and we call it the support of X , denoted by $\text{Supp}(X)$.

Statistical distance. Let X and Y be any two random variables over $\mathcal{X}$. Their statistical distance, or total-variational distance, is defined as follows:

$$SD(X, Y) := \sup_S |\Pr[X \in S] - \Pr[Y \in S]| ,$$

where the supremum is taken over all measurable⁹ subsets $S \subseteq \mathcal{X}$. If $\mathcal{X}$ is countable, then it can be rewritten as

$$SD(X, Y) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr[X = x] - \Pr[Y = x]| .$$

As a short-hand notation, for algorithms $\mathcal{A}$ and $\mathcal{B}$, we denote by $SD(\mathcal{A}, \mathcal{B})$ the statistical distance of their outputs, i.e. $SD(X, Y)$ for $X \leftarrow \mathcal{A}$ and $Y \leftarrow \mathcal{B}$.

When X and Y are within the same probability space, or when they can be unambiguously re-defined in a single probability space, we denote by $SD(X, Y | E)$, for any event E , the statistical distance conditioned on this event.

Guessing probability and min-entropy. Let X be any random variable of finite support $\mathcal{X}$. Its guessing probability is defined by

$$\text{guess}(X) := \max_{x \in \mathcal{X}} \Pr[X = x] ,$$

and its min-entropy is simply $H_{\min}(X) := -\log \text{guess}(X)$. For any event E , we denote by $\text{guess}(X | E)$ and $H_{\min}(X | E)$ the analogous quantity in the conditional probability space conditioned on E .

Let Y be any other random variable of finite support $\mathcal{Y}$, the conditional guessing probability of X given Y is defined by

$$\text{guess}(X | Y) := \sum_{y \in \mathcal{Y}} \Pr[Y = y] \text{guess}(X | Y = y) = \max_{f: \mathcal{Y} \rightarrow \mathcal{X}} \Pr[f(Y) = X] .$$

Likewise, we define $H_{\min}(X | Y) := -\log \text{guess}(X | Y)$ for min-entropy, and $\text{guess}(X | Y, E)$ and $H_{\min}(X | Y, E)$ for the analogous version conditioned on an event E .

⁹ As is the standard probability theory, we implicitly assume here that the sigma algebras underlying X and Y are the same, i.e. $\Pr[X \in S]$ is well-defined whenever $\Pr[Y \in S]$ is, and vice versa.

2.4 Security of Signature Schemes with Advice

A digital signature scheme $\mathcal{S} = (\text{KGen}, \text{Sign}, \text{Ver})$ consists of three algorithms as listed. The standard security notion to consider is UF-CMA, and typically one also considers a weaker notion known as UF-NMA.

In this paper, we are interested in the setting where algorithms underlying $\mathcal{S}$ may use a hash function H , which we model as a random oracle. An attacker would then be given quantum query access to the oracle H , and additionally, we give it as input a quantum state (advice) ρ that may be arbitrarily dependent on H . In Fig. 1 below, we recall the UF-CMA and UF-NMA security games, when H and ρ are in presence.

UF-CMA(ρ):	UF-NMA(ρ):
1: $(sk, pk) \leftarrow \text{KGen}$	1: $(sk, pk) \leftarrow \text{KGen}$
2: $(m, \sigma) \leftarrow \mathcal{A}^{H, \text{Sign}(sk, \cdot)}(\rho, pk)$	2: $(m, \sigma) \leftarrow \mathcal{A}^H(\rho, pk)$
3: $\{m_i := i\text{th signing query}\}$	3: return $\text{Ver}(pk, m, \sigma)$
4: return $\text{Ver}(pk, m, \sigma) \wedge m \notin \{m_i\}_i$	

Fig. 1. The UF-CMA and UF-NMA games with an H -dependent quantum advice ρ .

Informally, we say that the scheme $\mathcal{S}$ satisfies UF-CMA and UF-NMA with quantum advice, if for every efficient attacker $\mathcal{A}$, and for every H -dependent quantum state ρ that is bounded in (qubit) length, the corresponding advantages

$$\begin{aligned} \text{Adv}_{\mathcal{S}}^{\text{UF-CMA}}(\mathcal{A}, \rho) &:= \Pr[1 \leftarrow \text{UF-CMA}(\rho)] , \\ \text{Adv}_{\mathcal{S}}^{\text{UF-NMA}}(\mathcal{A}, \rho) &:= \Pr[1 \leftarrow \text{UF-NMA}(\rho)] , \end{aligned}$$

are small respectively.

2.5 Security of Symmetric Encryption Scheme with Advice

Likewise, we recall the IND-CPA security game for a symmetric encryption $\mathcal{E} := (\text{KGen}, \text{Enc}, \text{Dec})$ scheme below, when such advice ρ is in presence. In addition, we allow an attacker to query the challenge oracle multiple times, this would not only simplify our exposition, but also make our positive result stronger.

IND-CPA $_b$ (ρ)	Chall(m_0, m_1)
1: $K \leftarrow \text{KGen}$	1: $C \leftarrow \text{Enc}(K, m_b)$
2: return $\mathcal{A}^{H, \text{Chall}}(\rho)$	2: return C

Fig. 2. The IND-CPA security game, when quantum advice is present.

Informally, we say that a scheme $\mathcal{E}$ is IND-CPA secure if, for every efficient attacker $\mathcal{A}$ (of binary output), the following advantage is small:

$$\text{Adv}_{\mathcal{E}}^{\text{IND-CPA}}(\mathcal{A}, \rho) := SD(\text{IND-CPA}_0(\rho), \text{IND-CPA}_1(\rho)) .$$

2.6 Compressed Oracle Technique

The compressed oracle technique [Zha19] (see [CFHL21] for a simpler exposition) provides a perfect, efficient, and stateful simulator cO of a random oracle H , say from $\mathcal{X} \rightarrow \mathcal{Y}$. Let $\mathbf{D}$ be internal working register of cO , whose state is supported by all functions of the form $D : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$, which we call *databases*. The register $\mathbf{D}$ consists of a family of subregisters $\mathbf{D}_x$, each supported by $\mathcal{Y} \cup \{\perp\}$ and stores the value of $D(x)$, so that we can rewrite

$$|D\rangle_{\mathbf{D}} = \bigotimes_{x \in \mathcal{X}} |D(x)\rangle_{\mathbf{D}_x} .$$

In each of the (quantum) queries, the simulator cO implements the following unitary

$$\text{cO} : |x, y, D\rangle \mapsto |x\rangle \otimes \text{comp}_x \cdot \mathbf{O}_x \cdot \text{comp}_x |y, D\rangle$$

where the *compression operator* comp_x operates on $\mathbf{D}_x$, swaps the states $|\perp\rangle$ and $|\hat{0}\rangle = (1/\sqrt{|\mathcal{Y}|}) \sum_{y \in \mathcal{Y}} |y\rangle$, and fixes the orthogonal complements, and the operator $\mathbf{O}_x$ simply performs a CNOT from $\mathbf{D}_x$ to the register containing y .

Intuitively, a database D captures all the query input-output pairs that “the attacker has made,” and the special symbol $\perp$ indicates a yet-to-be specified value of the random oracle.

Useful measurements for compressed oracle. At any time in a security game, one may perform a measurement $\mathcal{M}^{\text{full}}(\mathbf{D})$ that decompresses $\mathbf{D}$ and then performs a full measurement in the computational basis. Namely, the projector that corresponds to the output H is defined by

$$\Pi^{\text{full}} := \text{comp} |H\rangle \langle H| \text{comp} ,$$

where $\text{comp} := \bigotimes_x \text{comp}_x$. It is then guaranteed that, whenever this measurement is performed, one can use the output H as the random oracle to continue the security game, and the attacker cannot notice such a measurement at all.

There is another useful measurement, which we call $\mathcal{M}_x^{\perp}$, on $\mathbf{D}$ checking whether an input x is recorded in the database. Specifically, the positive outcome of $\mathcal{M}_x^{\perp}(\mathbf{D})$ can be described by the following (orthogonal) projector

$$\Pi_x^{\perp} := \mathbb{I} - |\perp\rangle \langle \perp|_{\mathbf{D}_x} .$$

Conditioned on the outcome being 0, the state of $\mathbf{D}$ would collapse to one such that if you then perform $H \leftarrow \mathcal{M}^{\text{full}}(\mathbf{D})$, the outcome is so that $H(x)$ is uniform

and independent of the remaining values of H . Equivalently, consider the following quantum procedure that prepares a register $Z := |Y\rangle$ for a uniformly chosen $Y \leftarrow \mathcal{Y}$, and then performs

$$\text{comp}_x \cdot \text{SWAP}_{D_x Z} \cdot \text{comp}_x \quad (1)$$

to D_x and Z , after which throws away the register Z . Then conditioned on the outcome of $\mathcal{M}_x^Z(D)$ being 0, whether or not the above procedure is done subsequently is perfectly indistinguishable.

3 Reprogramming with Advice

In this section we are concerned with the following setting. Let $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle, and $\{\mathcal{D}_i\}_i$ be a family of distributions over $\mathcal{X}$, such that for every i ,

$$\text{guess}(X) := \max_{X \leftarrow \mathcal{D}_i} \Pr_{X \leftarrow \mathcal{D}_i} [X = x] \leq \epsilon \quad (2)$$

for some small but fixed $\epsilon > 0$. An algorithm $\mathcal{A}$ is given as input an H -dependent quantum advice ρ of L qubits, at most q_H quantum queries to H , and in addition tasked to distinguish, given at most q_S classical queries, an oracle O_b between the cases $b = 0$ and $b = 1$, as specified below.

$O_0(i)$:	$O_1(i)$:
1: $X \leftarrow \mathcal{D}_i$	1: $X \leftarrow \mathcal{D}_i$
2: $Y := H(X)$	2: $H(X) := Y \leftarrow \mathcal{Y}$
3: return (X, Y)	3: return (X, Y)

Fig. 3. The Reprogramming Oracles

Remark 1. Our reprogramming oracle is similar to that of [GHHM21, Theorem 1], except that it also outputs Y_i and $\mathcal{D}_i$ provides no side information about X_i . This does not weaken our result: the analysis of [GHHM21, Appendix A] applies equally to Theorem 1. In particular, since our theorem holds for arbitrary query-bounded adversaries, the additional side information can be simulated by a reduction after receiving X_i .

As one of our main technical contributions, we show that O_0 and O_1 are indistinguishable, with a concrete bound detailed as follows.

Theorem 1. *Let $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle, ρ be an H -dependent quantum advice of $L \geq 1$ qubits; let $\{\mathcal{D}_i\}_i$ be a family of distributions such that (2) holds for some $\epsilon > 0$, and O_0 and O_1 be oracles as specified in Fig. 3. Then for every oracle algorithm $\mathcal{A}$ given at most q_H quantum queries to H , and at most q_S classical queries to either O_b , we have*

$$SD(\mathcal{A}^{H, O_0}(\rho), \mathcal{A}^{H, O_1}(\rho)) \leq 12 \sqrt[3]{q_S(q_H + q_S)L\epsilon}.$$

Our proof consists of two parts:

First, we spell out the main proof in Section 3.2. Similar to the proof strategy of [HHM⁺26, Theorem 5], we construct from $\mathcal{A}$ an algorithm $\mathcal{B}$ that makes bounded-many (quantum) queries to H , so that for $b \leftarrow \{0, 1\}$, the advantage $SD(\mathcal{A}^{H, O_0}(\rho), \mathcal{A}^{H, O_1}(\rho))$ is small, whenever

$$SD(\mathcal{A}^{H, O_0}(\sigma_1), \mathcal{A}^{H, O_1}(\sigma_1) \mid (1; \sigma_1) \leftarrow \mathcal{B}^H(\rho))$$

is small. The precise way this step is done significantly differs from [HHM⁺26] which uses *multi-instance game* as in [CGLQ20]. Rather, we follow a technique known as *alternating-measurement game* put forward by [Liu23].¹⁰

Second, in Section 3.3, we control the conditional statistical distance via a dedicated quantum argument. This step proceeds through a hybrid sequence similar to that of [FH23, Appendix A]: we purify the use of random oracle H and perform a (gentle) measurement before every O_b query to check whether the programmed point X is “fresh”, and abort if that is not the case. With this measurement in place, O_0 and O_1 behave identically, completing the analysis. To obtain a tighter control over the disturbance caused by these measurements, we borrow a similar idea from the analyses of [LGXZ26, Theorem 4].

3.1 Some Facts About Alternating Measurements

Below, we state and prove some auxiliary facts that are useful for reasoning about alternating measurements. These facts have also been used in [Liu23], up to differences in presentation.

Lemma 1. *Let ρ be a density operator, and Π_0 and Π_1 be two orthogonal projectors so that $\Pi_0 \rho = \rho$. Then for every $k \in \mathbb{Z}_{>0}$ we have*

$$\text{tr}((\Pi_0 \Pi_1)^k \rho) = \text{tr}((\Pi_0 \Pi_1 \Pi_0)^k \rho) \quad (3)$$

$$(\text{tr}(\Pi_1 \rho))^{2k} \leq \text{tr}((\Pi_0 \Pi_1)^k \rho (\Pi_1 \Pi_0)^k) . \quad (4)$$

Proof of (3). Via the fact that $\rho = \Pi_0 \rho$ and $\Pi_0 \Pi_0 = \Pi_0$, we obtain

$$(\Pi_0 \Pi_1)^k \rho = (\Pi_0 \Pi_0 \Pi_1)^k \Pi_0 \rho = \Pi_0 (\Pi_0 \Pi_1 \Pi_0)^k \rho ,$$

with the last equality coming from re-parenthesization. Taking trace on both sides would then give us

$$\begin{aligned} \text{tr}((\Pi_0 \Pi_1)^k \rho) &= \text{tr}(\Pi_0 (\Pi_0 \Pi_1 \Pi_0)^k \rho) \\ &= \text{tr}((\Pi_0 \Pi_1 \Pi_0)^k \rho \Pi_0) = \text{tr}((\Pi_0 \Pi_1 \Pi_0)^k \rho) , \end{aligned}$$

where the second equality is via cyclicity of trace, and the last equality is via $\rho \Pi_0 = \rho$. This concludes the proof. $\square$

¹⁰ Note that the bit-fixing reduction of [Liu23] does *not* apply, as O_0 and O_1 are accessed classically.

Proof of (4). Given that $\Pi_0 \Pi_1 \Pi_0$ is a semidefinite operator, we can rewrite it as a (non-negative) linear combination of mutually orthogonal projectors $\{E_i\}_i$,

$$\Pi_0 \Pi_1 \Pi_0 = \sum_i \lambda_i E_i \quad (5)$$

such that $\sum_i E_i = \mathbb{I}$, where $\lambda_i \geq 0$ for each i . This would then give us

$$\begin{aligned} (\text{tr}(\Pi_1 \rho))^{2k} &= (\text{tr}(\Pi_1 \Pi_0 \rho \Pi_0))^{2k} = (\text{tr}(\Pi_0 \Pi_1 \Pi_0 \rho))^{2k} \\ &= \left(\sum_i \lambda_i \text{tr}(E_i \rho) \right)^{2k} \leq \sum_i \lambda_i^{2k} \text{tr}(E_i \rho), \end{aligned} \quad (6)$$

where the first equality is via the fact that $\Pi_0 \rho = \rho$, the second equality is via cyclicity of trace, and the last inequality is via Jensen's inequality.

Further observe the following:

$$\begin{aligned} \text{tr}((\Pi_0 \Pi_1)^k \rho (\Pi_1 \Pi_0)^k) &= \text{tr}((\Pi_0 \Pi_1)^k \rho \Pi_0 (\Pi_1 \Pi_0)^k) \\ &= \text{tr}(\Pi_0 (\Pi_1 \Pi_0)^k (\Pi_0 \Pi_1)^k \rho) = \text{tr}((\Pi_0 \Pi_1)^{2k} \rho) = \text{tr}((\Pi_0 \Pi_1 \Pi_0)^{2k} \rho) \\ &= \text{tr} \left(\left(\sum_i \lambda_i E_i \right)^{2k} \rho \right) = \sum_i \lambda_i^{2k} \text{tr}(E_i \rho), \end{aligned} \quad (7)$$

where the first equality is via the fact that $\rho = \rho \Pi_0$, the second equality is via cyclicity of trace, the third equality is via re-parenthesization and $\Pi_0 \Pi_0 = \Pi_0$, the fourth equality is via substituting k in (3) with $2k$, the fifth equality is via (5), and the last equality is via the mutual orthogonality of $\{E_i\}$. Combining (6) and (7), the proof is concluded. $\square$

Lemma 2. Let $\mathcal{M}_0$ and $\mathcal{M}_1$ be two binary projective measurements. Let ρ_0 be a density operator so that $\Pr[\mathcal{M}_0(\rho_0)] = 1$, and let $b_1, \dots, b_{2k}$ be the outcomes of applying $\mathcal{M}_0$ and $\mathcal{M}_1$ to the initial state ρ_0 in the alternating order, starting from $\mathcal{M}_1$, i.e.

$$(b_i; \rho_i) \leftarrow \mathcal{M}_{i \bmod 2}(\rho_{i-1}).$$

Then the conditional probabilities $p_i := \Pr[b_i \mid b_1 \wedge \dots \wedge b_{i-1}]$ are non-decreasing when i increases.

Proof. Let Π_0 and Π_1 denote the orthogonal¹¹ projectors onto the outcome 1 of $\mathcal{M}_0$ and $\mathcal{M}_1$, respectively. Since $\Pr[\mathcal{M}_0(\rho_0)] = 1$, we have $\Pi_0 \rho_0 = \rho_0$. Fix any integer $i \geq 1$. The probability that all the outcomes of the first i measurements

¹¹ By orthogonality we meant $\Pi_b \Pi_b = \Pi_b$, but not necessarily orthogonal to each other.

are 1 can be written as:

$$\begin{aligned}
& \Pr[b_1 \wedge \cdots \wedge b_i] \\
&= \text{tr}(\Pi_{(i \bmod 2)} \cdots \Pi_{(1 \bmod 2)} \rho_0 \Pi_{(1 \bmod 2)} \cdots \Pi_{(i \bmod 2)}) \\
&= \text{tr}(\Pi_{(1 \bmod 2)} \cdots \Pi_{(i \bmod 2)} \Pi_{(i \bmod 2)} \cdots \Pi_{(1 \bmod 2)} \rho_0) \\
&= \text{tr}(\Pi_{(1 \bmod 2)} \cdots \Pi_{(i \bmod 2)} \cdots \Pi_{(1 \bmod 2)} \rho_0 \Pi_0) \\
&= \text{tr}((\Pi_0 \Pi_1)^i \rho_0) = \text{tr}((\Pi_0 \Pi_1 \Pi_0)^i \rho_0), \tag{8}
\end{aligned}$$

where the second equality is via the cyclicity of trace, the third equality is via $\Pi_b \Pi_b = \Pi_b$ for each $b \in \{0, 1\}$ and $\rho_0 = \rho_0 \Pi_0$, the fourth equality is again via cyclicity of trace, and the last equality is via substituting k in (3) with i . Since Π_0 and Π_1 are orthogonal projectors, $\Pi_0 \Pi_1 \Pi_0$ is semidefinite, in which case the power $(\Pi_0 \Pi_1 \Pi_0)^\alpha$ for each $\alpha \geq 0$ is a well-defined semidefinite operator. Likewise, $\sqrt{\rho_0}$ is also well-defined. We thus have:

$$\begin{aligned}
\text{tr}((\Pi_0 \Pi_1 \Pi_0)^i \rho_0)^2 &= \text{tr}\left(\left((\Pi_0 \Pi_1 \Pi_0)^{(i-1)/2} \sqrt{\rho_0}\right)^\dagger \left((\Pi_0 \Pi_1 \Pi_0)^{(i+1)/2} \sqrt{\rho_0}\right)\right)^2 \\
&\leq \text{tr}((\Pi_0 \Pi_1 \Pi_0)^{i-1} \rho_0) \text{tr}((\Pi_0 \Pi_1 \Pi_0)^{i+1} \rho_0), \tag{9}
\end{aligned}$$

where the equality is via cyclicity of trace, in addition to $\Pi_0 \Pi_1 \Pi_0$ and ρ_0 being semidefinite, and the (non-equality) inequality follows from Cauchy-Schwarz for the Hilbert-Schmidt inner product $\langle X, Y \rangle := \text{tr}(X^\dagger Y)$. Combining (8) and (9) we obtain:

$$\begin{aligned}
p_i &= \frac{\Pr[b_1 \wedge \cdots \wedge b_i]}{\Pr[b_1 \wedge \cdots \wedge b_{i-1}]} = \frac{\text{tr}((\Pi_0 \Pi_1 \Pi_0)^i \rho_0)}{\text{tr}((\Pi_0 \Pi_1 \Pi_0)^{i-1} \rho_0)} \\
&\leq \frac{\text{tr}((\Pi_0 \Pi_1 \Pi_0)^{i+1} \rho_0)}{\text{tr}((\Pi_0 \Pi_1 \Pi_0)^i \rho_0)} = \frac{\Pr[b_1 \wedge \cdots \wedge b_{i+1}]}{\Pr[b_1 \wedge \cdots \wedge b_i]} = p_{i+1}.
\end{aligned}$$

Hence, the conditional probabilities p_i are non-decreasing in i , which concludes the proof. $\square$

3.2 Proof of Theorem 1

Theorem 1. *Let $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle, ρ be an H -dependent quantum advice of $L \geq 1$ qubits; let $\{\mathcal{D}_i\}_i$ be a family of distributions such that (2) holds for some $\epsilon > 0$, and O_0 and O_1 be oracles as specified in Fig. 3. Then for every oracle algorithm $\mathcal{A}$ given at most q_H quantum queries to H , and at most q_S classical queries to either O_b , we have*

$$SD(\mathcal{A}^{H, O_0}(\rho), \mathcal{A}^{H, O_1}(\rho)) \leq 12 \sqrt[3]{q_S(q_H + q_S)} L \epsilon.$$

Proof. For every fixed choice h of H , let $\mathcal{M}_\mathcal{A}^h$ be a binary projective measurement that operates on the register $\mathbf{z}$ storing the advice, and an auxiliary working register $\mathbf{w}$, so that

$$\Pr[\mathcal{M}_\mathcal{A}^h(\rho_{\mathbf{z}} \otimes |0\rangle\langle 0|_{\mathbf{w}})] = \Pr_{b \leftarrow \{0,1\}}[\mathcal{A}^{h, O_b}(\rho) = b].$$

In more detail, this measurement is done via purifying the randomness and b used in running $\mathcal{A}^{H,O_b}$, running $\mathcal{A}$ coherently and measuring whether the output matches b , and finally uncomputing the working register. It is therefore not too hard to see that the measurement can be implemented via making at most $2(q_H + q_S)$ quantum queries to H , since each query to H and to O_b by $\mathcal{A}$ would result in two queries to H in the measurement, respectively.

Now, we specify the algorithm $\mathcal{B}$, which runs $(y; \sigma_y) \leftarrow \mathcal{G}_{2k}^H(\mathbb{I}/2^L)$ with the input being a maximally mixed state over L qubits, and produces $(y; \tau_y := \text{tr}_w \sigma_y)$, where w is the internal working register of $\mathcal{G}_{2k}$, and k is a parameter to be determined later. On input $z := \rho$, the algorithm $\mathcal{G}_\ell$ (including and beyond the case when $\ell = 2k$) is defined as follows: initialize the working register to $w := |0\rangle\langle 0|$. Then, apply two binary projective measurements $\mathcal{M}_{\mathcal{A}}^H$ and $\mathcal{M}_{\text{reset}}$ in an alternating order (starting from $\mathcal{M}_{\mathcal{A}}^H$) for at most ℓ times, and output 1 if all measurement outcomes are 1. Here, the 1-outcome of $\mathcal{M}_{\text{reset}} = \{\Pi_{\text{reset}}, \mathbb{I} - \Pi_{\text{reset}}\}$ corresponds to $\Pi_{\text{reset}} := |0\rangle\langle 0|_w$, and we likewise write $\Pi_{\mathcal{A}}^H$ for the projector corresponding to the 1-outcome of $\mathcal{M}_{\mathcal{A}}^H$.

Without loss of generality, assume $\Pr_b[\mathcal{A}^{H,O_b} = b] \geq 1/2$, so that

$$SD(\mathcal{A}^{H,O_0}(\rho), \mathcal{A}^{H,O_1}(\rho)) = 2 \Pr_{b \leftarrow \{0,1\}}[\mathcal{A}^{H,O_b}(\rho) = b] - 1. \quad (10)$$

Toward bounding the above, observe that for every fixed choice h of H , and for every (h -dependent) quantum state ρ^h over L qubits, we have

$$\begin{aligned} \Pr_{b \leftarrow \{0,1\}}[\mathcal{A}^{h,O_b}(\rho^h) = b]^{2k} &= \Pr[\mathcal{M}_{\mathcal{A}}^h(\rho^h \otimes |0\rangle\langle 0|)]^{2k} = \text{tr}(\Pi_{\mathcal{A}}^h(\rho^h \otimes |0\rangle\langle 0|))^{2k} \\ &\leq \text{tr}((\Pi_{\text{reset}} \Pi_{\mathcal{A}}^h)^k(\rho^h \otimes |0\rangle\langle 0|)(\Pi_{\mathcal{A}}^h \Pi_{\text{reset}})^k) = \Pr[\mathcal{G}_{2k}^h(\rho^h)], \end{aligned}$$

where the only non-equality inequality is via (4) in Lemma 1. Averaging over $h \sim H$ on both sides with $\rho^h := \rho|_{H=h}$, Jensen's inequality gives

$$\Pr_{b \leftarrow \{0,1\}}[\mathcal{A}^{H,O_b}(\rho) = b]^{2k} \leq \Pr[\mathcal{G}_{2k}^H(\rho)] \leq 2^L \cdot \Pr[\mathcal{G}_{2k}^H(\mathbb{I}/2^L)]. \quad (11)$$

Let b_i be the outcome of the i th measurement in $\mathcal{G}_{2k}^H(\mathbb{I}/2^L)$, the above probability can be rewritten into

$$\Pr[\mathcal{G}_{2k}^H(\mathbb{I}/2^L)] = \Pr[b_1 \wedge \cdots \wedge b_{2k}] = \prod_{i \in [2k]} \Pr[b_i \mid b_1 \wedge \cdots \wedge b_{i-1}].$$

Exploiting the fact that, by Lemma 2, the terms in the product are non-decreasing in i , the above is further bounded by

$$\begin{aligned}
\Pr [\mathcal{G}_{2k}^H(\mathbb{I}/2^L)] &\leq \left(\Pr [b_{2k} \mid b_1 \wedge \dots \wedge b_{2k-1}] \right)^{2k} \\
&\leq \left(\Pr [b_{2k+1} \mid b_1 \wedge \dots \wedge b_{2k}] \right)^{2k} \\
&= \Pr [\mathcal{M}_{\mathcal{A}}^H(\sigma_1) \mid (1, \sigma_1) \leftarrow \mathcal{G}_{2k}^H(\mathbb{I}/2^L)]^{2k} \\
&= \Pr_{b \leftarrow \{0,1\}} [\mathcal{A}^{H,O_b}(\tau_1) = b \mid (1, \tau_1) \leftarrow \mathcal{B}^H]^{2k} \\
&\leq \left(\frac{1}{2} + \frac{1}{2} SD(\mathcal{A}^{H,O_0}(\sigma_1), \mathcal{A}^{H,O_1}(\sigma_1) \mid (1; \sigma_1) \leftarrow \mathcal{B}^H) \right)^{2k}. \quad (12)
\end{aligned}$$

Combining the obtained bounds (10), (11), and (12), we obtain

$$\begin{aligned}
&SD(\mathcal{A}^{H,O_0}(\rho), \mathcal{A}^{H,O_1}(\rho)) \\
&\leq 2^{L/(2k)} (1 + SD(\mathcal{A}^{H,O_0}(\sigma_1), \mathcal{A}^{H,O_1}(\sigma_1) \mid (1; \sigma_1) \leftarrow \mathcal{B}^H)) - 1 \\
&\leq 2^{L/(2k)} \left(1 + q_S^2 \epsilon + 2\sqrt{(2k+1)q_S q_H \epsilon} + 2kq_S^2 \epsilon \right) - 1. \quad (13)
\end{aligned}$$

where the last inequality is via applying Lemma 3 with the fact that $\mathcal{B}$ makes up to $2k(q_H + q_S)$ quantum queries to H , and then simplifying the obtained bound. As long as $L \leq 2k$, we have $2^{L/(2k)} \leq 1 + L/(2k)$, which further simplifies (13) to the following form¹²

$$SD(\mathcal{A}^{H,O_0}(\rho), \mathcal{A}^{H,O_1}(\rho)) \leq \left(1 + \frac{L}{2k} \right) \left(1 + q_S^2 \epsilon + 2\sqrt{(2k+1)q_S q_H \epsilon} + 2kq_S^2 \epsilon \right) - 1.$$

Plugging in $k := \left\lceil \sqrt[3]{\frac{L^2}{q_S(q_H + q_S)\epsilon}} \right\rceil$ to the above, which (roughly) equates the two parentheses, we get the following bound for $\alpha := q_S(q_H + q_S)L\epsilon$

$$SD(\mathcal{A}^{H,O_0}(\rho), \mathcal{A}^{H,O_1}(\rho)) \leq \left(1 + (1/2)\alpha^{1/3} \right) \left(1 + \alpha + 2\sqrt{2}\alpha^{1/3} + 2\sqrt{3}\alpha^{1/2} \right) - 1$$

From the observation that, whenever $\alpha \leq 1$, the above is at most $12\alpha^{1/3}$, and that otherwise the upperbound exceeds one and is trivially satisfied, the proof is concluded. $\square$

3.3 A Two-Stage Reprogramming Game

In the remainder, we prove the following statement. This is an analogue of [HHM⁺26, Lemma 6] in our setting, with a crucial difference that we have to deal with, and provide a tighter bound for, multiple reprogramming points.

Lemma 3. *Let $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle, and $\{\mathcal{D}_i\}_i$ be a family of distributions such that (2) holds for some $\epsilon > 0$, and O_0 and O_1 be oracles*

¹² When $L > 2k$, this bound trivially holds anyway.

as specified in Fig. 3. Then, for every oracle algorithm $\mathcal{A}$ making at most q_H quantum queries to H , at most q_S classical queries to O_b , and for every oracle algorithm $\mathcal{B}$ making at most q_B quantum queries to H , we have

$$SD\left(\mathcal{A}^{H,O_0}(\sigma_1), \mathcal{A}^{H,O_1}(\sigma_1) \mid (1; \sigma_1) \leftarrow \mathcal{B}^H\right) \leq q_S^2 \epsilon + 2\sqrt{q_S(q_H + q_B)\epsilon}.$$

Proof. In the following, we assume that the reader is familiar with the compressed oracle technique introduced in [Zha19]. Without loss of generality, assume that $\mathcal{A}$ makes exactly q_H queries to H and q_S queries to O_b , padding with dummy queries if necessary.

Rewriting the game as Hyb_0^b . Similar to the proof of Lemma 6 in [HHM⁺26], we consider the following game: repeatedly sample a random oracle H and run $\mathcal{B}^H$ until the measurement outcome is $b_B = 1$, obtaining the post-measurement state σ_B ; then run $\mathcal{A}^{H,O_b}$ on input σ_B .

Hyb₀^b:
 1: **repeat**
 2: let H be an RO
 3: $(b_B; \sigma_B) \leftarrow \mathcal{B}^H$
 4: **until** $b_B = 1$
 5: $b_A \leftarrow \mathcal{A}^{H,O_b}(\sigma_B)$
 6: **return** b_A

Fig. 4. A rewrite of the two-stage reprogramming game.

By construction, the output distribution of Hyb_0^b is identical to that of $\mathcal{A}^{H,O_b}(\sigma_1)$ conditioned on $(1, \sigma_1) \leftarrow \mathcal{B}^H$. Hence,

$$SD\left(\mathcal{A}^{H,O_0}(\sigma_1), \mathcal{A}^{H,O_1}(\sigma_1) \mid (1, \sigma_1) \leftarrow \mathcal{B}^H\right) = SD\left(\text{Hyb}_0^0, \text{Hyb}_0^1\right), \quad (14)$$

From Hyb_0^b to Hyb_1^b . Next, we consider game Hyb_1^b , which differs from Hyb_0^b in two respects. First, queries to H are simulated using a compressed oracle cO (see Section 2.6), whose database register $D = (D_x)_x$ is initialized to the empty database $|\perp\rangle$. Second, for each $i \in [q_S]$, the i th query to O_b is simulated as follows:

- O_0 : After $X_i \leftarrow D_i$ is sampled, prepare a response register in state $|0\rangle$, apply one (compressed-oracle) query on the classical input X_i , and measure only the response register in the computational basis. If the measurement outcome is Y_i , return (X_i, Y_i) .
- O_1 : First, sample $X_i \leftarrow D_i$ and $Y_i \leftarrow \mathcal{Y}$, and initialize an auxiliary register Z_i to $|Y_i\rangle$. Then, decompress the oracle at X_i , swap Z_i with the sub-register D_{X_i} which contains the restored value $H(X_i)$, decompress the oracle at X_i again, and finally return (X_i, Y_i) . Similar to (1), we denote the resulting operation by

$$R_{X_i Y_i} := \text{Comp}_{X_i} \cdot \text{SWAP}_{D_{X_i} Z_i} \cdot \text{Comp}_{X_i}.$$

Since all auxiliary registers introduced in the simulation of O_b are inaccessible to $\mathcal{A}$, the above construction provides a perfect simulation of Hyb_0^b . Therefore,

$$SD(\text{Hyb}_0^b, \text{Hyb}_1^b) = 0. \quad (15)$$

From Hyb_1^b to $\text{Hyb}_{\text{col}}^b$. Next, define the game $\text{Hyb}_{\text{col}}^b$, which differs from Hyb_1^b as follows: during each i th query to O_b , right after sampling $X_i \leftarrow \mathcal{D}_i$, classically check whether $X_i \in \{X_1, \dots, X_{i-1}\}$ and abort if so. Thus, conditioned on any non-aborting execution history prior to the i th query to O_b , game $\text{Hyb}_{\text{col}}^b$ guarantees that the previously sampled values $X_1, \dots, X_{i-1}$ are pairwise distinct. Since $X_i \leftarrow \mathcal{D}_i$ is freshly sampled, (2) implies

$$\Pr[X_i \in \{X_1, \dots, X_{i-1}\}] \leq (i-1)\epsilon.$$

Recall that $\mathcal{A}$ makes at most q_S queries to O_b , we have

$$SD(\text{Hyb}_1^b, \text{Hyb}_{\text{col}}^b) \leq \Pr\left[\bigcup_{i=1}^{q_S} X_i \in \{X_1, \dots, X_{i-1}\}\right] \leq \sum_{i=1}^{q_S} (i-1)\epsilon \leq \frac{q_S^2 \epsilon}{2}. \quad (16)$$

From $\text{Hyb}_{\text{col}}^b$ to $\text{Hyb}_{\text{hit}}^b$. Next, we consider $\text{Hyb}_{\text{hit}}^b$, which differs from $\text{Hyb}_{\text{col}}^b$ by an additional measurement during each execution of O_b : In the i th query to O_b , right after the classical test in $\text{Hyb}_{\text{col}}^b$, $\text{Hyb}_{\text{hit}}^b$ performs an additional binary projective measurement $\mathcal{M}_i^{\text{hit}} := \Pi_i^0, \Pi_i^1$ on the sub-register D_{X_i} , where

$$\Pi_i^0 := |\perp\rangle\langle\perp|_{D_{X_i}}, \quad \Pi_i^1 := I - \Pi_i^0.$$

This measurement tests whether X_i was already recorded in the compressed database prior to its sampling during the i th query to O_b .

Consider the hypothetical experiment, in which we run $\text{Hyb}_{\text{col}}^b$ except that in the i th O_b query, we perform the measurement $\mathcal{M}_i^{\text{hit}}$ right after the collision test is performed. Let b_i^{hit} be the output of such a measurement, and p_i^{hit} be the probability that it equals 1. Toward controlling the probability p_i^{hit} , we introduce (yet another) hypothetical measurement $D \leftarrow \mathcal{M}^{\text{hyp}}$ (right before $\mathcal{M}_i^{\text{hit}}$) that measures the database register in computational basis, i.e. with the projector associated to the outcome database $D : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ being $|D\rangle\langle D|$. This measurement is not a gentle measurement, but it does not change the probability p_i^{hit} . If the i th measurement ($\mathcal{M}_i^{\text{hit}}$) is reached, then none of the preceding i collision tests has caused the game to abort, in which case $X_j \notin \{X_1, \dots, X_{j-1}\}$ for all $j \in [i]$. Hence, even though the previous $i-1$ queries to O_b may have introduced non- $\perp$ records at positions $X_1, \dots, X_{i-1}$ in the database D , none of these positions would have caused $b_i^{\text{hit}} = 1$. Formally, we then have

$$\begin{aligned} p_i^{\text{hit}} &= \Pr[D(X_i) \neq \perp \wedge X_i \notin \{X_1, \dots, X_{i-1}\}] \\ &\leq \mathbb{E}_D \left[|\{x \in \mathcal{X} \mid D(x) \neq \perp\} \setminus \{X_1, \dots, X_{i-1}\}| \cdot \epsilon \right] \leq (q_B + q_H) \epsilon, \end{aligned}$$

where the first (non-equality) inequality is via X_i having $\log(1/\epsilon)$ bits of min-entropy, and the last inequality is via the fact that $\{x \in \mathcal{X} \mid D(x) \neq \perp\} \setminus \{X_1, \dots, X_{i-1}\}$ is of size at most $q_B + q_H$ with certainty.

We next bound the disturbance caused by $\mathcal{M}_1^{\text{hit}}, \dots, \mathcal{M}_{q_S}^{\text{hit}}$. Instead of applying the *gentle measurement lemma* [Win99] separately to each measurement, which would incur a loss of $\sum_i \sqrt{p_i^{\text{hit}}}$, similar to the case of [LGXZ26], we analyze the entire measurement sequence jointly. In particular, the *Gentle Sequential Measurement lemma* [OV22] bounds the distance between $\text{Hyb}_{\text{col}}^b$ and $\text{Hyb}_{\text{hit}}^b$ conditioned on $b^{\text{hit}} := b_1^{\text{hit}} \vee \dots \vee b_{q_S}^{\text{hit}} = 0$ in the latter game:

$$SD\left(\text{Hyb}_{\text{col}}^b, \text{Hyb}_{\text{hit}}^b \mid b^{\text{hit}}=0\right) \leq \sqrt{\sum_{i=1}^{q_S} p_i^{\text{hit}}} \leq \sqrt{q_S (q_B + q_H)} \epsilon. \quad (17)$$

Wrapping up. Finally, conditioned on non-abort (namely $X_1, \dots, X_{q_S}$ being distinct) and $b^{\text{hit}} = 0$, each reprogramming point X_i has not been previously recorded in the database. Namely, measuring out the database D at the i th O_b query would have implied that $D(X_i) = \perp$. Hence, in both $\text{Hyb}_{\text{hit}}^0$ and $\text{Hyb}_{\text{hit}}^1$, whether the reprogramming is done does not have any noticeable difference, and so the two games behave identically (see Section 2.6). Therefore, we obtain

$$SD\left(\text{Hyb}_{\text{hit}}^0, \text{Hyb}_{\text{hit}}^1 \mid b^{\text{hit}} = 0\right) = 0. \quad (18)$$

Summarizing the bounds so far, we have

$$\begin{aligned} \mathcal{A}^{H, O_0}(\sigma_1) \big|_{(1, \sigma_1) \leftarrow \mathcal{B}^H} &\stackrel{0}{\approx} \text{Hyb}_0^0 \stackrel{0}{\approx} \text{Hyb}_1^0 \stackrel{\frac{q_S^2 \epsilon}{2}}{\approx} \text{Hyb}_{\text{col}}^0 \stackrel{\sqrt{q_S(q_B + q_H)} \epsilon}{\approx} \text{Hyb}_{\text{hit}}^0 \big|_{b^{\text{hit}}=0} \\ &\stackrel{0}{\approx} 0 \\ \mathcal{A}^{H, O_1}(\sigma_1) \big|_{(1, \sigma_1) \leftarrow \mathcal{B}^H} &\stackrel{0}{\approx} \text{Hyb}_0^1 \stackrel{0}{\approx} \text{Hyb}_1^1 \stackrel{\frac{q_S^2 \epsilon}{2}}{\approx} \text{Hyb}_{\text{col}}^1 \stackrel{\sqrt{q_S(q_B + q_H)} \epsilon}{\approx} \text{Hyb}_{\text{hit}}^1 \big|_{b^{\text{hit}}=0} \end{aligned}$$

This concludes the proof. $\square$

3.4 On (Almost) Tightness

In prior subsections, our reprogramming theorem (Theorem 1) provides a bound of order $O(\sqrt[3]{q_S q_H L \epsilon})$, assuming as is typical that $q_H \geq q_S$. In this (sub)section, we show that the reprogramming games considered in our theorem can be distinguished with advantage $\Omega(\sqrt{q_S q_H L \epsilon})$. In other words, there exists an attacker that matches our lowerbound in the resources q_S, q_H and L taken to reach constant advantage, with $q_S q_H L = \Theta(1/\epsilon)$.

Remark 2. The cubic-root versus square-root gap suggests that there is still room to improve either (or both) our reprogramming theorem or our attack. Though, it is unclear which exact case we are at.

Theorem 2. *There exists a distribution $\mathcal{D}_1$ satisfying (2) for some $\epsilon > 0$, an attacker $\mathcal{A}$ making q_H and q_S classical queries to a random oracle H , and the reprogramming oracle O_b (as in Fig. 3 with respect to the singleton family $\{\mathcal{D}_1\}$) respectively, and an H -dependent classical advice z of L bits, such that*

$$SD(\mathcal{A}^{H, O_0}(z), \mathcal{A}^{H, O_1}(z)) \geq 0.0803 \cdot \min(1, \sqrt{q_S q_H L \epsilon}) .$$

Proof. In this proof, we write $\lceil x \rceil_{\text{odd}}$ for the smallest odd integer upperbound of x . Let $K := \lceil 1/(q_S q_H L \epsilon) \rceil_{\text{odd}}$, and let $\{S_{j,k}\}_{j,k}$ be a family of size- q_H , disjoint subsets of $\{0, 1\}^*$ indexed by $(j, k) \in [L] \times [K]$, and let $\tilde{S} \subseteq \{0, 1\}^*$ be a size- $\lceil 1/\epsilon \rceil$ superset or subset of $\bigcup_{j,k} S_{j,k}$. In addition, let $\text{Maj} : \{0, 1\}^* \rightarrow \{0, 1\}$ be the function that outputs 1 if there are more 1's in the input, and 0 otherwise.

We define the distribution $\mathcal{D}_1$ to be the uniform distribution over the set $\tilde{S}$. We then define the advice z and the attacker $\mathcal{A}$ as follows.

Definition of z :

```

1: for  $j \in [L]$ 
2:   for  $k \in [K]$ 
3:      $b_{j,k} := \bigoplus_{x \in S_{j,k}} H(x)$ 
4:    $B_j := \text{Maj}(b_{j,1}, \dots, b_{j,K})$ 
5: return  $(B_1, \dots, B_L)$ 
```

$\mathcal{A}^{H, O_b}(z = (B_1, \dots, B_L))$:

```

1: for  $i \in [q_S]$ 
2:    $(X_i, Y_i) \leftarrow O_b(1)$ 
3: for  $(i, j, k) \in [q_S] \times [L] \times [K]$ 
4:   if  $X_i \in S_{j,k}$ 
5:     return  $B_j = \bigoplus_{x \in S_{j,k}} H(x)$ 
6: abort
```

Fig. 5. An attack against the reprogramming game.

First, we observe that the probability $\mathcal{A}$ aborts is at most

$$\begin{aligned} \Pr[\mathcal{A}^{H, O_b} \text{ abort}] &\leq \Pr[X_i \notin S_{j,k} \forall i, j, k] \leq \left(1 - q_H L K / \lceil 1/\epsilon \rceil\right)^{q_S} \\ &\leq \left(1 - q_H L K \epsilon / 2\right)^{q_S} \leq \left(1 - 1/(2q_S)\right)^{q_S} \leq 1/\sqrt{e} , \end{aligned}$$

whenever $q_H L K \epsilon \leq 1$ (so that $\tilde{S} \supseteq \bigcup_{j,k} S_{j,k}$), where in the third inequality we assume $0 < \epsilon \leq 1$ and so $1/\lceil 1/\epsilon \rceil \geq \epsilon/2$, the fourth inequality is via $K = \lceil 1/(q_S q_H L \epsilon) \rceil_{\text{odd}}$, and the last inequality is via $(1 - 1/(2x))^x \leq 1/\sqrt{e}$ for all $x \geq 1$. On the other hand, when $q_H L K \epsilon > 1$, the set $\tilde{S}$ is a subset of $\bigcup_{j,k} S_{j,k}$, so the abort probability is zero and the final bound on $1/\sqrt{e}$ still holds.

Conditioned on $\mathcal{A}$ not aborting, we work out the explicit probability that $\mathcal{A}$ outputs 1. Starting with $b = 1$, we have

$$\Pr[1 \leftarrow \mathcal{A}^{H, O_1} \mid \text{no abort}] = \frac{1}{2},$$

since B_j is independent of the parity bit $\bigoplus_{x \in S_{j,k}} H(x)$ computed *after* the reprogramming. For the non-abort case of $b = 0$, the attacker $\mathcal{A}$ would enter line 5 for some $X_i \in S_{j,k}$, in which case the probability of outputting 1 is that of the majority bit $B_j = \text{Maj}(b_{j,1}, \dots, b_{j,K})$ coinciding with $b_{j,k}$. Formally, let $(\mathbf{i}, \mathbf{j}, \mathbf{k})$ be the first tuple of indices (in the order of the for-loop) such that $X_i \in S_{j,k}$.

Then

$$\Pr [1 \leftarrow \mathcal{A}^{H, O_0} \mid \text{no abort}] = \Pr [B_{\mathbf{j}} = b_{\mathbf{j}, \mathbf{k}} \mid \text{no abort}] \geq \frac{1}{2} + \frac{1}{2\sqrt{2K}},$$

where the last inequality follows from Lemma 4 and the fact that, given $\mathbf{j}$, the bit string $b_{\mathbf{j}, 1}, \dots, b_{\mathbf{j}, K}$ is just a uniform random string independent of $\mathbf{k}$.

Collecting the bounds so far, we have that

$$\begin{aligned} SD(\mathcal{A}^{H, O_0}, \mathcal{A}^{H, O_1}) &\geq \Pr[\text{no abort}] \cdot SD(\mathcal{A}^{H, O_0}, \mathcal{A}^{H, O_1} \mid \text{no abort}) \\ &\geq (1 - 1/\sqrt{e}) \cdot \frac{1}{2\sqrt{2K}}. \end{aligned}$$

Observe that, when $q_{sq_H} L\epsilon \leq 1$, we have $K = \lceil 1/q_{sq_H} L\epsilon \rceil_{\text{odd}} \leq 3/q_{sq_H} L\epsilon$, in which case the above is at least

$$\geq \frac{(1 - 1/\sqrt{e})}{2\sqrt{6}} \cdot \sqrt{q_{sq_H} L\epsilon} \geq 0.0803 \cdot \min(1, \sqrt{q_{sq_H} L\epsilon}),$$

while for $q_{sq_H} L\epsilon > 1$, we can simply plug in $K = 1$ so the bound holds trivially. This concludes the proof. $\square$

Lemma 4. *Let K be an odd integer and $\mathbf{k} \in [K]$. Let Maj be the majority function as defined in the proof above, $b_1, \dots, b_K \leftarrow \{0, 1\}$ be uniformly random bits that are sampled independently. Then, we have*

$$\Pr[\text{Maj}(b_1, \dots, b_K) = b_{\mathbf{k}}] \geq \frac{1}{2} + \frac{1}{2\sqrt{2K}}.$$

Proof. Since K is odd, write $K = 2r + 1$ for some integer $r \geq 0$. If $r = 0$, the inequality holds trivially. Otherwise, let $M := \sum_{i \in [K] \setminus \mathbf{k}} b_i$ denote the number of 1 bits among the remaining $2r$ bits, then we have $M \sim \text{Bin}(2r, 1/2)$. If $M = r$, then the remaining bits are equally split between 0 and 1, and hence $b_{\mathbf{k}}$ determines the majority. Therefore,

$$\Pr[\text{Maj}(b_1, \dots, b_K) = b_{\mathbf{k}} \mid M = r] = 1.$$

If $M \neq r$, then among the remaining $2r$ bits, the numbers of 0's and 1's differ by at least 2. Hence, the majority bit among these $2r$ bits cannot be changed by taking the single bit $b_{\mathbf{k}}$ into account. Therefore, $\text{Maj}(b_1, \dots, b_K)$ is independent of the uniformly random bit $b_{\mathbf{k}}$, and

$$\Pr[\text{Maj}(b_1, \dots, b_K) = b_{\mathbf{k}} \mid M \neq r] = 1/2.$$

Consequently,

$$\begin{aligned} \Pr[\text{Maj}(b_1, \dots, b_K) = b_{\mathbf{k}}] &= \Pr[M = r] \cdot 1 + \Pr[M \neq r] \cdot \frac{1}{2} \\ &= \frac{1}{2} + \frac{1}{2} \Pr[M = r] = \frac{1}{2} + \frac{1}{2} \frac{\binom{2r}{r}}{2^{2r}}. \end{aligned}$$

Let $d_r := \binom{2r}{r} / 2^{2r}$. Since $r \geq 1$,

$$\frac{d_{r+1}}{d_r} = \frac{2r+1}{2r+2} \geq \sqrt{\frac{r}{r+1}},$$

Together with $d_1 = 1/2$, this implies by induction that $d_r \geq 1/(2\sqrt{r})$. Therefore,

$$\Pr[\text{Maj}(b_1, \dots, b_K) = b_K] = \frac{1}{2} + \frac{1}{2}d_r \geq \frac{1}{2} + \frac{1}{2\sqrt{2K}}.$$

□

4 Pseudorandomness from Random Oracle

In this section, we provide two immediate applications of our reprogramming theorem with advice. In Section 4.2, we are going to provide a tighter PRG security bound for random oracle, and in Section 4.3 we provide IND-CPA security proof of salted pseudo one-time pad, both in the presence of quantum advice.

Before stating and proving these two results though, we present an indistinguishability bound in Section 4.1 that will be useful later. This bound is widely believed true, and may be considered a folklore in the community. However, to the best of our knowledge, no explicit statements nor proofs can be found in the literature. We therefore provide a self-contained proof from scratch, which may be of independent interest.

4.1 A Folklore Indistinguishability Bound

For any set S , we define the membership oracle as follows

$$\text{MEMBER}_S(k) := \begin{cases} 1 & \forall k \in S \\ 0 & \forall k \notin S. \end{cases}$$

What we are interested in is the distinguishing advantage between $\text{MEMBER}_\emptyset$ and $\text{MEMBER}_{\{K\}}$ for a hidden key $K \leftarrow \mathcal{D}$ of high min-entropy, namely,

$$\text{guess}(K) = \max_{K \leftarrow \mathcal{D}} \Pr[K = k] \leq \epsilon, \quad (19)$$

for some small $\epsilon > 0$.

Lemma 5. *Let $K \leftarrow \mathcal{D}$ satisfy (19) for some $\epsilon > 0$, and let $\mathcal{A}$ be an oracle algorithm making at most q quantum queries to $\text{MEMBER}_\emptyset$ or $\text{MEMBER}_{\{K\}}$. Then, we have*

$$SD(\mathcal{A}^{\text{MEMBER}_\emptyset}, \mathcal{A}^{\text{MEMBER}_{\{K\}}}) \leq 8q^2\epsilon.$$

To begin with, we analyze the special case where $\mathcal{D}$ is a uniform distribution over a finite set $\mathcal{K}$. Then, we will reduce the general case to the uniform case.

Proof for uniform $\mathcal{D}$. For every sequence of binary bits $\mathbf{y} := (y_k)_{k \in \mathcal{K}}$, let $O_{\mathbf{y}}(k)$ be the oracle that returns y_k for every k , and let

$$f(\mathbf{y}) := \Pr [1 \leftarrow \mathcal{A}^{O_{\mathbf{y}}}] .$$

Standard polynomial method (e.g. [BBC⁺98, Lemma 4.2]) shows that $f(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]$ is a multilinear polynomial of total degree $\deg(f) \leq 2q$ in the bits of $\mathbf{y}$. Consider now the univariate polynomial $g : \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$g(x) := \sum_{\mathbf{y} \in \{0,1\}^{\mathcal{K}}} f(\mathbf{y}) x^{\|\mathbf{y}\|_1} (1-x)^{|\mathcal{K}| - \|\mathbf{y}\|_1}$$

with $\|\mathbf{y}\|_1$ denoting the Hamming weight of $\mathbf{y}$. Note that, for $0 \leq x \leq 1$, we have $g(x) = \mathbb{E}[f(\mathbf{y})]$, where the expectation is taken over $\mathbf{y}$, with each bit sampled from the Bernoulli distribution $\text{Ber}(x)$ independently. Therefore, it follows that $\deg(g) \leq \deg(f) \leq 2q$. Moreover, since

$$\begin{aligned} g'(0) &= \sum_{\mathbf{y} \text{ s.t. } \|\mathbf{y}\|_1=1} f(\mathbf{y}) - |\mathcal{K}| f(0, \dots, 0) \\ &= |\mathcal{K}| \Pr [1 \leftarrow \mathcal{A}^{\text{MEMBER}_{\{\mathcal{K}\}}}] - |\mathcal{K}| \Pr [1 \leftarrow \mathcal{A}^{\text{MEMBER}_{\emptyset}}] , \end{aligned} \quad (20)$$

it suffices to bound $|g'(0)|$, which we do via Markov brothers' inequality as follows

$$|g'(0)| \leq 2(\deg g)^2 \max_{0 \leq x \leq 1} |g(x)| \leq 8q^2 . \quad (21)$$

Combining (20) and (21) and rearranging terms yields the desired bound. $\square$

Proof for general $\mathcal{D}$. Without loss of generality, we assume that $\mathcal{D}$ is over a finite support $\mathcal{K}$. Otherwise, the desired bound would still be achieved via discarding every point with a zero probability mass, and then applying the same bound to a sequence of finite-support distributions that converges to $\mathcal{D}$. Since the uniform $\mathcal{D}$ case has been addressed, here we assume that $\mathcal{D}$ is non-uniform. We also assume that $\mathcal{K}$ is non-singleton, since otherwise the desired bound also holds trivially.

For notational convenience, we write as a short-hand notation the following:

$$\Delta(\mathcal{D}) := \mathbb{E}_{K \leftarrow \mathcal{D}} (\mathcal{A}^{\text{MEMBER}_{\emptyset}}, \mathcal{A}^{\text{MEMBER}_{\{\mathcal{K}\}}}) .$$

Let $k^* \in \mathcal{K}$ be the element with smallest (non-zero) probability mass, i.e.

$$\gamma := \Pr [K = k^*] = \min_{k \in \mathcal{K}} \Pr [K = k] .$$

We can rewrite $\mathcal{D}$ as a convex combination of the uniform distribution over $\mathcal{K}$, which we denote by $\mathcal{U}_{\mathcal{K}}$, and a distribution $\mathcal{D}'$ over the punctured set $\mathcal{K} \setminus \{k^*\}$. Specifically, we have $\mathcal{D} = \gamma|\mathcal{K}|\mathcal{U}_{\mathcal{K}} + (1 - \gamma|\mathcal{K}|)\mathcal{D}'$ where $\mathcal{D}'$ is uniquely defined

by $\mathcal{D}'(k) := \frac{\mathcal{D}(k) - \gamma}{1 - \gamma|\mathcal{K}|}$. Therefore, we have that

$$\begin{aligned}
\Delta(\mathcal{D}) &\leq \gamma|\mathcal{K}| \cdot \Delta(\mathcal{U}_{\mathcal{K}}) + (1 - \gamma|\mathcal{K}|) \cdot \Delta(\mathcal{D}') \\
&\leq 8q^2\gamma + (1 - \gamma|\mathcal{K}|) \cdot \Delta(\mathcal{D}') \\
&\leq 8q^2\gamma + (1 - \gamma|\mathcal{K}|) \cdot 8q^2 \underset{K' \leftarrow \mathcal{D}'}{\text{guess}}(K') \\
&\leq 8q^2\gamma + (1 - \gamma|\mathcal{K}|) \cdot 8q^2 \frac{\epsilon - \gamma}{1 - \gamma|\mathcal{K}|} = 8q^2\epsilon,
\end{aligned}$$

where the second inequality is via the desired bound for the uniform distribution, and the third inequality is via assuming the induction hypothesis that the desired bound holds for all distributions over $\mathcal{K} \setminus \{k^*\}$. This concludes the proof. $\square$

4.2 Multi-challenge PRG Security

Theorem 3. *Let $H : \mathcal{K} \rightarrow \mathcal{Y}$ be a random oracle, ρ be an H -dependent quantum advice of $L \geq 1$ qubits, and $K_1, \dots, K_n$ be independent random variables sampled according to some distribution over $\mathcal{K}$, so that each satisfies $\text{guess}(K_i) \leq \epsilon$ for some $\epsilon > 0$. Then, for every oracle algorithm $\mathcal{A}$ making at most q quantum queries to H , and for $Y_1, \dots, Y_n \leftarrow \mathcal{Y}$ sampled uniformly and independently from $\mathcal{Y}$, we have*

$$SD\left(\mathcal{A}^H(\rho, H(K_1), \dots, H(K_n)), \mathcal{A}^H(\rho, Y_1, \dots, Y_n)\right) \leq 12\sqrt[3]{n(q+n)L\epsilon} + 32nq^2\epsilon.$$

Proof. For notational convenience, we write H_i for the oracle H reprogrammed at $K_1, \dots, K_i$ to output $Y_1, \dots, Y_i$ respectively, in order, with the convention that H_0 is the original oracle H .

Our proof proceeds through the following hybrid sequence

$$\mathcal{A}^H(\rho, H(K_1), \dots, H(K_n)) \approx \mathcal{A}^{H_n}(\rho, Y_1, \dots, Y_n) \approx \mathcal{A}^H(\rho, Y_1, \dots, Y_n).$$

The first step is via a direct application of our reprogramming theorem (Theorem 1), substituting q_S and q_H with n and q , giving us

$$SD\left(\mathcal{A}^H(\rho, H(K_1), \dots, H(K_n)), \mathcal{A}^{H_n}(\rho, Y_1, \dots, Y_n)\right) \leq 12\sqrt[3]{n(q+n)L\epsilon}, \quad (22)$$

The second step is via another hybrid (sub)sequence

$$\mathcal{A}^{H_n}(\rho, Y_1, \dots, Y_n) \approx \dots \approx \mathcal{A}^{H_0}(\rho, Y_1, \dots, Y_n),$$

where, for each transition from H_i to H_{i-1} , we construct a reduction $\mathcal{B}_i$ (Fig. 6) that distinguishes $\text{MEMBER}_{\emptyset}$ from $\text{MEMBER}_{\{K_i\}}$, as defined in Lemma 5.

In more detail, $\mathcal{B}_i$ would internally sample the random oracle H , the advice ρ , the inputs $Y_1, \dots, Y_n$, the keys $K_1, \dots, K_{i-1}$ (hence H_{i-1}), and run $\mathcal{A}$ under the hood. Upon each of $\mathcal{A}$'s random-oracle queries on input k , the reduction $\mathcal{B}_i$ would query MEMBER_S (for $S \in \{\emptyset, \{K_i\}\}$), and respond Y_i if $\text{MEMBER}_S(k)$ returns 1 and $H_{i-1}(k)$ otherwise. We illustrate the reduction $\mathcal{B}_i$ in the pseudocode below,

noting on the side that the sub-procedure $\mathcal{O}$ can be implemented coherently, thereby supporting quantum superposition queries made by $\mathcal{A}$.

$\mathcal{B}_i^{\text{MEMBER}_S}$: 1: sample $H, \rho, Y_1, \dots, Y_n$ 2: sample $K_1, \dots, K_{i-1}$ hence H_{i-1} 3: return $\mathcal{A}^{\mathcal{O}}(\rho, Y_1, \dots, Y_n)$	$\mathcal{O}(k)$: 1: if $\text{MEMBER}_S(k)$ then 2: return Y_i 3: else return $H_{i-1}(k)$
---	---

Fig. 6. A reduction $\mathcal{B}_i$ that distinguishes the membership oracles.

Observe that $\mathcal{B}_i$ can implement $\mathcal{O}$ via querying MEMBER_S twice per $\mathcal{A}$'s (quantum) query, summing up to $2q$ queries to MEMBER_S in total. Furthermore, given that $\mathcal{B}_i^{\text{MEMBER}_{\{\kappa_i\}}}$ and $\mathcal{B}_i^{\text{MEMBER}_{\emptyset}}$ perfectly simulate $\mathcal{A}^{H_i}$ and $\mathcal{A}^{H_{i-1}}$ respectively (on input $(\rho, Y_1, \dots, Y_n)$), we have

$$\begin{aligned}
& SD(\mathcal{A}^{H_n}(\rho, Y_1, \dots, Y_n), \mathcal{A}^H(\rho, Y_1, \dots, Y_n)) \\
& \leq \sum_{i \in [n]} SD(\mathcal{A}^{H_i}(\rho, Y_1, \dots, Y_n), \mathcal{A}^{H_{i-1}}(\rho, Y_1, \dots, Y_n)) \\
& \leq \sum_{i \in [n]} SD\left(\mathcal{B}_i^{\text{MEMBER}_{\{\kappa_i\}}}, \mathcal{B}_i^{\text{MEMBER}_{\emptyset}}\right) \leq 32nq^2\epsilon, \tag{23}
\end{aligned}$$

where the last inequality is via substituting q in Lemma 5 with $2q$. Combining (22) and (23), the proof is concluded. $\square$

4.3 Security of Salted Pseudo One-Time Pad

In this (sub)section, we consider post-quantum security of the natural salted version of pseudo one-time pad, viewed as a symmetric encryption scheme with secret key $K \in \mathcal{K}$ sampled from a high-min-entropy distribution $\mathcal{D}$, i.e.,

$$\max_k \Pr_{K \leftarrow \mathcal{D}}[K = k] \leq \epsilon, \tag{24}$$

for some small $\epsilon > 0$. We recall the exact construction in Fig. 7 below, which relies on a hash function H that we model as a random oracle.

KGen 1: $K \leftarrow \mathcal{D}$ 2: return K	Enc (K, m) 1: $R \leftarrow \{0, 1\}^\ell$ 2: $c := H(K, R) \oplus m$ 3: return $C := (R, c)$	Dec (K, C): 1: $(R, c) := C$ 2: $m := c \oplus H(K, R)$ 3: return m
--	--	--

Fig. 7. The salted pseudo one-time pad $\text{\$-pseudoOTP} = (\text{KGen}, \text{Enc}, \text{Dec})$ defined with respect to a hash function $H : \{0, 1\}^* \rightarrow \mathcal{M}$, where the message space $\mathcal{M}$ consists of bit strings of a certain fixed length.

As an application of our reprogramming theorem, we show that $\$$ -pseudoOTP satisfies the standard IND-CPA security, even when the attacker is given an arbitrary quantum advice ρ dependent on the random oracle H (see Section 2.5).

Theorem 4. *Let H be a random oracle, ρ be an H -dependent quantum advice of length $L \geq 1$ qubits, and $\$$ -pseudoOTP be the salted pseudo one-time pad scheme of salt length ℓ that satisfies (24) for some $\epsilon > 0$. Then, for every IND-CPA attacker $\mathcal{A}$ making at most q_H quantum queries to the random oracle H , and at most q_E classical queries to the challenge oracle, we have the following two incomparable bounds:*

$$\text{Adv}_{\$-\text{pseudoOTP}}^{\text{IND-CPA}}(\mathcal{A}, \rho) \leq 24 \sqrt[3]{q_E(q_H + q_E)L/2^\ell} + 64q_H^2\epsilon, \quad (25)$$

$$\text{Adv}_{\$-\text{pseudoOTP}}^{\text{IND-CPA}}(\mathcal{A}, \rho) \leq 24 \sqrt[3]{(q_H + 1)L\epsilon} + 64q_H^2\epsilon + q_E^2/2^{\ell-1}. \quad (26)$$

Proof of (25). Our proof proceeds through the following hybrid sequence

$$\mathcal{A}^{H, \text{Hyb}_0^b}(\rho) \approx \mathcal{A}^{H, \text{Hyb}_1^b}(\rho) \approx \mathcal{A}^{H, \text{Hyb}_2^b}(\rho),$$

where Hyb_i are as specified below. We note that the starting game $\mathcal{A}^{H, \text{Hyb}_0^b}$ is just a rewrite of the IND-CPA_b game, and hence it suffices to show that the $b = 0$ case is indistinguishable to the $b = 1$ case.

$\text{Hyb}_0^b(m_0, m_1)$:	$\text{Hyb}_1^b(m_0, m_1)$:	$\text{Hyb}_2^b(m_0, m_1)$:
1: $R \leftarrow \{0, 1\}^\ell$	1: $R \leftarrow \{0, 1\}^\ell$	1: $R \leftarrow \{0, 1\}^\ell$
2: $Y := H(K, R)$	2: $H(K, R) := Y \leftarrow \mathcal{M}$	2: $Y \leftarrow \mathcal{M}$
3: $c := Y \oplus m_b$	3: $c := Y \oplus m_b$	3: $c := Y \oplus m_b$
4: return $C := (R, c)$	4: return $C := (R, c)$	4: return $C := (R, c)$

Fig. 8. The hybrid oracles for the proof of Eq. (25).

From Hyb_0^b to Hyb_1^b . The only difference between Hyb_0^b and Hyb_1^b is that the former evaluates $Y := H(K, R)$, while the latter reprograms the same point with $H(K, R) := Y \leftarrow \mathcal{M}$. Since R is drawn freshly with ℓ bits of min-entropy, Theorem 1 would directly give us

$$SD(\mathcal{A}^{H, \text{Hyb}_0^b}, \mathcal{A}^{H, \text{Hyb}_1^b}) \leq 12 \sqrt[3]{q_E(q_H + q_E)L/2^\ell}.$$

From Hyb_1^b to Hyb_2^b . The only difference between Hyb_1^b and Hyb_2^b is that the former reprograms $H(K, R) := Y \leftarrow \mathcal{M}$, while the latter samples $Y \leftarrow \mathcal{M}$ without reprogramming it to H . Then, via the same reasoning as in the proof of Theorem 3, there exists a reduction algorithm $\mathcal{B}$ distinguishing MEMBER_S for $S = \emptyset$ or $S = \{K\}$, such that

$$SD(\mathcal{A}^{H, \text{Hyb}_1^b}(\rho), \mathcal{A}^{H, \text{Hyb}_2^b}(\rho)) \leq SD(\mathcal{B}^{H, \text{MEMBER}_{\{K\}}}, \mathcal{B}^{H, \text{MEMBER}_{\emptyset}}) \leq 32q_H^2\epsilon,$$

where the last inequality follows from the fact that $\mathcal{B}$ makes at most $2q_H$ quantum queries to MEMBER_S .

Wrapping up. We note that the oracles Hyb_2^b behave identically independent of b , since the only difference is when masking the message via $c := Y \oplus m_b$ for a freshly sampled Y . Therefore, with $SD\left(\mathcal{A}^{H, \text{Hyb}_0^b}(\rho), \mathcal{A}^{H, \text{Hyb}_0^{-b}}(\rho)\right)$ bounded by

$$\leq SD\left(\mathcal{A}^{H, \text{Hyb}_2^b}(\rho), \mathcal{A}^{H, \text{Hyb}_2^{-b}}(\rho)\right) + \sum_{b,i} SD\left(\mathcal{A}^{H, \text{Hyb}_{i-1}^b}(\rho), \mathcal{A}^{H, \text{Hyb}_i^b}(\rho)\right) \quad (27)$$

where summation goes over $(b, i) \in \{0, 1\} \times [2]$, we conclude the proof via plugging in the obtained bounds so far. $\square$

Proof of (26). Similarly, our proof of the second bound proceeds through the following hybrid sequence

$$\mathcal{A}^{H, \text{Hyb}_0^b}(\rho) \approx \mathcal{A}^{H, \text{Hyb}_1^b}(\rho) \approx \mathcal{A}^{H, \text{Hyb}_2^b}(\rho),$$

where Hyb_i are as specified below, with F being a random function that is independent of H . Again, the starting game $\mathcal{A}^{H, \text{Hyb}_0^b}$ is identical to IND-CPA_b game, and hence it suffices to show that the $b = 0$ case is indistinguishable to the $b = 1$ case.

$\text{Hyb}_0^b(m_0, m_1):$	$\text{Hyb}_1^b(m_0, m_1):$	$\text{Hyb}_2^b(m_0, m_1):$
1: $R \leftarrow \{0, 1\}^\ell$	1: $R \leftarrow \{0, 1\}^\ell$	1: $R \leftarrow \{0, 1\}^\ell$
2: $Y := H(K, R)$	2: $Y := F(R)$	2: $Y \leftarrow \mathcal{M}$
3: $c := Y \oplus m_b$	3: $c := Y \oplus m_b$	3: $c := Y \oplus m_b$
4: return $C := (R, c)$	4: return $C := (R, c)$	4: return $C := (R, c)$

Fig. 9. The hybrid oracles for the proof of Eq. (26).

From Hyb_0^b to Hyb_1^b . The only difference between Hyb_0^b and Hyb_1^b is that the former evaluates $Y := H(K, R)$ through the use of secret key K , while the latter evaluates another random function F that is independent of H . This follows from the fact that $H(K, \cdot)$ is a quantum-secure PRF, a direct consequence of $G(K) := (H(K, m))_m$ being a quantum-secure PRG. Note that we have shown the latter in Theorem 3 (as a special case when $n = 1$), and the obtained concrete bound is as follows:

$$SD\left(\mathcal{A}^{H, \text{Hyb}_0^b}, \mathcal{A}^{H, \text{Hyb}_1^b}\right) \leq 12 \sqrt[3]{(q_H + 1)L\epsilon} + 32q_H^2\epsilon.$$

From Hyb_1^b to Hyb_2^b . The only difference between Hyb_1^b and Hyb_2^b is that the former evaluates $Y := F(R)$ using a random function, while the latter samples a fresh $Y \leftarrow \mathcal{M}$. Therefore, conditioned on R being distinct in all of the q_E queries, which almost always happens except with probability at most $q_E^2/2^\ell$, the two oracles behave identically. This gives us

$$SD\left(\mathcal{A}^{H, \text{Hyb}_1^b}, \mathcal{A}^{H, \text{Hyb}_2^b}\right) \leq q_E^2/2^\ell.$$

Wrapping up. Given that the definition of Hyb_0^b and Hyb_2^b in this proof is the same as that of (25), the inequality (27) still applies. Plugging in the bounds that we have obtained so far, we conclude the proof. $\square$

5 Fiat-Shamir and Hash-and-Sign Security with Advice

In this section, we analyze signature schemes based on Fiat-Shamir transform and hash-and-sign design principle, when attackers are given quantum advice.

5.1 A Generalized Fiat-Shamir Signature Scheme

Similar to [HHM⁺26], we treat Fiat-Shamir and hash-and-sign signature schemes via a unified framework that captures both.¹³ Recall that such signature schemes are built up from a Sigma protocol, or a trapdoor function. We capture them in terms of a scheme called public-key Sigma protocol $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$. As its name suggests, it is a three-round, public-coin, interactive proof system, in which the prover $\mathcal{P}$ tries to convince the verifier $\mathcal{V}$ that he knows the secret-key sk corresponding to a public-key pk . We illustrate in Fig. 10 how Σ works when executed honestly. With a slight abuse of notation, we write $(a, y, z) \leftarrow \Sigma(sk)$ for the transcript (a, y, z) generated honestly given the secret-key sk .




Fig. 10. The public-key Sigma protocol $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$, when both the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ behave honestly.

Naturally, we would like Σ to be hard to win if a malicious prover $\mathcal{A}$ only has pk but not sk . We formalize this via the *identification* (ID) advantage

$$\text{Adv}_{\Sigma}^{\text{ID}}(\mathcal{A}, \rho) := \Pr \left[1 \leftarrow \text{Ver}(pk, a, y, z) \mid \begin{array}{l} (sk, pk) \leftarrow \text{KGen} \\ (a; st) \leftarrow \mathcal{A}(\rho, pk) \\ z \leftarrow \mathcal{A}(st, y \leftarrow \mathcal{Y}) \end{array} \right], \quad (28)$$

¹³ Note that [HHM⁺26] covers a more generalized class of signature schemes that involves aborting, which we do not cover here.

where we further allow $\mathcal{A}$ to be given a quantum advice ρ . Informally, we say that Σ is identification-secure if the advantage is small for every efficient $\mathcal{A}$ and every ρ that is bounded in (qubit) length.

Below we define the (generalized) Fiat-Shamir signature scheme $\text{FS}^H[\Sigma]$ using a cryptographic hash function H , which is then modeled as a random oracle.

KGen' :	$\text{Sign}'(sk, m)$:	$\text{Ver}'(pk, m, \sigma)$:
1: $(sk, pk) \leftarrow \text{KGen}$	1: $(a, st) \leftarrow \text{Com}(sk)$	1: $(a, z) := \sigma$
2: return (sk, pk)	2: $y := H(a, m)$	2: $y := H(a, m)$
	3: $z \leftarrow \text{Resp}(st, y)$	3: return $\text{Ver}(pk, a, y, z)$
	4: return $\sigma := (a, z)$	

Fig. 11. The (generalized) Fiat-Shamir signature scheme $\text{FS}^H[\Sigma] = (\text{KGen}', \text{Sign}', \text{Ver}')$ for a public-key Sigma protocol $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$ and hash function H .

Several other security properties of Σ are of relevance here. The high min-entropy property (Definition 1) and the HVZK property (Definition 2) are standard when providing a security reduction from UF-CMA to UF-NMA, which we will discuss in Section 5.2.

Definition 1. A public-key Sigma protocol $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$ has $\log(1/\epsilon)$ bits of min-entropy, if

$$\text{guess}(a \mid sk, pk) := \max_f \Pr[a = f(sk, pk)] \leq \epsilon$$

for $(sk, pk) \leftarrow \text{KGen}$ and $(a, st) \leftarrow \text{Com}(sk)$.

Definition 2. A public-key Sigma protocol $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$ is (ζ, T) -honest-verifier-zero-knowledge (or (ζ, T) -HVZK for short), if there is a simulator Sim_Σ of running time T , such that

$$\mathbb{E}_{(sk, pk) \leftarrow \text{KGen}} \left[\left((sk, pk, \text{Sim}_\Sigma(pk)), (sk, pk, \Sigma(sk)) \right) \right] \leq \zeta.$$

In addition, since we are interested in the setting in which an attacker is given an advice as input, we introduce an additional nontriviality condition (Definition 3). As we will see in Section 5.3, this condition is necessary for UF-NMA security in the presence of advice. More surprisingly, it also turns out to be sufficient, thereby providing a full characterization of security in this setting.

Definition 3. Let $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$ be a public-key Sigma-protocol. We define the non-triviality advantage of an algorithm $\mathcal{A}$ given a quantum advice ρ against Σ by

$$\text{Adv}_\Sigma^{\neg \text{Triv}}(\mathcal{A}, \rho) := \max_{a^\circ, y^\circ} \Pr_{(sk, pk) \leftarrow \text{KGen}} [\text{Ver}(pk, a^\circ, y^\circ, \mathcal{A}(\rho, pk, a^\circ, y^\circ))] .$$

Informally, we say that Σ is nontrivial, if for every efficient algorithm $\mathcal{A}$ and every ρ bounded in (qubit) length, the advantage is small.

5.2 Achieving UF-CMA

Theorem 5. *Let H be a random oracle, ρ be H -dependent quantum advice of $L \geq 1$ qubits, and Σ be a public-key Sigma-protocol with $\log(1/\epsilon)$ bits of min-entropy and (ζ, T) -HVZK. Let $\text{FS}^H[\Sigma]$ be its Fiat-Shamir signature scheme obtained from Σ using H . Then, for every UF-CMA attacker $\mathcal{A}$ making at most q_H quantum queries to H and q_S classical signing queries, there exists a UF-NMA attacker $\mathcal{B}$ making at most q_H quantum queries to H such that*

$$\text{Adv}_{\text{FS}^H[\Sigma]}^{\text{UF-CMA}}(\mathcal{A}, \rho) \leq \text{Adv}_{\text{FS}^H[\Sigma]}^{\text{UF-NMA}}(\mathcal{B}, \rho) + 12\sqrt[3]{q_S(q_H + q_S + 1)L\epsilon} + q_S\zeta.$$

Moreover, the $\text{TIME}(\mathcal{B})$ is on the order of $\text{TIME}(\mathcal{A}) + O(q_S T + (q_S + q_H) \log q_S)$ and $\text{TIME}(\mathcal{A}) + O(q_S T + q_H q_S + q_S^2)$ with and without (quantum-read, classical-write) QRAM respectively, when we count basic arithmetics as unit time.

The proof of Theorem 5 follows the standard UF-CMA-to-UF-NMA reduction for Fiat-Shamir signature schemes without advice, while the additional difficulty arising from oracle-dependent advice is handled by Theorem 1. The proof first reprograms the random oracle at the points associated with signing queries, allowing the corresponding signing transcripts to be replaced by simulated ones using the HVZK property. Once all signing queries have been simulated in this way, the signing oracle can be removed, yielding an adversary for the UF-NMA experiment.

Proof. Our proof proceeds through the following hybrid sequence

$$\text{UF-CMA} \approx \text{Hyb}_1 \approx \text{Hyb}_2 \approx \text{Hyb}_3,$$

where we recall the UF-CMA game and specify the Hyb_i games below. The oracles used in these games will be further specified in Fig. 13.

UF-CMA (ρ):	Hyb₁ (ρ):
1: $H_0 := H$	1: $H_0 := H$
2: $(sk, pk) \leftarrow \text{KGen}$	2: $(sk, pk) \leftarrow \text{KGen}$
3: $(m, \sigma) \leftarrow \mathcal{A}^{H, \text{Sign}(sk, \cdot)}(\rho, pk)$	3: $(m, \sigma) \leftarrow \mathcal{A}^{H, \text{Trans}(sk, \cdot)}(\rho, pk)$
4: $\{m_i := \text{ith signing query}\}$	4: $\{m_i := \text{ith signing query}\}$
5: return $\text{Ver}^H(pk, m, \sigma) \wedge m \notin \{m_i\}_i$	5: return $\text{Ver}^H(pk, m, \sigma) \wedge m \notin \{m_i\}_i$
Hyb₂ (ρ):	Hyb₃ (ρ):
1: $H_0 := H$	1: $H_0 := H$
2: $(sk, pk) \leftarrow \text{KGen}$	2: $(sk, pk) \leftarrow \text{KGen}$
3: $(m, \sigma) \leftarrow \mathcal{A}^{H, \text{Sim}(pk, \cdot)}(\rho, pk)$	3: $(m, \sigma) \leftarrow \mathcal{A}^{H, \text{Sim}(pk, \cdot)}(\rho, pk)$
4: $\{m_i := \text{ith signing query}\}$	4: $\{m_i := \text{ith signing query}\}$
5: return $\text{Ver}^H(pk, m, \sigma) \wedge m \notin \{m_i\}_i$	5: return $\text{Ver}^{H_0}(pk, m, \sigma) \wedge m \notin \{m_i\}_i$

Fig. 12. The hybrid games for the UF-CMA to UF-NMA reduction.

Note that, in each game except for Hyb_3 , we grayed out the (irrelevant) first line $H_0 := H$ that denotes the initial state of the random oracle H before any reprogramming takes place.

$\text{Sign}(sk, m)$:	$\text{Trans}(sk, m)$:	$\text{Sim}(pk, m)$:
1: $(a, st) \leftarrow \text{Com}(sk)$	1: $(a, st) \leftarrow \text{Com}(sk)$	1: $(a, y, z) \leftarrow \text{Sim}_\Sigma(pk)$
2: $y := H(a, m)$	2: $H(a, m) := y \leftarrow \mathcal{Y}$	2: $H(a, m) := y$
3: $z \leftarrow \text{Resp}(st, y)$	3: $z \leftarrow \text{Resp}(st, y)$	3:
4: return (a, z)	4: return (a, z)	4: return (a, z)

Fig. 13. The oracles in Sign , Trans , and Sim . The first hop changes how the challenge is obtained; the second hop changes how the transcript is generated.

From UF-CMA to Hyb_1 . The only difference between UF-CMA and Hyb_1 is that we replace the oracle Sign with Trans , which replaces the hash evaluation $y := H(a, m)$ with $H(a, m) := y \leftarrow \mathcal{Y}$. Since a has high min-entropy the two should be indistinguishable. Formally, denoting the event $(sk, pk) = (sk^\circ, pk^\circ)$ by E_{sk°, pk° for every fixed (sk°, pk°) , and substituting q_H in Theorem 1 with $q_H + 1$ (where the extra query comes from verification), we obtain

$$SD(\text{UF-CMA}(\rho), \text{Hyb}_1(\rho) \mid E_{sk^\circ, pk^\circ}) \leq 12 \sqrt[3]{q_S(q_H + q_S + 1)L \frac{\text{guess}}{a \leftarrow \text{Com}(sk^\circ)}(a)}.$$

Averaging both sides over $(sk^\circ, pk^\circ) \sim (sk, pk)$, Jensen's inequality yields

$$\begin{aligned} SD(\text{UF-CMA}(\rho), \text{Hyb}_1(\rho)) &\leq 12 \sqrt[3]{q_S(q_H + q_S + 1)L \frac{\text{guess}}{(sk, pk) \leftarrow \text{KGen}, a \leftarrow \text{Com}(sk)}(a \mid sk, pk)} \\ &\leq 12 \sqrt[3]{q_S(q_H + q_S + 1)L\epsilon}, \end{aligned}$$

where the second inequality follows from the min-entropy property of Σ .

From Hyb_1 to Hyb_2 . The only difference between Hyb_1 and Hyb_2 is that we replace the oracle Trans with Sim , which replaces the transcript (a, y, z) generated by Σ with the HVZK simulator $\text{Sim}_\Sigma(pk)$. By the (ζ, T) -HVZK property of Σ , and taking into account q_S signing queries in total, we have

$$SD(\text{Hyb}_1(\rho), \text{Hyb}_2(\rho)) \leq q_S \zeta.$$

From Hyb_2 to Hyb_3 . The only difference between Hyb_2 and Hyb_3 is that we replace the random oracle H queried by the verifier with the oracle H_0 that is not reprogrammed. By construction of our simulator Sim , on input a message m , it only programs entries of the form $H(a, m)$, meaning that a fresh forgery $m \notin \{m_i\}_i$ that is not among prior signing queries must satisfy $H_0(a, m) = H(a, m)$ for every choice of a . Therefore, this step is just a syntactic change, with $SD(\text{Hyb}_2(\rho), \text{Hyb}_3(\rho)) = 0$.

Wrapping up. Assume without loss of generality that $\mathcal{A}$ never produces a message queried to the signing oracle before. In this case, observe that for a UF-NMA reduction algorithm $\mathcal{B}^H(\rho)$ that internally simulates $\mathcal{A}^{H, \text{Sim}}(\rho)$ and locally book-keeps the reprogrammed points, its advantage is identical to the winning probability of Hyb_3 , namely,

$$\text{Adv}_{\text{FS}^H[\Sigma]}^{\text{UF-NMA}}(\mathcal{B}, \rho) = \Pr[1 \leftarrow \text{Hyb}_3] .$$

Collecting the bounds so far, we have

$$\text{Adv}_{\text{FS}^H[\Sigma]}^{\text{UF-CMA}}(\mathcal{A}, \rho) \leq \text{Adv}_{\text{FS}^H[\Sigma]}^{\text{UF-NMA}}(\mathcal{B}, \rho) + 12\sqrt[3]{q_S(q_H + q_S + 1)L\epsilon} + q_S\zeta .$$

Finally, we note that the reduction algorithm $\mathcal{B}$ is standard and the same as that of [BBD⁺23, FFH26], so it inherits the same running time as has been summarized in [FFH26, Table 1]. This concludes the proof. $\square$

5.3 An UF-NMA Characterization

Theorem 6. *Let H be a random oracle, ρ be an H -dependent quantum advice of $L \geq 1$ qubits, and Σ be a public-key Sigma-protocol. Let $\text{FS}^H[\Sigma]$ be the Fiat-Shamir signature scheme of Σ using H . Then, for every UF-NMA attacker $\mathcal{A}$ given at most q_H quantum queries to H , there are an attacker $\mathcal{B}$ against identification security of Σ , an attacker $\mathcal{C}$ against the nontriviality condition of Σ , and two H -dependent advices $\rho_{\mathcal{B}}$ and $\rho_{\mathcal{C}}$, such that*

$$\text{Adv}_{\text{FS}[\Sigma]}^{\text{UF-NMA}}(\mathcal{A}, \rho) \leq 4(2q_H + 1)^2 \text{Adv}_{\Sigma}^{\text{ID}}(\mathcal{B}, \rho_{\mathcal{B}}) + 4L(q_H + 1) \text{Adv}_{\Sigma}^{\neg \text{Trvl}}(\mathcal{C}, \rho_{\mathcal{C}}) .$$

Moreover, both the running time $\text{TIME}(\mathcal{B})$ and $\text{TIME}(\mathcal{C})$ are at most $\text{TIME}(\mathcal{A}) + O(Lq_H^2)$, and $\rho_{\mathcal{B}}$ and $\rho_{\mathcal{C}}$ are of length at most $O(Lq_H(\log|\mathcal{X}| + \log|\mathcal{Y}|))$ qubits.

Proof. Without loss of generality, assume that $\mathcal{A}$ makes exactly q_H queries to H , padding with dummy queries if necessary.

Reducing UF-NMA(ρ) to the P-Bit-Fixing QROM. We first apply the generic reduction of [Liu23, Theorem 6.1] from the non-uniform UF-NMA(ρ) game to the P-Bit-Fixing-QROM model. In its notation, $S = L$, $T = q_H$, $T_{\text{Samp}} = 0$, and $T_{\text{Verify}} = 1$ since KGen is independent of H , while the final verification queries $H(a, m)$ once. Consequently, the P-Bit-Fixing-QROM parameter is

$$P = S(T + T_{\text{Samp}} + T_{\text{Verify}}) = L(q_H + 1) .$$

We then define the corresponding bit-fixing UF-NMA game as follows. An offline oracle algorithm $\mathcal{D}$ making at most P quantum queries is fixed before the game starts. In the rejection-sampling stage, a random oracle H is sampled uniformly at random and $(b; \sigma) \leftarrow \mathcal{D}^H$ is executed. If $b \neq 1$, a fresh random oracle H is sampled and the procedure is restarted. Once $b = 1$, the resulting state σ is passed to the online stage. The challenger samples $(sk, pk) \leftarrow \text{KGen}$, and the online adversary runs $(m, (a, z)) \leftarrow \mathcal{A}^H(\sigma, pk)$, and wins iff $\text{Ver}(pk, a, H(a, m), z) = 1$.

We refer to the resulting game as the bit-fixing UF-NMA security game. Let $\mathbf{Adv}_{\text{FS}[\Sigma]}^{\text{BF-UF-NMA}}(P, q_H)$ denote the maximum winning probability over all offline algorithms $\mathcal{D}$ making at most P quantum queries and online adversaries making at most q_H quantum queries. Then, by [Liu23, Theorem 6.1], we have

$$\begin{aligned} \mathbf{Adv}_{\text{FS}[\Sigma]}^{\text{UF-NMA}}(\rho) &\leq 2\mathbf{Adv}_{\text{FS}[\Sigma]}^{\text{BF-UF-NMA}}(L(q_H + 1), q_H) \\ &= 2\Pr[1 \leftarrow \text{UF-NMA}(\sigma) \mid (1; \sigma) \leftarrow \mathcal{D}^H]. \end{aligned} \quad (29)$$

Hence, it remains to bound the winning probability of an arbitrary bit-fixing UF-NMA adversary satisfying the above query bounds.

G_x :	$\text{Hyb}_x^{\text{yes/no}}$:
1: repeat	1: repeat
2: $D := \perp\rangle$	2: $D := \perp\rangle$
3: $(b; \sigma) \leftarrow \mathcal{D}^{\text{cO}}$	3: $(b; \sigma) \leftarrow \mathcal{D}^{\text{cO}}$
4: until $b = 1$	4: until $b = 1$
5:	5: abort unless $\text{yes/no} \leftarrow \mathcal{M}_x^{\text{f}}(D)$
6: $(sk, pk) \leftarrow \text{KGen}$	6: $(sk, pk) \leftarrow \text{KGen}$
7: $(m, (a, z)) \leftarrow \mathcal{A}^{\text{cO}}(\sigma, pk)$	7: $(m, (a, z)) \leftarrow \mathcal{A}^{\text{cO}}(\sigma, pk)$
8: $y \leftarrow \mathcal{M}_x^{\text{eval}}(D)$	8: $y \leftarrow \mathcal{M}_x^{\text{eval}}(D)$
9: return $\text{Ver}(pk, a, y, z) \wedge x = (a, m)$	9: return $\text{Ver}(pk, a, y, z) \wedge x = (a, m)$

Fig. 14. The fixed-point game G_x and its yes/no branches.

From bit-fixing UF-NMA to G_x . From now on, we keep $\mathcal{D}$ fixed throughout the proof. For each $x = (a^\circ, m^\circ) \in \mathcal{X}$, let G_x be the bit-fixing UF-NMA experiment in Fig. 14, with the additional winning condition $(a, m) = x$. We simulate H via the compressed oracle cO with database register D , where $\mathcal{M}_x^{\text{eval}}(D)$ queries cO at the classical input x and measures the response in the computational basis. Since $(a, m) = x$ upon winning, this evaluates precisely $H(a, m)$. Consequently, decomposing the winning probability over x gives

$$\Pr[1 \leftarrow \text{UF-NMA}(\sigma) \mid (1; \sigma) \leftarrow \mathcal{D}^H] = \sum_{x \in \mathcal{X}} \Pr[1 \leftarrow G_x]. \quad (30)$$

From G_x to its yes/no branches. After $\mathcal{D}$ outputs $b = 1$ and before running KGen , we test whether x has been queried using the projective measurement $\mathcal{M}_x^{\text{f}} := \{\Pi_x^{\text{no}}, \Pi_x^{\text{yes}}\}$ where $\Pi_x^{\text{no}} := |\perp\rangle\langle\perp|_{D(x)}$, $\Pi_x^{\text{yes}} := I - \Pi_x^{\text{no}}$. For each $s \in \{\text{yes}, \text{no}\}$, Hyb_x^s aborts unless $\mathcal{M}_x^{\text{f}}$ outputs s , and otherwise proceeds as G_x ; see Figure 14. Let $|\Psi\rangle$ be the state immediately before $\mathcal{M}_x$, and let F_x denote the remaining computation of G_x . Since $\Pi_x^{\text{yes}} + \Pi_x^{\text{no}} = I$, we have

$$\begin{aligned} \sqrt{\Pr[1 \leftarrow G_x]} &= \|F_x |\Psi\rangle\| \leq \|F_x \Pi_x^{\text{yes}} |\Psi\rangle\| + \|F_x \Pi_x^{\text{no}} |\Psi\rangle\| \\ &= \sqrt{\Pr[1 \leftarrow \text{Hyb}_x^{\text{yes}}]} + \sqrt{\Pr[1 \leftarrow \text{Hyb}_x^{\text{no}}]}. \end{aligned} \quad (31)$$

Squaring (31), summing over x , and applying the Cauchy-Schwarz inequality to the resulting cross term gives

$$\begin{aligned} \sum_x \Pr[1 \leftarrow G_x] &\leq \left(\sqrt{\sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{yes}}]} + \sqrt{\sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{no}}]} \right)^2 \\ &\leq 2 \sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{yes}}] + 2 \sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{no}}], \end{aligned} \quad (32)$$

where the last inequality follows from $(a + b)^2 \leq 2a^2 + 2b^2$.




Fig. 15. The no-branch hybrid games.

The no branch from Hyb_x^{no} to $\text{Hyb}_{x,1}^{\text{no}}$. The game $\text{Hyb}_{x,1}^{\text{no}}$ is similar to Hyb_x^{no} except that, after the measurement $\mathcal{M}_x$ outputs **no** and before KGen generates sk, pk , we apply $\mathcal{M}^{\text{full}}$ to obtain the oracle H to be used later. From this point on, all oracle queries of $\mathcal{A}$ are answered directly using H , and the final value y is computed as $y := H(x)$. This is a perfect simulation, and this change preserves the winning probability, which gives

$$\Pr[1 \leftarrow \text{Hyb}_x^{\text{no}}] = \Pr[1 \leftarrow \text{Hyb}_{x,1}^{\text{no}}]. \quad (33)$$

The no branch from $\text{Hyb}_{x,1}^{\text{no}}$ to $\text{Hyb}_{x,2}^{\text{no}}$. Next, in game $\text{Hyb}_{x,2}^{\text{no}}$, we introduce two modifications to $\text{Hyb}_{x,1}^{\text{no}}$. First, in line 7, we additionally reprogram $H(x) := y \leftarrow \mathcal{Y}$ with a fresh uniform randomness y . Since in $\text{Hyb}_{x,1}^{\text{no}}$ game we abort unless $\mathcal{M}_x^\mathcal{Y}$ produces **no**, the state of $\mathbf{D}$ already collapses to one such that the reprogramming causes no difference (see Section 2.6). Second, in line 5, we remove the abort condition entirely, which only increases (or preserves) the winning probability of the game. Hence

$$\Pr[1 \leftarrow \text{Hyb}_{x,1}^{\text{no}}] \leq \Pr[1 \leftarrow \text{Hyb}_{x,2}^{\text{no}}]. \quad (34)$$

Note that in this game, we have ~~grayed out~~ the measurement $\mathcal{M}_x^\mathcal{Y}$ since it has become irrelevant (hence removable) due to the reprogramming.

The no branch from $\text{Hyb}_{x,2}^{\text{no}}$ to $\text{Hyb}_{x,3}^{\text{no}}$. Next, we use the measure-and-reprogramming technique of [DFM20] to postpone the reprogramming in $\text{Hyb}_{x,2}^{\text{no}}$ to a randomly selected query of $\mathcal{A}$. Specifically, after key generation, $\text{Hyb}_{x,3}^{\text{no}}$ samples

$$(i, \beta) \leftarrow I_q := (\{0, \dots, q_H - 1\} \times \{0, 1\}) \cup \{(q_H, 0)\}.$$

Before the i th query to H , apply $\mathcal{M}^{\text{in}}$ to the input register to obtain (a', m') , and reprogram $H(a', m')$ to a uniformly sampled $y \leftarrow \mathcal{Y}$ before or after the query according to $\beta = 0$ or 1 respectively. Applying [DFM20, Lemma 1] then gives

$$\Pr[1 \leftarrow \text{Hyb}_{x,2}^{\text{no}}] \leq (2q_H + 1)^2 \Pr[1 \leftarrow \text{Hyb}_{x,3}^{\text{no}}]. \quad (35)$$

Wrapping up the no branch. Consider the identification attacker $\mathcal{B}$ (see Eq. (28)) which, on input pk , samples $(i, \beta) \leftarrow I_q$ and runs $\mathcal{A}$ until before its i th query. It then measures the input register to obtain (a', m') and sends a' to the identification verifier. Upon receiving the verifier's challenge y , it reprograms $H(a', m') := y$ before the $(i + \beta)$ th query, and continues the execution of $\mathcal{A}$. Finally, it outputs $\mathcal{A}$'s output z as its response. For every fixed x , the joint event that the measurement yields $(a', m') = x$ and that the above execution wins coincides with the winning event of $\text{Hyb}_{x,3}^{\text{no}}$. Summing over all x , we therefore obtain

$$\begin{aligned} \sum_x \Pr[1 \leftarrow \text{Hyb}_{x,3}^{\text{no}}] &= \sum_x \Pr[1 \leftarrow \text{Ver}(pk, a, y, z) \wedge (a', m') = x = (a, m)] \\ &= \Pr[1 \leftarrow \text{Ver}(pk, a, y, z) \wedge (a', m') = (a, m)] \\ &\leq \Pr[1 \leftarrow \text{Ver}(pk, a', y, z)] = \mathbf{Adv}_\Sigma^{\text{ID}}(\mathcal{B}, \rho_{\mathcal{B}}) \end{aligned} \quad (36)$$

The offline alternating-measurement procedure leaves the database register $\mathbf{D}$ of the compressed oracle containing at most $P = L(q_H + 1)$ recorded elements, each of $O(\log|\mathcal{X}| + \log|\mathcal{Y}|)$ bits. Using the sparse representation of the compressed oracle, during the online execution, the j th query of $\mathcal{A}$ is simulated with a database of size at most $P + j - 1$, and hence costs $O(P + j)$ running time. Summing over all q_H queries would yield

$$O\left(\sum_{j=1}^{q_H} (P + j)\right) = O(Pq_H + q_H^2) = O(Lq_H^2),$$

assuming each of the basic arithmetics takes unit time. Therefore, the running time of $\mathcal{B}$ satisfies $\text{TIME}(\mathcal{B}) \leq \text{TIME}(\mathcal{A}) + O(Lq_H^2)$. Moreover, since each database record requires $\log|\mathcal{X}| + \log|\mathcal{Y}|$ qubits, the advice $\rho_{\mathcal{B}}$ has length at most $O(L + P(\log|\mathcal{X}| + \log|\mathcal{Y}|)) = O(Lq_H(\log|\mathcal{X}| + \log|\mathcal{Y}|))$ qubits.

Combining (33) to (36) would give us

$$\sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{no}}] \leq (2q_H + 1)^2 \text{Adv}_{\Sigma}^{\text{ID}}(\mathcal{B}, \rho_{\mathcal{B}}). \quad (37)$$




Fig. 16. The yes-branch hybrid games.

The yes branch from Hyb_x^{yes} to Hyb₁^{yes}. As illustrated in Fig. 16, the only difference between Hyb_x^{yes} and Hyb₁^{yes} is that the former fixes x from the beginning of the game, while the latter samples $x \leftarrow \mathcal{X}$ uniformly at random. Conditioned on the value of x sampled in Hyb₁^{yes}, the two games are identical. Hence,

$$\Pr[1 \leftarrow \text{Hyb}_1^{\text{yes}}] = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \Pr[1 \leftarrow \text{Hyb}_x^{\text{yes}}], \quad (38)$$

The yes branch from Hyb₁^{yes} to Hyb₂^{yes}. Game Hyb₂^{yes} is obtained from Hyb₁^{yes} by moving the evaluation $y \leftarrow \mathcal{M}_x^{\text{eval}}(D)$ before key generation. Conditioned on $\text{yes} \leftarrow \mathcal{M}_x^{\mathcal{Y}}(D)$, measuring $D(x)$ to obtain y commutes with all subsequent oracle queries. Moreover, key generation does not modify the database. Hence, the evaluation $y \leftarrow \mathcal{M}_x^{\text{eval}}(D)$ can be performed earlier without changing the winning probability. Therefore,

$$\Pr[1 \leftarrow \text{Hyb}_1^{\text{yes}}] = \Pr[1 \leftarrow \text{Hyb}_2^{\text{yes}}]. \quad (39)$$

Wrapping up the yes branch. Let E_{yes} denote the event that the yes test in line 6 accepts in Hyb₂^{yes}. Note that the query bound of the offline algorithm $\mathcal{D}$ in the P -Bit-Fixing-QROM model is $P = L(q_H + 1)$, before the $\mathcal{M}_x(D)$ measurement

every compressed-database basis state in its support contains at most P non-empty elements. Since x is sampled uniformly,

$$\Pr[\mathbf{E}_{\text{yes}}] \leq \frac{P}{|\mathcal{X}|} = \frac{L(q_H + 1)}{|\mathcal{X}|}.$$

Moreover, conditioned on $\mathbf{E}_{\text{yes}}$, the values $x = (a^\circ, m^\circ)$ and $y = y^\circ$ are fixed before key generation. Hence, there exists a non-triviality adversary $\mathcal{C}$ with advice $\rho_{\mathcal{C}}$, independent of the freshly generated key pair, that continues the execution on input pk and outputs $\mathcal{A}$'s output z such that $\text{Ver}(pk, a^\circ, y^\circ, z) = 1$ whenever the original game wins. Dropping the additional requirement $x = (a, m)$ will not decrease the success probability. Therefore, by the non-triviality of the Sigma-protocol Σ (see Definition 3),

$$\Pr[1 \leftarrow \text{Hyb}_2^{\text{yes}}] = \Pr[\mathbf{E}_{\text{yes}}] \Pr[1 \leftarrow \text{Hyb}_2^{\text{yes}} \mid \mathbf{E}_{\text{yes}}] \leq \Pr[\mathbf{E}_{\text{yes}}] \mathbf{Adv}_{\Sigma}^{\neg \text{Trvl}}(\mathcal{C}, \rho_{\mathcal{C}}). \quad (40)$$

Moreover, by the same analysis as for the adversary $\mathcal{B}$ in the no branch, $\text{TIME}(\mathcal{C}) \leq \text{TIME}(\mathcal{A}) + O(Lq_H^2)$, while $\rho_{\mathcal{C}}$ has length at most $O(Lq_H(\log|\mathcal{X}| + \log|\mathcal{Y}|))$ qubits.

Combining (38) to (40) gives

$$\sum_x \Pr[1 \leftarrow \text{Hyb}_x^{\text{yes}}] \leq L(q_H + 1) \mathbf{Adv}_{\Sigma}^{\neg \text{Trvl}}(\mathcal{C}, \rho_{\mathcal{C}}). \quad (41)$$

Collecting the final bound. Combining (29),(30),(32),(37) and (41) yields:

$$\mathbf{Adv}_{\text{FS}[\Sigma]}^{\text{UF-NMA}}(\mathcal{A}) \leq 4(2q_H + 1)^2 \mathbf{Adv}_{\Sigma}^{\text{ID}}(\mathcal{B}, \rho_{\mathcal{B}}) + 4L(q_H + 1) \mathbf{Adv}_{\Sigma}^{\neg \text{Trvl}}(\mathcal{C}, \rho_{\mathcal{C}}).$$

□

We note that the converse of Theorem 6 is also true, and hence the statement characterizes when $\text{FS}[\Sigma]$ is UF-NMA and when it is not.

Indeed, every attacker $\mathcal{B}$ against ID security of Σ , can be trivially transformed into a UF-NMA attacker $\mathcal{A}_{\mathcal{B}}$ via replacing the challenge y given to the attacker with a Fiat-Shamir hash output $H(a, m)$.

On the other hand, every nontriviality attacker $\mathcal{C}$ can be transformed into a UF-NMA attacker $\mathcal{A}_{\mathcal{C}}$, via selecting the optimal a°, y° that maximizes

$$\Pr_{sk, pk} [\text{Ver}(pk, a^\circ, y^\circ, \mathcal{C}(\rho, pk, a^\circ, y^\circ))] ,$$

and then embedding a preimage m such that $H(a^\circ, m) = y^\circ$ as an additional part of the advice. Such a message would exist in λ bits with probability at least $1 - \exp(-2^\lambda/|\mathcal{Y}|)$, where $\mathcal{Y}$ is the challenge space of Σ .

We summarize the above two observations below.

Theorem 7. *Let $\Sigma = (\text{KGen}, \text{Com}, \mathcal{Y}, \text{Resp}, \text{Ver})$ be a public-key Sigma protocol, and ρ be a quantum advice of L qubits; let H be a random oracle and $\text{FS}^H[\Sigma]$ be the signature scheme obtained from applying Fiat-Shamir transformation to Σ .*

Then, for every attacker $\mathcal{B}$ and $\mathcal{C}$ against identification and nontriviality security of Σ (see Eq. (28) and Definition 3), there are similarly efficient UF-NMA attackers $\mathcal{A}_{\mathcal{B}}, \mathcal{A}_{\mathcal{C}}$, and an H -dependent message m of λ bits, such that

$$\begin{aligned}\mathbf{Adv}_{\mathbf{FS}[\Sigma]}^{\text{UF-NMA}}(\mathcal{A}_{\mathcal{B}}, \rho) &\geq \mathbf{Adv}_{\Sigma}^{\text{ID}}(\mathcal{B}, \rho) \\ \mathbf{Adv}_{\mathbf{FS}[\Sigma]}^{\text{UF-NMA}}(\mathcal{A}_{\mathcal{C}}, (\rho, m)) &\geq \mathbf{Adv}_{\Sigma}^{\text{Trvl}}(\mathcal{C}, \rho) - \exp(-2^{\lambda}/|\mathcal{Y}|) .\end{aligned}$$

for every $\lambda \in \mathbb{Z}_{>0}$.

Remark 3. Here we assume that the (partial) transcript (a°, y°) , which is a classical string chosen independent of the random oracle, is hard-coded into $\mathcal{A}_{\mathcal{C}}$. In the uniform setting where a° and y° cannot be hard-coded, we can embed them into the advice instead.

5.4 Instantiations

We briefly discuss, for three existing instantiations of (public-key) Sigma protocol Σ , whether they are secure in the sense detailed in Section 5.3. In all three examples, we consider hash and sign signature schemes, meaning that for a trapdoor function $f(x)$ sampled as public key, the prover would first send over a salt, the verifier then responds with a random image $y \leftarrow \mathcal{Y}$, and then the prover would need to respond with a preimage of y in order to convince the verifier.

In this case, nontriviality translates to the hardness of recovering an x such that $f(x) = y$, for any fixed image y chosen independently of f .

Non-key-hashing Falcon. Recall that the Falcon public key is a ring element $h \in R := \mathbb{Z}[X]/(X^n + 1)$, with the corresponding trapdoor function $f(s_1, s_2) := s_1 + s_2 h \bmod q$ defined on $s_1, s_2 \in R$ satisfying $\|(s_1, s_2)\| \leq \beta$ for a sufficiently small bound β (and some prime q). For every sufficiently small $y \in R$, the vector $(s_1, s_2) := (y, 0)$ would already be a small preimage (in norm) of y , thereby breaking the nontriviality condition of Σ . Hence, by Theorem 7, Falcon without hashing public key does not achieve UF-NMA security in the presence of advice.

Remark 4. It is conceivable that, by rejecting y that are too small, we can restore nontriviality of Σ and hence the security of $\mathbf{FS}[\Sigma]$. However in such a case the corresponding computational assumption would change, so we consider it beyond the scope of this work.

Remark 5. We emphasize here that we do not claim that Falcon is insecure. Rather, our discussion here should be seen as a supporting argument why hashing public key is necessary in Falcon.

GPV-based Signature. Alternatively, consider a public key $\mathbf{A} \in \mathbb{F}_q^{n \times m}$ ($m \geq n$) whose distribution is indistinguishable from uniform, defining a trapdoor function $f(x) := \mathbf{A}x$ on nonzero short vectors $x \in \mathbb{F}_q^m$ with $\|x\| \leq \beta$. Then, for

every fixed output $y \in \mathbb{F}_q^n$, if $y = 0$, then finding a preimage would already imply solving a SIS. Otherwise, since $\mathbf{A}$ is indistinguishable to a random matrix, we construct a reduction that, on input an SIS instance $\mathbf{B} \in \mathbb{F}_q^{(n-1) \times m}$, embeds it into the public key matrix via substituting $\mathbf{A}$ with

$$\mathbf{A}' := P^{-1} \begin{pmatrix} r \\ \mathbf{B} \end{pmatrix}, \quad r \leftarrow \mathbb{F}_q^{1 \times m}$$

where $P \in \mathbb{F}_q^{n \times n}$ is an invertible matrix that transforms Py to all-zero except for the first row. If a solution given by the nontriviality attacker x is valid, then that would mean $\mathbf{A}'x = y$, and so

$$Py = P\mathbf{A}'x = \begin{pmatrix} r \\ \mathbf{B} \end{pmatrix} x.$$

Since all but the first row of Py are zero, we have $\mathbf{B}x = 0$ being a valid SIS solution. Therefore, in this case Σ satisfies the nontriviality condition under SIS, and hence our generic security analyses in Theorems 5 and 6 apply, assuming standard assumptions on HVZK and identification security of Σ .

Wave. The same reasoning above also applies to code-based signature schemes. Taking Wave [DST19] for example, it follows a similar design principle, except that input x to the trapdoor function is restricted to a small, fixed Hamming weight. The corresponding Σ then satisfies nontriviality condition, assuming hardness of syndrome decoding that is analogous to SIS. Therefore, again, assuming standard assumptions on identification security of Σ , our UF-NMA generic analysis Theorem 6 also applies.¹⁴

AI Disclosure

The authors of this paper use AI for editorial discussion, and to generate TikZ code, from which they then refine. The attack presented in Theorem 2 is developed through multiple iterations of human discussion, with the aid of AI for brainstorming, testing ideas, and sometimes tackling easier subproblems.

Acknowledgement

The authors thank Akshima for useful discussion. Heming Liao is supported by the National Natural Science Foundation of China (Grant Nos. 62302496 and 62172405).

¹⁴ Wave uses rejection sampling which may result in a non-standard HVZK property so that our UF-CMA to UF-NMA reduction does not apply.

References

- ADGL23. Akshima, Xiaoqi Duan, Siyao Guo, and Qipeng Liu. On time-space lower bounds for finding short collisions in Sponge hash functions. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part III*, volume 14371 of *LNCS*, pages 237–270. Springer, Cham, November / December 2023.
- AGL24. Akshima, Siyao Guo, and Qipeng Liu. Time-space lower bounds for finding collisions in Merkle-Damgård hash functions. *Journal of Cryptology*, 37(2):10, April 2024.
- AHU19. Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Cham, August 2019.
- BBC⁺98. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *39th FOCS*, pages 352–361. IEEE Computer Society Press, November 1998.
- BBD⁺23. Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu. Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 358–389. Springer, Cham, August 2023.
- BCD⁺24. Manuel Barbosa, Deirdre Connolly, João Diogo Duarte, Aaron Kaiser, Peter Schwabe, Karolin Varner, and Bas Westerbaan. X-wing: The hybrid KEM you’ve been looking for. *IACR Communications in Cryptology*, 1(1), 2024.
- BDF⁺11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Berlin, Heidelberg, December 2011.
- BDK⁺18. Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Kyber: A CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy*, pages 353–367. IEEE Computer Society Press, April 2018.
- BHH⁺19. Nina Bindel, Mike Hamburg, Kathrin Hövelmanns, Andreas Hülsing, and Edoardo Persichetti. Tighter proofs of CCA security in the quantum random oracle model. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 61–90. Springer, Cham, December 2019.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- CDG18. Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. Cryptology ePrint Archive, Report 2018/226, 2018.
- CDGS18. Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 227–258. Springer, Cham, April / May 2018.

- CFHL21. Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 598–629. Springer, Cham, October 2021.
- CGH98. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998.
- CGLQ20. Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *61st FOCS*, pages 673–684. IEEE Computer Society Press, November 2020.
- CPZ24. Joseph Carolan, Alexander Poremba, and Mark Zhandry. (Quantum) indifferentiability and pre-computation. Cryptology ePrint Archive, Report 2024/1727, 2024.
- DFH⁺24. Jelle Don, Serge Fehr, Yu-Hsuan Huang, Jyun-Jie Liao, and Patrick Struck. Hide-and-seek and the non-resignability of the BUFF transform. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 347–370. Springer, Cham, December 2024.
- DFM20. Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round Fiat-Shamir and more. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 602–631. Springer, Cham, August 2020.
- DFMS19. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 356–383. Springer, Cham, August 2019.
- DFMS22. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 677–706. Springer, Cham, May / June 2022.
- DGK17. Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 473–495. Springer, Cham, April / May 2017.
- DST19. Thomas Debris-Alazard, Nicolas Sendrier, and Jean-Pierre Tillich. Wave: A new family of trapdoor one-way preimage sampleable functions based on codes. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 21–51. Springer, Cham, December 2019.
- DTT10. Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 649–665. Springer, Berlin, Heidelberg, August 2010.
- Ebr22. Ehsan Ebrahimi. Post-quantum security of plain OAEP transform. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part I*, volume 13177 of *LNCS*, pages 34–51. Springer, Cham, March 2022.
- FFH26. Pouria Fallahpour, Serge Fehr, and Yu-Hsuan Huang. Tighter quantum security for fiat-shamir-with-aborts and hash-and-sign-with-retry signatures. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, pages 603–631, Cham, 2026. Springer Nature Switzerland.

- FGdG⁺26. Pierre-Alain Fouque, Phillip Gajland, Hubert de Groote, Jonas Janneck, and Eike Kiltz. A closer look at falcon. In Joan Daemen and Emmanuel Thomé, editors, *EUROCRYPT 2026, Part IV*, volume 16545 of *LNCS*, pages 3–31. Springer, Cham, May 2026.
- FGK22. Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Time-space trade-offs for Sponge hashing: Attacks and limitations for short collisions. *Cryptography ePrint Archive*, Report 2022/1009, 2022.
- FH23. Serge Fehr and Yu-Hsuan Huang. On the quantum security of HAWK. In Thomas Johansson and Daniel Smith-Tone, editors, *PQCrypto 2023*, pages 405–416. Springer, Cham, August 2023.
- Gao15. Jingliang Gao. Quantum union bounds for sequential projective measurements. *Physical Review A*, 92(5):052331, 2015.
- GHHM21. Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the QROM. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 637–667. Springer, Cham, December 2021.
- GHP18. Federico Giacon, Felix Heuer, and Bertram Poettering. KEM combiners. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part I*, volume 10769 of *LNCS*, pages 190–218. Springer, Cham, March 2018.
- GLLZ21. Siyao Guo, Qian Li, Qipeng Liu, and Jiapeng Zhang. Unifying presampling via concentration bounds. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 177–208. Springer, Cham, November 2021.
- GLX24. Jiangxia Ge, Heming Liao, and Rui Xue. Measure-rewind-extract: Tighter proofs of one-way to hiding and CCA security in the quantum random oracle model. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 3–34. Springer, Singapore, December 2024.
- GO94. Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.
- GSX23. Jiangxia Ge, Tianshu Shan, and Rui Xue. Tighter QCCA-secure key encapsulation mechanism with explicit rejection in the quantum random oracle model. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 292–324. Springer, Cham, August 2023.
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- HBD⁺22. Andreas Hülsing, Daniel J. Bernstein, Christoph Dobraunig, Maria Eichlseder, Scott Fluhrer, Stefan-Lukas Gazdag, Panos Kampanakis, Stefan Kölbl, Tanja Lange, Martin M. Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, Peter Schwabe, Jean-Philippe Aumasson, Bas Westerbaan, and Ward Beullens. SPHINCS⁺. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- Hel80. Martin E. Hellman. A cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, 26(4):401–406, 1980.

- HHK17. Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, Cham, November 2017.
- HHM22. Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Failing gracefully: Decryption failures and the Fujisaki-Okamoto transform. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 414–443. Springer, Cham, December 2022.
- HHM⁺26. Yu-Hsuan Huang, Andreas Hülsing, Varun Maram, Silvia Ritsch, and Abishek Saha. Feistel tools: Reprogramming and query-recording for QRPs. Cryptology ePrint Archive, Paper 2026/146, 2026.
- Hua26. Yu-Hsuan Huang. *Post-Quantum Security of Cryptographic Transformations in the Random Oracle Model*. Phd thesis, Leiden University, 2026.
- Jae25. Joseph Jaeger. Nonadaptive one-way to hiding implies adaptive quantum reprogramming. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part VIII*, volume 16252 of *LNCS*, pages 287–319. Springer, Singapore, December 2025.
- JZC⁺18. Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. IND-CCA-secure key encapsulation mechanism in the quantum random oracle model, revisited. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 96–125. Springer, Cham, August 2018.
- KHN⁺14. Charlie Kaufman, Paul Hoffman, Yoav Nir, Pasi Eronen, and Tero Kivinen. Internet key exchange protocol version 2 (IKEv2). RFC 7296, 2014.
- KLS18. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Cham, April / May 2018.
- KSS⁺20. 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 Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 703–728. Springer, Cham, May 2020.
- LDK⁺22. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- LG CX25. Heming Liao, Jiangxia Ge, Shujiao Cao, and Rui Xue. New proof for plain OAEP: Post-quantum security without parameter restrictions or collision-resistance. Cryptology ePrint Archive, Report 2025/1288, 2025.
- LGXZ26. Heming Liao, Jiangxia Ge, Rui Xue, and Xiaogang Zhou. Non-adaptive one-way to hiding not only implies adaptive quantum reprogramming, but also does better. In Shi Bai and Edoardo Persichetti, editors, *PKC 2026, Part II*, volume 16552 of *LNCS*, pages 96–124. Springer, Cham, May 2026.
- Liu23. Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 117–143. Springer, Cham, April 2023.

- LLHC26. You Lyu, Shengli Liu, Shuai Han, and Bohang Chen. OAEP[†] transform in the post-quantum world. Cryptology ePrint Archive, Report 2026/1951, 2026.
- MW20. Tal Moran and Daniel Wichs. Incompressible encodings. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 494–523. Springer, Cham, August 2020.
- Nie02. Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126. Springer, Berlin, Heidelberg, August 2002.
- NIST17. National Institute of Standards and Technology. Post-quantum cryptography standardization. <https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization>, 2017.
- NIST26. National Institute of Standards and Technology. Post-quantum cryptography: Additional digital signature schemes. <https://csrc.nist.gov/Projects/pqc-dig-sig/standardization>, 2026.
- OV22. Ryan O’Donnell and Ramgopal Venkateswaran. The quantum union bound made easy. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 314–320. SIAM, 2022.
- PFH⁺22. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- PZ24. Jiaxin Pan and Runzhi Zeng. Selective opening security in the quantum random oracle model, revisited. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part II*, volume 14603 of *LNCS*, pages 92–122. Springer, Cham, April 2024.
- Res26. Eric Rescorla. The transport layer security (TLS) protocol version 1.3. RFC 9846, 2026.
- SAB⁺22. Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- SGX23. Tianshu Shan, Jiangxia Ge, and Rui Xue. QCCA-secure generic transformations in the quantum random oracle model. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 36–64. Springer, Cham, May 2023.
- Sho94. Peter W Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science*, pages 124–134. Ieee, 1994.
- TU16. Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 192–216. Springer, Berlin, Heidelberg, October / November 2016.
- Unr07a. Dominique Unruh. Random oracles and auxiliary input. Cryptology ePrint Archive, Report 2007/168, 2007.

- Unr07b. Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 205–223. Springer, Berlin, Heidelberg, August 2007.
- Unr14a. Dominique Unruh. Quantum position verification in the random oracle model. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 1–18. Springer, Berlin, Heidelberg, August 2014.
- Unr14b. Dominique Unruh. Revocable quantum timed-release encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 129–146. Springer, Berlin, Heidelberg, May 2014.
- Unr15. Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Berlin, Heidelberg, April 2015.
- Win99. Andreas J. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theory*, 45(7):2481–2485, 1999.
- Zha19. Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, Cham, August 2019.