



Quantum-Access-Secure Message Authentication via Blind-Unforgeability

Gorjan Alagic^{1(✉)}, Christian Majenz^{2(✉)}, Alexander Russell³, and Fang Song⁴

¹ QuICS, University of Maryland, and NIST, Gaithersburg, MD, USA
galagic@gmail.com

² QuSoft and Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
c.majenz@uva.nl

³ Department of Computer Science and Engineering, University of Connecticut,
Storrs, CT, USA
acr@cse.uconn.edu

⁴ Department of Computer Science and Engineering, Texas A&M University,
College Station, TX, USA
fang.song@tamu.ed

Abstract. Formulating and designing authentication of classical messages in the presence of adversaries with quantum query access has been a longstanding challenge, as the familiar classical notions of unforgeability do not directly translate into meaningful notions in the quantum setting. A particular difficulty is how to fairly capture the notion of “predicting an unqueried value” when the adversary can query in quantum superposition.

We propose a natural definition of unforgeability against quantum adversaries called *blind unforgeability*. This notion defines a function to be predictable if there exists an adversary who can use “partially blinded” oracle access to predict values in the blinded region. We support the proposal with a number of technical results. We begin by establishing that the notion coincides with EUF-CMA in the classical setting and go on to demonstrate that the notion is satisfied by a number of simple guiding examples, such as random functions and quantum-query-secure pseudo-random functions. We then show the suitability of blind unforgeability for supporting canonical constructions and reductions. We prove that the “hash-and-MAC” paradigm and the Lamport one-time digital signature scheme are indeed unforgeable according to the definition. To support our analysis, we additionally define and study a new variety of quantum-secure hash functions called *Bernoulli-preserving*.

Finally, we demonstrate that blind unforgeability is strictly stronger than a previous definition of Boneh and Zhandry [EUROCRYPT ’13, CRYPTO ’13] and resolve an open problem concerning this previous definition by constructing an explicit function family which is forgeable yet satisfies the definition.

1 Introduction

Large-scale quantum computers will break widely-deployed public-key cryptography, and may even threaten certain post-quantum candidates [6, 9, 10, 12, 23].

Even elementary symmetric-key constructions like Feistel ciphers and CBC-MACs become vulnerable in quantum attack models where the adversary is presumed to have quantum query access to some part of the cryptosystem [16–18, 22]. As an example, consider encryption in the setting where the adversary has access to the unitary operator $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus f_k(x)\rangle$, where f_k is the encryption or decryption function with secret key k . While it is debatable if this model reflects physical implementations of symmetric-key cryptography, it appears necessary in a number of generic settings, such as public-key encryption and hashing with public hash functions. It could also be relevant when private-key primitives are composed in larger protocols, e.g., by exposing circuits via obfuscation [21]. Setting down appropriate security definitions in this quantum attack model is the subject of several threads of recent research [8, 13].

In this article, we study authentication of classical information in the quantum-secure model. Here, the adversary is granted quantum query access to the signing algorithm of a message authentication code (MAC) or a digital signature scheme, and is tasked with producing valid forgeries. In the purely classical setting, we insist that the forgeries are fresh, i.e., distinct from previous queries to the oracle. When the function may be queried in superposition, however, it’s unclear how to meaningfully reflect this constraint that a forgery was previously “unqueried.” For example, it is clear that an adversary that simply queries with a uniform superposition and then measures a forgery—a feasible attack against any function—should not be considered successful. On the other hand, an adversary that uses the same query to discover some structural property (e.g., a superpolynomial-size period in the MAC) should be considered a break. Examples like these indicate the difficulty of the problem. How do we correctly “price” the queries? How do we decide if a forgery is fresh? Furthermore, how can this be done in a manner that is consistent with these guiding examples? In fact, this problem has a natural interpretation that goes well beyond cryptography: *What does it mean for a classical function to appear unpredictable to a quantum oracle algorithm?*¹

Previous approaches. The first approach to this problem was suggested by Boneh and Zhandry [7]. They define a MAC to be unforgeable if, after making q queries to the MAC, no adversary can produce $q + 1$ valid input-output pairs except with negligible probability. We will refer to this notion as “PO security” (PO for “plus one,” and k -PO when the adversary is permitted a maximum of k queries). Among a number of results, Boneh and Zhandry prove that this notion can be realized by a quantum-secure pseudorandom function (qPRF).

Another approach, due to Garg, Yuen and Zhandry [14] (GYZ), considers a function *one-time* unforgeable if only a trivial “query, measure in computational basis, output result” attack² is allowed. Unfortunately, it is not clear how to

¹ The related notion of “appearing *random* to quantum oracle algorithms” has a satisfying definition, which can be fulfilled efficiently [29].

² Technically, the *Stinespring dilation* [25] of a computational basis measurement is the most general attack.

extend GYZ to two or more queries. Furthermore, the single query is allowed in a limited query model with an non-standard restriction.³ Zhandry recently showed a separation between PO and GYZ by means of the powerful tool of obfuscation [31].

It is interesting to note that similar problems arise in encryption schemes of *quantum* data and a convincing solution was recently found [2,3]. However, it relies on the fact that for quantum messages, *authentication implies secrecy*. This enables “tricking” the adversary by replacing their queries with “trap” plaintexts to detect replays. As unforgeability and secrecy are orthogonal in the classical world, adversaries would easily recognize the spoofed oracle. This renders the approach of [2,3] inapplicable in this case.

Unresolved issues. PO security, the only candidate definition of quantum-secure unforgeability in the general, multi-query setting, appears to be insufficient for several reasons. First, as observed in [14], it is a priori unclear if PO security rules out forging on a message region A while making queries to a signing oracle supported on a disjoint message region B . Second, there may be unique features of quantum information, such as the destructiveness of quantum measurement, which PO does not capture. In particular, quantum algorithms must sometimes “consume” (i.e., fully measure) a state to extract some useful information, such as a symmetry in the oracle. There might be an adversary that makes one or more quantum queries but then must consume the post-query states completely in order to make a single, but convincing, forgery.

Surprisingly, prior to this work none of these plausible attack strategies have been exploited to give a separation between PO and “intuitive security.”

2 Summary of Results

A new definition: Blind-unforgeability. To address the above mentioned issues, and in light of the concrete “counterexample” presented below as Construction 8, we develop a new definition of many-time unforgeability we call “blind-unforgeability” (or BU). In this approach we examine the behavior of adversaries in the following experiment. The adversary is granted quantum oracle access to the MAC, “blinded” at a random region B . Specifically, we set B to be a random ϵ -fraction of the message space, and declare that the oracle function will output \perp on all of B .

$$B_\epsilon \text{Mac}_k(x) := \begin{cases} \perp & \text{if } x \in B_\epsilon, \\ \text{Mac}_k(x) & \text{otherwise.} \end{cases}$$

Given a MAC (Mac, Ver) , an adversary \mathcal{A} , and \mathcal{A} -selected parameter ϵ , the “blind forgery experiment” is:

³ Compared to the standard quantum oracle for a classical function, GYZ require the output register to be empty prior to the query.

1. Generate key k and random blinding B_ϵ ;
2. Produce candidate forgery $(m, t) \leftarrow \mathcal{A}^{B_\epsilon \text{Mac}_k}(1^n)$.
3. Output win if $\text{Ver}_k(m, t) = \text{acc}$ and $m \in B_\epsilon$; otherwise output rej .

Definition 1. A MAC is *blind-unforgeable* (BU) if for every adversary (\mathcal{A}, ϵ) , the probability of winning the blind forgery experiment is negligible.

In this work, BU will typically refer to the case where \mathcal{A} is an efficient quantum algorithm (QPT) and the oracle is quantum, i.e., $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus B_\epsilon \text{Mac}_k(x)\rangle$. We will also consider q -BU, the information-theoretic variant where the total number of queries is a priori fixed to q . We remark that the above definition is also easy to adapt to other settings, e.g., classical security against PPT adversaries, quantum or classical security for digital signatures, etc.

We remark that one could define a variant of the above where the adversary is allowed to describe the blinding distribution, rather than it being uniform. However, this is not a stronger notion. By a straightforward argument, an adversary wins in the chosen-blinding BU game if and only if it wins with a uniform ϵ -blinding for inverse-polynomial ϵ . Indeed, the adversary can just simulate its chosen blinding herself, and this still succeeds with inverse polynomial probability when interacting with a standard-blinded oracle (see Theorem 2 below).

Results about blind-unforgeability. To solidify our confidence in the new notion, we collect a series of results which we believe establish BU as a definition of unforgeability that captures the desired intuitive security requirement. In particular, we show that BU is strictly stronger than previous candidate definitions, and that it classifies a wide range of representative examples (in fact, all examples examined thus far) as either forgeable or unforgeable in a way that agrees with cryptographic intuition.

Relations and characterizations. First, we show that BU correctly classifies unforgeability in the classical-query setting: it is equivalent to the classical unforgeability notion of EUF-CMA (existential unforgeability against chosen-message attack). Then, we show that it implies PO.

Theorem 1. *If a function family is BU-unforgeable, then it is PO-unforgeable.*

One key technical component of the proof is a general simulation theorem, which tightly controls the deviation in the behavior of an algorithm when subjected to the BU experiment.

Theorem 2. *Let \mathcal{A} be a quantum query algorithm making at most T queries. Let $f : X \rightarrow Y$ be a function, B_ϵ a random ϵ -blinding subset of X , and for each $B \subset X$, let g_B a function with support B . Then*

$$\mathbb{E}_{B_\epsilon} \|\mathcal{A}^f(1^n) - \mathcal{A}^{f \oplus g_{B_\epsilon}}(1^n)\|_1 \leq 2T\sqrt{\epsilon}.$$

This result can be viewed as strong evidence that algorithms that produce “good forgeries” in any reasonable sense will also win the BU experiment.

Specifically, adversaries that produce “good forgeries” will not be disturbed too much by blinding, and will thus in fact also win the BU experiment with non-negligible probability.

We can formulate and prove this intuition explicitly for a wide class of adversaries, as follows. Given an oracle algorithm \mathcal{A} , we let $\text{supp}(\mathcal{A})$ denote the union of the supports of all the queries of \mathcal{A} , taken over all choices of oracle function.

Theorem 3 (informal). *Let \mathcal{A} be QPT and $\text{supp}(\mathcal{A}) \cap R = \emptyset$ for some $R \neq \emptyset$. Let Mac be a MAC, and suppose $\mathcal{A}^{\text{Mac}_k}(1^n)$ outputs a valid pair $(m, \text{Mac}_k(m))$ with $m \in R$ with noticeable probability. Then Mac is not BU secure.*

Blind-unforgeable MACs. Next, we show that several natural constructions satisfy BU. We first show that a random function is blind-unforgeable.

Theorem 4. *Let $R : X \rightarrow Y$ be a random function such that $1/|Y|$ is negligible. Then R is a blind-unforgeable MAC.*

By means of results of Zhandry [29] and Boneh and Zhandry [7], this leads to efficient BU-secure constructions.

Corollary 1. *Quantum-secure pseudorandom functions (qPRF) are BU-secure MACs, and $(4q+1)$ -wise independent functions are q -BU-secure MACs.*

We can then invoke a recent result about the quantum-security of domain-extension schemes such as NMAC and HMAC [24], and obtain variable-length BU-secure MACs from any qPRF.

In the setting of public verification, we show that the one-time Lamport signature scheme [19] is BU-secure, provided that the underlying hash function family $\mathcal{R} : X \rightarrow Y$ is modeled as a random oracle.

Theorem 5. *Let $\mathcal{R} : X \rightarrow Y$ be a random function family. Then the Lamport scheme $L_{\mathcal{R}}$ is BU against adversaries which make one quantum query to $L_{\mathcal{R}}$ and poly-many quantum queries to \mathcal{R} .*

Hash-and-MAC. Consider the following natural variation on the blind-forgery experiment. To blind $F : X \rightarrow Y$, we first select a hash function $h : X \rightarrow Z$ and a blinding set $B_\epsilon \subseteq Z$; we then declare that F will be blinded on $x \in X$ whenever $h(x) \in B_\epsilon$. We refer to this as “hash-blinding.” We say that a hash function h is a Bernoulli-preserving hash if, for every oracle function F , no QPT can distinguish between an oracle that has been hash-blinded with h , and an oracle that has been blinded in the usual sense. Recall the notion of *collapsing* from [27].

Theorem 6. *Let $h : X \rightarrow Y$ be a hash function. If h is Bernoulli-preserving hash, then it is also collapsing. Moreover, against adversaries with classical oracle access, h is a Bernoulli-preserving hash if and only if it is collision-resistant.*

We apply this new notion to show security of the Hash-and-MAC construction $\Pi^h = (\text{Mac}^h, \text{Ver}^h)$ with $\text{Mac}_k^h(m) := \text{Mac}_k(h(m))$.

Theorem 7. *Let $\Pi = (\text{Mac}_k, \text{Ver}_k)$ be a BU-secure MAC with $\text{Mac}_k : X \rightarrow Y$, and let $h : Z \rightarrow X$ a Bernoulli-preserving hash. Then Π^h is a BU-secure MAC.*

We also show that the Bernoulli-preserving property can be satisfied by pseudorandom constructions, as well as a (public-key) hash based on *lossy functions* from LWE [20, 26].

A concrete “counterexample” for PO. Supporting our motivation to devise a new unforgeability definition, we present a construction of a MAC which is forgeable (in a strong intuitive sense) and yet is classified by PO as secure.

Construction 8. Given a triple $k = (p, f, g)$ where $p \in \{0, 1\}^n$ and $f, g : \{0, 1\}^n \rightarrow \{0, 1\}^n$, define $M_k : \{0, 1\}^{n+1} \rightarrow \{0, 1\}^{2n}$ by

$$M_k(x) = \begin{cases} 0^{2n} & x = 0\|p, \\ 0^n\|f(x') & x = 0\|x', \ x' \neq p, \\ g(x' \bmod p)\|f(x') & x = 1\|x'. \end{cases}$$

Define $g_p(x) := g(x \bmod p)$ and consider an adversary that queries only on messages starting with 1, as follows:

$$\sum_{x,y} |1, x\rangle_X |0^n\rangle_{Y_1} |y\rangle_{Y_2} \mapsto \sum_{x,y} |1, x\rangle_X |g_p(x)\rangle_{Y_1} |y \oplus f(x)\rangle_{Y_2}; \tag{1}$$

discarding the first qubit and Y_2 then yields $\sum_x |x\rangle |g_p(x)\rangle$, as $\sum_y |y \oplus f(x)\rangle_{Y_2} = \sum_y |y\rangle_{Y_2}$. One can then recover p via period-finding and output $(0\|p, 0^{2n})$. We emphasize that the forgery was queried with *zero* amplitude. In practice, we can interpret it as, e.g., the attacker queries only on messages starting with “From: Alice” and then forges a message starting with “From: Bob”. Despite this, we can show that it is PO-secure.

Theorem 9. *The family M_k (for uniformly random $k = (p, f, g)$) is PO-secure.*

The PO security of M relies on a dilemma the adversary faces at each query: either learn an output of f , or obtain a superposition of $(x, g(x))$ -pairs for Fourier sampling. Our proof shows that, once the adversary commits to one of these two choices, the other option is irrevocably lost. Our result can thus be understood as a refinement of an observation of Aaronson: quantumly learning a property sometimes requires *uncomputing* some information [1]. Note that, while Aaronson could rely on standard (asymptotic) query complexity techniques, our problem is quite fragile: PO security describes a task which should be hard with q queries, but is completely trivial given $q + 1$ queries. Our proof makes use of a new quantum random oracle technique of Zhandry [30].



Unforgeability against classical adversaries Unforgeability against quantum adversaries

Fig. 1. Relationship between different unforgeability notions

A straightforward application of Theorem 3 shows that Construction 8 is BU-insecure. In particular, we have the following.

Corollary 2. *There exists a PO-secure MAC which is BU-insecure.*

The relationship between BU, PO some other notions are visualized in Fig. 1.

3 Preliminaries

Basic notation, conventions. Given a finite set X , the notation $x \in_R X$ will mean that x is a uniformly random element of X . Given a subset B of a set X , let $\chi_B : X \rightarrow \{0, 1\}$ denote the characteristic function of B , i.e., $\chi_B(x) = 1$ if $x \in B$ and $\chi_B(x) = 0$ else. When we say that a classical function F is efficiently computable, we mean that there exists a uniform family of deterministic classical circuits which computes F . We will consider three classes of algorithms: (i) unrestricted algorithms, modeling computationally unbounded adversaries, (ii) probabilistic poly-time algorithms (PPTs), modeling classical adversaries, and (iii) quantum poly-time algorithms (QPTs), modeling quantum adversaries. We assume that the latter two are given as polynomial-time uniform families of circuits. For PPTs, these are probabilistic circuits. For QPTs, they are quantum circuits, which may contain both unitary gates and measurements. We will often assume (without loss of generality) that the measurements are postponed to the end of the circuit, and that they take place in the computational basis. Given an algorithm \mathcal{A} , we let $\mathcal{A}(x)$ denote the (in general, mixed) state output by \mathcal{A} on input x . In particular, if \mathcal{A} has classical output, then $\mathcal{A}(x)$ denotes a probability distribution. Unless otherwise stated, the probability is taken over all random coins and measurements of \mathcal{A} , and any randomness used to select the input x . If \mathcal{A} is an oracle algorithm and F a classical function, then $\mathcal{A}^F(x)$ is the mixed state output by \mathcal{A} equipped with oracle F and input x ; the probability is now also taken over any randomness used to generate F .

We will distinguish between two ways of presenting a function $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ as an oracle. First, the usual “classical oracle access” simply means that each oracle call grants one classical invocation $x \mapsto F(x)$. This will always be the oracle model for PPTs. Second, “quantum oracle access” will mean that each oracle call grants an invocation of the $(n + m)$ -qubit unitary gate $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus F(x)\rangle$. For us, this will always be the oracle model for QPTs. Note that both QPTs and unrestricted algorithms could in principle receive either oracle type.

We will need the following lemma. We use the formulation from [8, Lemma 2.1], which is a special case of a more general “pinching lemma” of Hayashi [15].

Lemma 1. *Let \mathcal{A} be a quantum algorithm and $x \in \{0, 1\}^*$. Let \mathcal{A}_0 be another quantum algorithm obtained from \mathcal{A} by pausing \mathcal{A} at an arbitrary stage of execution, performing a partial measurement that obtains one of k outcomes, and then resuming \mathcal{A} . Then $\Pr[\mathcal{A}_0(1^n) = x] \geq \Pr[\mathcal{A}(1^n) = x]/k$.*

We denote the trace distance between states ρ and σ by $\delta(\rho, \sigma)$. Recall its definition via the trace norm, i.e., $\delta(\rho, \sigma) = (1/2)\|\rho - \sigma\|_1$. When ρ and σ are classical states, the trace distance is equal to the total variation distance.

Quantum-secure pseudorandomness. A quantum-secure pseudorandom function (qPRF) is a family of classical, deterministic, efficiently-computable functions which appear random to QPT adversaries with quantum oracle access.

Definition 2. *An efficiently computable function family $f : K \times X \rightarrow Y$ is a quantum-secure pseudorandom function (qPRF) if, for all QPTs \mathcal{D} ,*

$$\left| \Pr_{k \in_R K} [\mathcal{D}^{f_k}(1^n) = 1] - \Pr_{g \in_R \mathcal{F}_X^Y} [\mathcal{D}^g(1^n) = 1] \right| \leq \text{negl}(n).$$

Here \mathcal{F}_X^Y denotes the set of all functions from X to Y . The standard ‘‘GGM+GL’’ construction of a PRF yields a qPRF when instantiated with a quantum-secure one-way function [29]. One can also construct a qPRF directly from the Learning with Errors assumption [29]. If we have an a priori bound on the number of allowed queries, then a computational assumption is not needed.

Theorem 10. (Lemma 6.4 in [7]). *Let $q, c \geq 0$ be integers, and $f : K \times X \rightarrow Y$ a $(2q + c)$ -wise independent family of functions. Let \mathcal{D} be an algorithm making no more than q quantum oracle queries and c classical oracle queries. Then*

$$\Pr_{k \in_R K} [\mathcal{D}^{f_k}(1^n) = 1] = \Pr_{g \in_R \mathcal{F}_X^Y} [\mathcal{D}^g(1^n) = 1].$$

PO-unforgeability. Boneh and Zhandry define unforgeability (against quantum queries) for classical MACs as follows [7]. They also show that random functions satisfy this notion.

Definition 3. *Let $\Pi = (\text{KeyGen}, \text{Mac}, \text{Ver})$ be a MAC with message set X . Consider the following experiment with an algorithm \mathcal{A} :*

1. Generate key: $k \leftarrow \text{KeyGen}(1^n)$.
2. Generate forgeries: \mathcal{A} receives quantum oracle for Mac_k , makes q queries, and outputs a string s ;
3. Outcome: output win if s contains $q + 1$ distinct input-output pairs of Mac_k , and fail otherwise.

We say that Π is PO-secure if no adversary can succeed at the above experiment with better than negligible probability.

The Fourier Oracle. Our separation proof will make use of a new technique of Zhandry [30] for analyzing random oracles. We briefly describe this framework.

A random function f from n bits to m bits can be viewed as the outcome of a quantum measurement. More precisely, let $\mathcal{H}_F = \bigotimes_{x \in \{0,1\}^n} \mathcal{H}_{F_x}$, where $\mathcal{H}_{F_x} \cong \mathbb{C}^{2^m}$. Then set $f(x) \leftarrow \mathcal{M}_{F_x}(\eta_F)$ with $\eta_F = |\phi_0\rangle\langle\phi_0|^{\otimes 2^n}$, $|\phi_0\rangle = 2^{-\frac{m}{2}} \sum_{y \in \{0,1\}^m} |y\rangle$, and where \mathcal{M}_{F_x} denotes the measurement of the register F_x in the computational basis. This measurement commutes with any $\text{CNOT}_{A:B}$ gate with control qubit A in F_x and target qubit B outside F_x . It follows that, for any quantum algorithm making queries to a random oracle, the output distribution is identical if the algorithm is instead run with the following oracle:

1. Setup: prepare the state η_F .
2. Upon a query with query registers X and Y , controlled on X being in state $|x\rangle$, apply $(\text{CNOT}^{\otimes m})_{F_x:Y}$.
3. After the algorithm has finished, measure F to determine the success of the computation.

We denote the oracle unitary defined in step 2 above by U_{XYF}^O . Having defined this oracle representation, we are free to apply any unitary U_H to the oracle state, so long as we then also apply the conjugated query unitary $U_H(\text{CNOT}^{\otimes m})_{F_x:Y}U_H^\dagger$ in place of U_{XYF}^O . We choose $U_H = H^{\otimes m2^n}$, which means that the oracle register starts in the all-zero state now. Applying Hadamard to both qubits reverses the direction of CNOT, i.e., $H_A \otimes H_B \text{CNOT}_{A:B} H_A \otimes H_B = \text{CNOT}_{B:A}$, so the adversary-oracle-state after a first query with query state $|x\rangle_X |\phi_y\rangle_Y$ is

$$|x\rangle_X |\phi_y\rangle_Y |0^m\rangle^{\otimes 2^n} \longmapsto |x\rangle_X |\phi_y\rangle_Y |0^m\rangle^{\otimes (\text{lex}(x)-1)} |y\rangle_{F_x} |0^m\rangle^{\otimes (2^n - \text{lex}(x))}, \quad (2)$$

where $\text{lex}(x)$ denotes the position of x in the lexicographic ordering of $\{0,1\}^n$, and we defined the Fourier basis state $|\phi_y\rangle = H^{\otimes m}|y\rangle$. In the rest of this section, we freely change the order in which tensor products are written, and keep track of the tensor factors through the use of subscripts. This adjusted representation is called the *Fourier oracle* (FO), and we denote its oracle unitary by

$$U_{XYF}^{\text{FO}} = \left(H^{\otimes m2^n} \right)_F U_{XYF}^O \left(H^{\otimes m2^n} \right)_F.$$

An essential fact about the FO is that each query can only change the number of non-zero entries in the FO’s register by at most one. To formalize this idea, we define the “number operator” $N_F = \sum_{x \in \{0,1\}^n} (\mathbb{1} - |0\rangle\langle 0|)_{F_x} \otimes \mathbb{1}^{\otimes (2^n - 1)}$. The number operator can also be written in its spectral decomposition,

$$N_F = \sum_{l=0}^{2^n} l P_l \quad \text{where} \quad P_l = \sum_{r \in S_l} |r\rangle\langle r|,$$

$$S_l = \left\{ r \in (\{0,1\}^m)^{2^n} \mid |\{x \in \{0,1\}^n \mid r_x \neq 0\}| = l \right\}.$$

Note that the initial joint state of a quantum query algorithm and the oracle (in the FO-oracle picture described above) is in the image of P_0 . The following fact is essential in working with the Fourier Oracle; the proof is given in Appendix A.

Lemma 2. *The number operator satisfies $\| [N_F, U_{XYF}^{\text{FO}}] \|_\infty = 1$. In particular, the joint state of a quantum query algorithm and the oracle after the q -th query is in the kernel of P_l for all $l > q$.*

4 The New Notion: Blind-Unforgeability

Formal definition. For ease of exposition, we begin by introducing our new security notion in a form analogue to the standard notion of existential unforgeability under chosen-message attacks, EUF-CMA. We will also later show how to extend our approach to obtain a corresponding analogue of strong unforgeability. We begin by defining a “blinding” operation. Let $f : X \rightarrow Y$ and $B \subseteq X$. We let

$$Bf(x) = \begin{cases} \perp & \text{if } x \in B, \\ f(x) & \text{otherwise.} \end{cases}$$

We say that f has been “blinded” by B . In this context, we will be particularly interested in the setting where elements of X are placed in B independently at random with a particular probability ϵ ; we let B_ϵ denote this random variable. (It will be easy to infer X from context, so we do not reflect it in the notation.)

Next, we define a security game in which an adversary is tasked with using a blinded MAC oracle to produce a valid input-output pair in the blinded set.

Definition 4. Let $\Pi = (\text{KeyGen}, \text{Mac}, \text{Ver})$ be a MAC with message set X . Let \mathcal{A} be an algorithm, and $\epsilon : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ an efficiently computable function. The blind forgery experiment $\text{BlindForge}_{\mathcal{A}, \Pi}(n, \epsilon)$ proceeds as follows:

1. Generate key: $k \leftarrow \text{KeyGen}(1^n)$.
2. Generate blinding: select $B_\epsilon \subseteq X$ by placing each m into B_ϵ independently with probability $\epsilon(n)$.
3. Produce forgery: $(m, t) \leftarrow \mathcal{A}^{B_\epsilon \text{Mac}_k}(1^n)$.
4. Outcome: output 1 if $\text{Ver}_k(m, t) = \text{acc}$ and $m \in B_\epsilon$; otherwise output 0.

We say that a scheme is blind-unforgeable if, for any efficient adversary, the probability of winning the game is negligible. The probability is taken over the choice of key, the choice of blinding set, and any internal randomness of the adversary. We remark that specifying an adversary requires specifying (in a uniform fashion) both the algorithm \mathcal{A} and the blinding fraction ϵ .

Definition 5. A MAC Π is blind-unforgeable (BU) if for every polynomial-time uniform adversary (\mathcal{A}, ϵ) , $\Pr[\text{BlindForge}_{\mathcal{A}, \Pi}(n, \epsilon(n)) = 1] \leq \text{negl}(n)$.

We also define the “ q -time” variant of the blinded forgery game, which is identical to Definition 4 except that the adversary is only allowed to make q queries to $B_\epsilon \text{Mac}_k$ in step (3). We call the resulting game $\text{BlindForge}_{\mathcal{A}, \Pi}^q(n, \epsilon)$, and give the corresponding definition of q -time security (now against computationally unbounded adversaries).

Definition 6. A MAC Π is q -time blind-unforgeable (q -BU) if for every q -query adversary (\mathcal{A}, ϵ) , we have $\Pr[\text{BlindForge}_{\mathcal{A}, \Pi}^q(n, \epsilon(n)) = 1] \leq \text{negl}(n)$.

The above definitions are agnostic regarding the computational power of the adversary and the type of oracle provided. For example, selecting PPT adversaries and classical oracles in Definition 5 yields a definition of classical unforgeability; we will later show that this is equivalent to standard EUF-CMA. The main focus of our work will be on BU against QPTs with quantum oracle access, and q -BU against unrestricted adversaries with quantum oracle access.

Some technical details. We now remark on a few details in the usage of BU. First, strictly speaking, the blinding sets in the security games above cannot be generated efficiently. However, a pseudorandom blinding set will suffice. Pseudorandom blinding sets can be generated straightforwardly using an appropriate pseudorandom function, such as a PRF against PPTs or a qPRF against QPT. A precise description of how to perform this pseudorandom blinding is given in the proof of Corollary 3. Note that simulating the blinding requires computing and uncomputing the random function, so we must make two quantum queries for each quantum query of the adversary. Moreover, verifying whether the forgery is in the blinding set at the end requires one additional classical query. This means that $(4q + 1)$ -wise independent functions are both necessary and sufficient for generating blinding sets for q -query adversaries (see [7, Lemma 6.4]). In any case, an adversary which behaves differently in the random-blinding game versus the pseudorandom-blinding game immediately yields a distinguisher against the corresponding pseudorandom function.

The Blinding Symbol. There is some flexibility in how one defines the blinding symbol \perp . In situations where the particular instantiation of the blinding symbol might matter, we will adopt the convention that the blinded version Bf of $f : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ is defined by setting $Bf : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell+1}$, where $Bf(m) = 0^\ell \| 1$ if $m \in B$ and $Bf(m) = f(m) \| 0$ otherwise. One advantage of this convention (i.e., that $\perp = 0^\ell \| 1$) is that we can compute on and/or measure the blinded bit (i.e., the $(\ell + 1)$ -st bit) without affecting the output register of the function. This will also turn out to be convenient for uncomputation.

Strong Blind-Unforgeability. The security notion BU given in Definition 5 is an analogue of simple unforgeability, i.e., EUF-CMA, for the case of a quantum-accessible MAC/Signing oracle. It is, however, straightforward to define a corresponding analogue of strong unforgeability, i.e., SUF-CMA, as well.

The notion of strong blind-unforgeability, sBU, is obtained by a simple adjustment compared to BU: we blind (message, tag) pairs rather than just messages. We briefly describe this for the case of MACs. Let $\Pi = (\text{KeyGen}, \text{Mac}, \text{Ver})$ be a MAC with message set M , randomness set R and tag set T , so that $\text{Mac}_k : M \times R \rightarrow T$ and $\text{Ver}_k : M \times T \rightarrow \{\text{acc}, \text{rej}\}$ for every $k \leftarrow \text{KeyGen}$. Given a parameter ϵ and an adversary \mathcal{A} , the strong blind forgery game proceeds as follows:

1. Generate key: $k \leftarrow \text{KeyGen}$; generate blinding: select $B_\epsilon \subseteq M \times T$ by placing pairs (m, t) in B_ϵ independently with probability ϵ ;
2. Produce forgery: produce (m, t) by executing $\mathcal{A}(1^n)$ with quantum oracle access to the function

$$B_\epsilon \text{Mac}_{k;r}(m) := \begin{cases} \perp & \text{if } (m, \text{Mac}_k(m; r)) \in B_\epsilon, \\ \text{Mac}_k(m; r) & \text{otherwise.} \end{cases}$$

where r is sampled uniformly for each oracle call.

3. Outcome: output 1 if $\text{Ver}_k(m, t) = \text{acc} \wedge (m, t) \in B_\epsilon$; otherwise output 0.

Security is then defined as before: Π is sBU-secure if for all adversaries \mathcal{A} (and their declared ϵ), the success probability at winning the above game is negligible. Note that, for the case of canonical MACs, this definition coincides with Definition 5, just as EUF-CMA and SUF-CMA coincide in this case.

5 Intuitive Security and the Meaning of BU

In this section, we gather a number of results which build confidence in BU as a correct definition of unforgeability in our setting. We begin by showing that a wide range of “intuitively forgeable” MACs (indeed, all such examples we have examined) are correctly characterized by BU as insecure.

Intuitively forgeable schemes. As indicated earlier, BU security rules out any MAC schemes where an attacker can query a subset of the message space and forge outside that region. To make this claim precise, we first define the *query support* $\text{supp}(\mathcal{A})$ of an oracle algorithm \mathcal{A} . Let \mathcal{A} be a quantum query algorithm with oracle access to the quantum oracle O for a classical function from n to m bits. Without loss of generality \mathcal{A} proceeds by applying the sequence of unitaries $\mathcal{O}U_q\mathcal{O}U_{q-1}\dots U_1$ to the initial state $|0\rangle_{XYZ}$, followed by a POVM \mathcal{E} . Here, X and Y are the input and output registers of the function and Z is the algorithm’s workspace. Let $|\psi_i\rangle$ be the intermediate state of \mathcal{A} after the application of U_i . Then $\text{supp}(\mathcal{A})$ is defined to be the set of input strings x such that there exists a function $f : \{0,1\}^n \rightarrow \{0,1\}^m$ such that $\langle x|_X|\psi_i\rangle \neq 0$ for at least one $i \in \{1, \dots, q\}$ when $\mathcal{O} = \mathcal{O}_f$.

Theorem 11. *Let \mathcal{A} be a QPT such that $\text{supp}(\mathcal{A}) \cap R = \emptyset$ for some $R \neq \emptyset$. Let Mac be a MAC, and suppose $\mathcal{A}^{\text{Mac}_k}(1^n)$ outputs a valid pair $(m, \text{Mac}_k(m))$ with $m \in R$ with non-negligible probability. Then Mac is not BU-secure.*

To prove Theorem 11, we will need the following theorem, which controls the change in the output state of an algorithm resulting from applying a blinding to its oracle. Given an oracle algorithm \mathcal{A} and two oracles F and G , the trace distance between the output of \mathcal{A} with oracle F and \mathcal{A} with oracle G is denoted by $\delta(\mathcal{A}^F(1^n), \mathcal{A}^G(1^n))$. Given two functions $F, P : \{0,1\}^n \rightarrow \{0,1\}^m$, we define the function $F \oplus P$ by $(F \oplus P)(x) = F(x) \oplus P(x)$.

Theorem 12. *Let \mathcal{A} be a quantum query algorithm making at most T queries, and $F : \{0,1\}^n \rightarrow \{0,1\}^m$ a function. Let $B \subseteq \{0,1\}^n$ be a subset chosen by independently including each element of $\{0,1\}^n$ with probability ϵ , and $P : \{0,1\}^n \rightarrow \{0,1\}^m$ be any function with support B . Then*

$$\mathbb{E}_B[\delta(\mathcal{A}^F(1^n), \mathcal{A}^{F \oplus P}(1^n))] \leq 2T\sqrt{\epsilon}.$$

The proof is a relatively straightforward adaptation of a hybrid argument in the spirit of the lower bound for Grover search [5]. We provide the complete proof in the full version [4]. We are now ready to prove Theorem 11.

Proof (of Theorem 11). Let \mathcal{A} be a quantum algorithm with $\text{supp}(\mathcal{A})$ for any oracle. By our hypothesis,

$$\tilde{p} := \Pr_{k,(m,t) \leftarrow \mathcal{A}^{\text{Mac}_k(1^n)}} [\text{Mac}_k(m) = t \wedge m \notin \text{supp}(\mathcal{A})] \geq n^{-c},$$

for some $c > 0$ and sufficiently large n . Since $\text{supp}(A)$ is a fixed set, we can think of sampling a random B_ε as picking $B_0 := B_\varepsilon \cap \text{supp}(A)$ and $B_1 := B_\varepsilon \cap \overline{\text{supp}(A)}$ independently. Let “blind” denote the random experiment of \mathcal{A} running on Mac_k blinded by a random B_ε : $k, B_\varepsilon, (m, t) \leftarrow \mathcal{A}^{B_\varepsilon \text{Mac}_k(1^n)}$, which is equivalent to $k, B_0, B_1, (m, t) \leftarrow \mathcal{A}^{B_0 \text{Mac}_k(1^n)}$. The probability that \mathcal{A} wins the BU game is

$$\begin{aligned} p &:= \Pr_{\text{blind}} [f(m) = t \wedge m \in B_\varepsilon] \geq \Pr_{\text{blind}} [f(m) = t \wedge m \in B'] \\ &\geq \Pr_{\text{blind}} [f(m) = t \wedge m \in B' \mid m \notin \text{supp}(A)] \cdot \Pr_{\text{blind}} [m \notin \text{supp}(A)] \\ &= \Pr_{\substack{f, B_0 \\ (m,t) \leftarrow \mathcal{A}^{B_f}}} [f(m) = t \wedge m \notin \text{supp}(A)] \cdot \Pr_{\substack{f, B' \\ (m,t) \leftarrow \mathcal{A}^{B_f}}} [m \in B' \mid m \notin \text{supp}(A)] \\ &\geq (\tilde{p} - 2T\sqrt{\varepsilon}) \varepsilon \geq \frac{\tilde{p}^3}{27T^2}. \end{aligned}$$

Here the second-to-last step follows from Theorem 12; in the last step, we chose $\varepsilon = (\tilde{p}/3T)^2$. We conclude that \mathcal{A} breaks the BU security of the MAC. \square

Relationship to other definitions. As we will show in Sect. 7, PO fails to capture certain prediction algorithms. It does, however, capture a natural family of attacks and should hence be implied by a good security notion. In this section we show that our new definition, BU, indeed implies PO. To this end, we first introduce a natural weaker variant of BU that we call measured BU, or mBU.

Definition 7. *The measured ε -blinded oracle for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is the oracle that first applies the ε -blinded oracle for f and then performs the projective measurement $|\perp\rangle\langle\perp|$ vs. $\mathbb{1} - |\perp\rangle\langle\perp|$. A scheme Π is measured-BU, or mBU, secure, if for all $\varepsilon > 0$ and all QPT adversaries \mathcal{A} , the winning probability in the BU game when provided with a measured ε -blinded oracle instead of a ε -blinded oracle, is negligible.*

A straightforward reduction argument shows that BU implies mBU.

Proposition 1. *Let Π be a BU ($k - \text{BU}$)-secure MAC. Then Π is mBU ($k - \text{mBU}$)-secure.*

Proof. Let \mathcal{A} be an mBU-adversary against Π . We construct a BU-adversary \mathcal{A}' against Π as follows. \mathcal{A}' runs \mathcal{A} . For each query that \mathcal{A} makes to the measured ε -blinded oracle, \mathcal{A}' queries the ε -blinded oracle and performs the “blinded or not” measurement before returning the answer to \mathcal{A} . Clearly the probabilities for \mathcal{A}' winning the BU and for \mathcal{A} winning the mBU game are the same. \square

For the following proof we need a generalization of Zhandry’s superposition oracle technique to functions drawn from a non-uniform distribution. Such has been developed in detail in [11]. As for the proof of Theorem 21, we do not need the more complicated (but efficiently implementable) compressed oracle. Hence we introduce only the basic non-uniform superposition oracle. The generalization is straight-forward. In [30], a uniformly random function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is sampled by preparing 2^n m -qubit uniform superposition states. The measurement that performs the actual sampling is delayed, which allows for new ways of analyzing the behavior of a query algorithm by inspecting the oracle registers. Here, we would like to use the superposition oracle representation for the indicator function $\mathbb{1}_{B_\epsilon} : \{0, 1\}^n \rightarrow \{0, 1\}$ of the blinding set B_ϵ . This is a Boolean function with $\Pr[\mathbb{1}_{B_\epsilon}(x) = 1] = \epsilon$ independently for all $x \in \{0, 1\}^n$.

We will sample $\mathbb{1}_{B_\epsilon}$ by preparing 2^n qubits in the state

$$|\eta_0^\epsilon\rangle = \sqrt{1 - \epsilon}|0\rangle + \sqrt{\epsilon}|1\rangle, \tag{3}$$

i.e., we prepare the 2^n -qubit oracle register F in the state

$$\left(|\eta_0^\epsilon\rangle^{\otimes 2^n}\right)_F = \bigotimes_{x \in \{0,1\}^n} |\eta_0^\epsilon\rangle_{F_x}. \tag{4}$$

We will refrain from fourier-transforming any registers, so if the adversaries query registers are X and B (the input register and the blinding bit register), the oracle unitary is just given by

$$U_{\text{StO}} = \sum_{x \in \{0,1\}^n} |x\rangle\langle x|_X \otimes \text{CNOT}_{F_x:B}. \tag{5}$$

We can also define the generalization of the projectors P_ℓ . To this end we complete $|\eta_0^\epsilon\rangle$ to an orthonormal basis by introducing the state

$$|\eta_1^\epsilon\rangle = \sqrt{\epsilon}|0\rangle - \sqrt{1 - \epsilon}|1\rangle. \tag{6}$$

Let further U_ϵ be the unitary such that $U_\epsilon|i\rangle = |\eta_i^\epsilon\rangle$. The generalization of P_ℓ is now defined by $P_\ell^\epsilon = U_\epsilon P_\ell U_\epsilon^\dagger$. As U_{StO} is a sum of terms that each act non-trivially only on one out of the 2^n F_x registers, the analogue of Lemma 2 clearly holds, i.e., if $|\psi_q\rangle$ is the joint algorithm-oracle state after q queries to the superposition oracle for $\mathbb{1}_{B_\epsilon}$, then $P_\ell^\epsilon|\psi_q\rangle = 0$ for all $\ell > q$.

We are now ready to prove that BU security implies PO security.

Theorem 13. *Let Π be a BU-secure MAC. Then Π is PO-secure.*

Proof. According to Proposition 1, Π is mBU secure. It therefore suffices to find a reduction from breaking mBU to breaking PO. Let \mathcal{A} be a q query PO adversary against Π , i.e., an algorithm that makes q queries and outputs $q + 1$ pairs (x_i, t_i) with the goal that $t_i = \text{Mac}_k(x_i)$ for all $i = 1, \dots, q + 1$. We construct an mBU-adversary \mathcal{A}' as follows. The adversary \mathcal{A}' runs \mathcal{A} , answering the queries using the measured ϵ -blinded oracle for Mac_k . If for any of the queries the result is \perp ,

\mathcal{A}' aborts. In this case we formally define \mathcal{A}' 's output to be \perp . After (and if) \mathcal{A} has finished by outputting $q + 1$ candidate message-tag pairs (m_i, t_i) , \mathcal{A}' chooses $i \in_R \{1, \dots, q + 1\}$ and outputs (m_i, t_i) .

According to Theorem 2, the trace distance between the distribution of the $q + 1$ candidate message tag pairs that \mathcal{A} outputs only changes by $\delta = 2q\sqrt{\epsilon}$ in total variational distance when run with the measured ϵ -blinded oracle as done as a subroutine of \mathcal{A}' . It follows that with at least probability $p_{\text{succ}}^{\mathcal{A}} - \delta$, all $q + 1$ outputs of \mathcal{A} are valid message-tag-pairs, where $p_{\text{succ}}^{\mathcal{A}}$ is the probability with which \mathcal{A} wins the PO game when provided with an unblinded Mac_k -oracle.

For the rest of the proof we instantiate the blinding set using the superposition oracle described above. In this case, the measured ϵ -blinded oracle is implemented as follows. On input registers X and Y , create a blank qubit register B and query the blinding function $\mathbf{1}_{B_\epsilon}$ on XB . Measure B to obtain b (the blinding bit). If $b = 1$, query the Mac_k -oracle on XY , otherwise add \perp to Y . For the q -query algorithm \mathcal{A}' , q queries are made to the superposition blinding oracle. Afterwards the oracle register F is measured in the computational basis to determine whether the output is blinded or not.

We continue by finding a lower bound on the probability that the message output by \mathcal{A}' is blinded. To that end, consider the modified game, where after \mathcal{A}' has finished, but before measuring the oracle register F , we compute the smallest index $i \in \{1, \dots, q + 1\}$ such that F_{x_i} is in state $|\eta_0^{(\epsilon)}\rangle$ in superposition into an additional register. Such an index always exists. This is because $P_\ell^\epsilon|\psi\rangle = 0$ for all $\ell > q$, where $|\psi\rangle$ is the joint adversary-oracle state after the execution of \mathcal{A}' . Hence $|\psi\rangle$ is a superposition of states $|\beta\rangle = \bigotimes_{x \in \{0,1\}^n} |\eta_{\beta_x}^\epsilon\rangle$ for strings $\beta \in \{0,1\}^{2^n}$ of Hamming weight at most q . Now we measure the register to obtain an outcome i_0 . But given outcome i_0 , the register $F_{m_{i_0}}$ is in state $|\eta_0^\epsilon\rangle$. Now the oracle register is measured to determine the blinding set $B_\epsilon \subset \{0,1\}^n$. The computation together with the measurement implements a $(q + 1)$ -outcome projective measurement on F . The probability that m_{i_0} is blinded is ϵ independently, so the success probability in the modified game is

$$\tilde{p}_{\text{succ}}^{\mathcal{A}'} \geq \frac{\epsilon (p_{\text{succ}}^{\mathcal{A}} - 2q\sqrt{\epsilon})}{q + 1}. \tag{7}$$

Finally, we can apply Lemma 1 to conclude that adding the measurement has not decreased the success probability by more than a factor $1/(q + 1)$, to conclude that the success probability of \mathcal{A}' in the unmodified mBU game is lower-bounded by

$$p_{\text{succ}}^{\mathcal{A}'} \geq \frac{\epsilon (p_{\text{succ}}^{\mathcal{A}} - 2q\sqrt{\epsilon})}{(q + 1)^2}. \tag{8}$$

Choosing $\epsilon = (p_{\text{succ}}^{\mathcal{A}}/3q)^2$ we obtain

$$p_{\text{succ}}^{\mathcal{A}'} \geq \frac{(p_{\text{succ}}^{\mathcal{A}})^3}{27q^2(q+1)^2}. \tag{9}$$

In particular we have that $p_{\text{succ}}^{\mathcal{A}'}$ is non-negligible if p_{succ} was non-negligible. \square

1-BU also implies the notion by Garg et al. [14], see the full version [4].

In the purely classical setting, our notion is equivalent to EUF-CMA. Also, sBU from Sect. 4 implies SUF-CMA.

Proposition 2. *A MAC is EUF-CMA if and only if it is blind-unforgeable against classical adversaries.*

Proof. Set $F_k = \text{Mac}_k$. Consider an adversary \mathcal{A} which violates EUF-CMA. Such an adversary, given 1^n and oracle access to F_k (for $k \in_R \{0, 1\}^n$), produces a forgery (m, t) with non-negligible probability $s(n)$; in particular, $|m| \geq n$ and m is not among the messages queried by \mathcal{A} . This same adversary (when coupled with an appropriate ϵ) breaks the system under the blind-forgery definition. Specifically, let $p(n)$ be the running time of \mathcal{A} , in which case \mathcal{A} clearly makes no more than $p(n)$ queries, and define $\epsilon(n) = 1/p(n)$. Consider now a particular $k \in \{0, 1\}^n$ and a particular sequence r of random coins for $\mathcal{A}^{F_k}(1^n)$. If this run of \mathcal{A} results in a forgery (m, t) , observe that with probability at least $(1 - \epsilon)^{p(n)} \approx e^{-1}$ in the choice of B_ϵ , we have $F_k(q) = B_\epsilon F_k(q)$ for every query q made by \mathcal{A} . On the other hand, $B_\epsilon(m) = \perp$ with (independent) probability ϵ . It follows $\phi(n, \epsilon_n)$ is at least $\epsilon s(n)/e = \Omega(s(n)/p(n))$.

On the the other hand, suppose that (\mathcal{A}, ϵ) is an adversary that breaks blind-unforgeability. Consider now the EUF-CMA adversary $\mathcal{A}'^{F_k}(1^n)$ which simulates the adversary $\mathcal{A}^{(\cdot)}(1^n)$ by answering oracle queries according to a locally-simulated version of $B_\epsilon F_k$; specifically, the adversary \mathcal{A}' proceeds by drawing a subset $B_{\epsilon(n)} \subseteq \{0, 1\}^*$ as described above and answering queries made by \mathcal{A} according to $B_\epsilon F$. Two remarks are in order:

- When $x \in B_\epsilon$, this query is answered without an oracle call to $F(x)$.
- \mathcal{A}' can construct the set B_ϵ “on the fly,” by determining, when a particular query q is made by \mathcal{A} , whether $q \in B_\epsilon$ and “remembering” this information in case the query is asked again (“lazy sampling”).

With probability $\phi(n, \epsilon(n))$ \mathcal{A} produces a forgery on a point which was not queried by \mathcal{A}' , as desired. It follows that \mathcal{A} produces a (conventional) forgery with non-negligible probability when given F_k for $k \in_R \{0, 1\}^n$. \square

6 Blind-Unforgeable Schemes

Random schemes. We now show that suitable random and pseudorandom function families satisfy our notion of unforgeability.

Theorem 14. *Let $R : X \rightarrow Y$ be a uniformly random function such that $1/|Y|$ is negligible in n . Then R is a blind-forgery secure MAC.*

Proof. For simplicity, we assume that the function is length-preserving; the proof generalizes easily. Let \mathcal{A} be an efficient quantum adversary. The oracle $B_\epsilon R$ supplied to \mathcal{A} during the blind-forgery game is determined entirely by B_ϵ and the restriction of R to the complement of B_ϵ . On the other hand, the forgery event

$$\mathcal{A}^{B_\epsilon F_k}(1^n) = (m, t) \wedge |m| \geq n \wedge F_k(m) = t \wedge B_\epsilon F_k(m) = \perp$$

depends additionally on values of R at points in B_ϵ . To reflect this decomposition, given R and B_ϵ define $R_\epsilon : B_\epsilon \rightarrow Y$ to be the restriction of R to the set B_ϵ and note that—conditioned on $B_\epsilon R$ and B_ϵ —the random variable R_ϵ is drawn uniformly from the space of all (length-preserving) functions from B_ϵ into Y . Note, also, that for every n the purported forgery $(m, t) \leftarrow \mathcal{A}^{B_\epsilon R}(1^n)$ is a (classical) random variable depending only on $B_\epsilon R$. In particular, conditioned on B_ϵ , (m, t) is independent of R_ϵ . It follows that, conditioned on $m \in B_\epsilon$, that $t = R_\epsilon(m)$ with probability no more than $1/2^n$ and hence $\phi(n, \epsilon) \leq 2^{-n}$, as desired. \square

Next, we show that a qPRF is a blind-unforgeable MAC.

Corollary 3. *Let m and t be $\text{poly}(n)$, and $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^t$ a qPRF. Then F is a blind-forgery-secure fixed-length MAC (with length $m(n)$).*

Proof. For a contradiction, let \mathcal{A} be a QPT which wins the blind forgery game for a certain blinding factor $\epsilon(n)$, with running time $q(n)$ success probability $\delta(n)$. We will use \mathcal{A} to build a quantum oracle distinguisher \mathcal{D} between the qPRF F and the perfectly random function family \mathcal{F}_m^t with the same domain and range.

First, let $k = q(n)$ and let \mathcal{H} be a family of $(4k + 1)$ -wise independent functions with domain $\{0, 1\}^m$ and range $\{0, 1, \dots, 1/\epsilon(n)\}$. The distinguisher \mathcal{D} first samples $h \in_R \mathcal{H}$. Set $B_h := h^{-1}(0)$. Given its oracle \mathcal{O}_f , \mathcal{D} can implement the function $B_h f$ (quantumly) as follows:

$$\begin{aligned} |x\rangle|y\rangle &\mapsto |x\rangle|y\rangle|H_x\rangle|\delta_{h(x),0}\rangle \mapsto |x\rangle|y\rangle|H_x\rangle|\delta_{h(x),0}\rangle|f(x)\rangle \\ &\mapsto |x\rangle|y \oplus f(x) \cdot (1 - \delta_{h(x),0})\rangle|H_x\rangle|\delta_{h(x),0}\rangle|f(x)\rangle \\ &\mapsto |x\rangle|y \oplus f(x) \cdot (1 - \delta_{h(x),0})\rangle. \end{aligned}$$

Here we used the CCNOT (Toffoli) gate from step 2 to 3 (with one control bit reversed), and uncomputed both h and f in the last step. After sampling h , the distinguisher \mathcal{D} will execute \mathcal{A} with the oracle $B_h f$. If \mathcal{A} successfully forges a tag for a message in B_h , \mathcal{A}' outputs “pseudorandom”; otherwise “random.”

Note that the function $B_h f$ is perfectly ϵ -blinded if h is a perfectly random function. Note also that the entire security experiment with \mathcal{A} (including the final check to determine if the output forgery is blind) makes at most $2k$ quantum queries and 1 classical query to h , and is thus (by Theorem 10) identically distributed to the perfect-blinding case.

Finally, by Theorem 14, the probability that \mathcal{D} outputs “pseudorandom” when $f \in_R \mathcal{F}_m^t$ is negligible. By our initial assumption about \mathcal{A} , the probability that \mathcal{D} outputs “pseudorandom” becomes $\delta(n)$ when $f \in_R F$. It follows that \mathcal{D} distinguishes F from perfectly random. \square

Next, we give a information-theoretically secure q -time MACs (Definition 6).

Theorem 15. *Let \mathcal{H} be a $(4q + 1)$ -wise independent function family with range Y , such that $1/|Y|$ is a negligible function. Then \mathcal{H} is a q -time BU-secure MAC.*

Proof. Let (\mathcal{A}, ϵ) be an adversary for the q -time game $\text{BlindForge}_{\mathcal{A},h}^q(n, \epsilon(n))$, where h is drawn from \mathcal{H} . We will use \mathcal{A} to construct a distinguisher \mathcal{D} between \mathcal{H} and a random oracle. Given access to an oracle \mathcal{O} , \mathcal{D} first runs \mathcal{A} with the blinded oracle $B\mathcal{O}$, where the blinding operation is performed as in the proof of Corollary 3 (i.e., via a $(4q + 1)$ -wise independent function with domain size $1/\epsilon(n)$). When \mathcal{A} is completed, it outputs (m, σ) . Next, \mathcal{D} queries \mathcal{O} on the message m and outputs 1 if and only if $\mathcal{O}(m) = \sigma$ and $m \in B$. Let $\gamma_{\mathcal{O}}$ be the probability of the output being 1.

We consider two cases: (i) \mathcal{O} is drawn as a random oracle R , and (ii) \mathcal{O} is drawn from the family \mathcal{H} . By Theorem 10, since \mathcal{D} makes only $2q$ quantum queries and one classical query to \mathcal{O} , its output is identical in the two cases. Observe that γ_R (respectively, $\gamma_{\mathcal{H}}$) is exactly the success probability of \mathcal{A} in the blind-forgery game with random oracle R (respectively, \mathcal{H}). We know from Theorem 14 that γ_R is negligible; it follows that $\gamma_{\mathcal{H}}$ is as well. \square

Several domain-extension schemes, including NMAC (a.k.a. encrypted cascade), HMAC, and AMAC, can transform a fixed-length qPRF to a qPRF that takes variable-length inputs [24]. As a corollary, starting from a qPRF, we also obtain a number of quantum blind-unforgeable variable-length MACs.

Lamport one-time signatures. The Lamport signature scheme [19] is a EUF-1-CMA-secure signature scheme, specified as follows.

Construction 16 (Lamport signature scheme, [19]). *For the Lamport signature scheme using a hash function family $h : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, the algorithms KeyGen , Sign and Ver are specified as follows. KeyGen , on input 1^n , outputs a pair (pk, sk) with*

$$\text{sk} = (s_i^j)_{i \in \{1, \dots, n\}, j=0,1}, \text{ with } s_i^j \in_R \{0, 1\}^n, \text{ and} \tag{10}$$

$$\text{pk} = \left(k, \left(p_i^j \right)_{i \in \{1, \dots, n\}, j=0,1} \right), \text{ with } k \in \{0, 1\}^n \text{ and } p_i^j = h_k \left(s_i^j \right). \tag{11}$$

The signing algorithm is defined by $\text{Sign}_{\text{sk}}(x) = (s_i^{x_i})_{i \in \{1, \dots, n\}}$ where $x_i, i = 1, \dots, n$ are the bits of x . The verification procedure checks the signature's consistency with the public key, i.e., $\text{Ver}_{\text{pk}}(x, s) = 0$ if $p_i^{x_i} = h_k(s_i)$ and $\text{Ver}_{\text{pk}}(x, s) = 0$ else.

We now show that the Lamport scheme is 1-BU secure in the quantum random oracle model.

Theorem 17. *If in Construction 16, h is modeled as a quantum-accessible random oracle, it is 1-BU secure.*

We give a brief sketch of the proof; for details, see Appendix A. The proof uses arguments analogous to the classical proof. This is made possible through the use of the Fourier oracle technique (from Sect. 3) for both h and sk . The latter can be understood as a uniformly random function $\text{sk} : \{0, \dots, n - 1\} \times \{0, 1\} \rightarrow \{0, 1\}^n$. Subsequently we perform “forensics” on the oracle database after the adversary has finished, in a similar way as in the proof of Theorem 21. Let us first argue that an adversary \mathcal{A} that makes a signing query but no queries to h and outputs (m, σ) does not succeed except with negligible probability. If m is blinded, then there is at least one bit of m where the corresponding part of sk (and hence the correct signature) is independent of σ . While this is only true in superposition, we can break this superposition using an n -outcome measurement on the sk -register, which does not change the success probability by much according to Lemma 1.

For the general case, we observe that queries to h do not help, because they will only have negligible support on the unqueried parts of sk . Concretely, we show that the commutator of the oracle unitary for h and the projector on the uniform superposition state (the initial state of the oracle register holding a part of sk) is small in operator norm, which implies that an untouched sk register remains untouched except with negligible amplitude, even in superposition.

A simple proof of the PO-security of a random function can be given using a similar idea; see the full version [4].

Hash-and-MAC. To authenticate messages of arbitrary length with a fixed-length MAC, it is common practice to first compress a long message by a *collision-resistant* hash function and then apply the MAC. This is known as Hash-and-MAC. However, when it comes to BU-security, collision-resistance may not be sufficient. We therefore propose a new notion, Bernoulli-preserving hash, generalizing collision-resistance in the quantum setting, and show that it is sufficient for Hash-and-MAC with BU security. Recall that, given a subset B of a set X , $\chi_B : X \rightarrow \{0, 1\}$ denotes the characteristic function of B .

Definition 8 (Bernoulli-preserving hash). *Let $\mathcal{H} : X \rightarrow Y$ be an efficiently computable function family. Define the following distributions on subsets of X :*

1. \mathcal{B}_ϵ : generate $B_\epsilon \subseteq X$ by placing $x \in B_\epsilon$ independently with probability ϵ . Output B_ϵ .
2. $\mathcal{B}_\epsilon^{\mathcal{H}}$: generate $C_\epsilon \subseteq Y$ by placing $y \in C_\epsilon$ independently with probability ϵ . Sample $h \in \mathcal{H}$ and define $B_\epsilon^h := \{x \in X : h(x) \in C_\epsilon\}$. Output B_ϵ^h .

We say that \mathcal{H} is a Bernoulli-preserving hash if for all adversaries (\mathcal{A}, ϵ) ,

$$\left| \Pr_{B \leftarrow \mathcal{B}_\epsilon} [\mathcal{A}^{X^B}(1^n) = 1] - \Pr_{B \leftarrow \mathcal{B}_\epsilon^{\mathcal{H}}} [\mathcal{A}^{X^B}(1^n) = 1] \right| \leq \text{negl}(n).$$

The motivation for the name Bernoulli-preserving hash is simply that selecting \mathcal{B}_ϵ can be viewed as a Bernoulli process taking place on the set X , while \mathcal{B}_ϵ^h can be viewed as the pullback (along h) of a Bernoulli process taking place on Y .

We show that the standard, so-called “Hash-and-MAC” construction will work w.r.t. to BU security, if we instantiate the hash function with a Bernoulli-preserving hash. Recall that, given a MAC $\Pi = (\text{Mac}_k, \text{Ver}_k)$ with message set X and a function $h : Z \rightarrow X$, there is a MAC $\Pi^h := (\text{Mac}_k^h, \text{Ver}_k^h)$ with message set Z defined by $\text{Mac}_k^h = \text{Mac}_k \circ h$ and $\text{Ver}_k^h(m, t) = \text{Ver}_k(h(m), t)$.

Theorem 18 (Hash-and-MAC with Bernoulli-preserving hash).

Let $\Pi = (\text{Mac}_k, \text{Ver}_k)$ be a BU-secure MAC with $\text{Mac}_k : X \rightarrow Y$, and let $h : Z \rightarrow X$ a Bernoulli-preserving hash. Then Π^h is a BU-secure MAC.

The proof follows in a straightforward way from the definitions of BU and Bernoulli-preserving hash; the details are in the full version [4].

In Appendix B, we also provide a number of additional results about Bernoulli-preserving hash functions. These results can be summarized as follows.

Theorem 19. *We prove the following about Bernoulli-preserving hash functions.*

- If H is a random oracle or a qPRF, then it is a Bernoulli-preserving hash.
- If H is $4q$ -wise independent, then it is a Bernoulli-preserving hash against q -query adversaries.
- Under the LWE assumption, there is a (public-key) family of Bernoulli-preserving hash functions.
- If we only allow classical oracle access, then the Bernoulli-preserving property is equivalent to standard collision-resistance.
- Bernoulli-preserving hash functions are collapsing (another quantum generalization of collision-resistance proposed in [27]).

7 The Problem with PO-Unforgeability

Our search for a new definition of unforgeability for quantum-secure authentication is partly motivated by concerns about the PO security notion [7]. In this section, we make these concerns concrete by pointing out a significant security concern not addressed by this definition. Specifically, we demonstrate a MAC which is readily broken with an efficient attack, and yet is PO secure. The attack queries the MAC with a superposition over a particular subset S of the message space, and then forges a valid tag for a message lying outside S .

One of the intuitive issues with PO is that it might rule out adversaries that have to measure, and thereby destroy, one or more post-query states to produce an interesting forgery. Constructing such an example seems not difficult at first. For instance, let us look at one-time PO, and construct a MAC from a qPRF f by sampling a key k for f and a superpolynomially-large prime p , and setting

$$\text{Mac}_{k,p}(m) = \begin{cases} 0^n & \text{if } m = p, \\ (f_k(m \bmod p)) & \text{otherwise.} \end{cases} \tag{12}$$

This MAC is forgeable: a quantum adversary can use a single query to perform period-finding on the MAC, and then forge at 0^n . Intuitively, it seems plausible

that the MAC is 1-PO secure as period-finding uses a full measurement. This is incorrect for a somewhat subtle reason: identifying the hidden symmetry does not necessarily consume the post-query state completely, so an adversary can learn the period and a random input-output-pair of the MAC simultaneously. As shown in the full version [4] this is a special case of a fairly general situation, which makes establishing a proper PO “counterexample” difficult.

A counterexample to PO. Another intuitive problem with PO is that using the contents of a register can necessitate *uncomputing* the contents of another one. We exploit this insufficiency in the counterexample below. Consider the following MAC construction.

Construction 20. Given $k = (p, f, g, h)$ where $p \in \{0, 1\}^n$ is a random period and $f, g, h : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are random functions, define $M_k : \{0, 1\}^{n+1} \rightarrow \{0, 1\}^{2n}$ by

$$M_k(x) = \begin{cases} g(x' \bmod p) \| f(x') & x = 1 \| x', \\ 0^n \| h(x') & x = 0 \| x', \ x' \neq p, \\ 0^{2n} & x = 0 \| p. \end{cases}$$

Consider an adversary that queries as follows

$$\sum_{x,y} |1, x\rangle_X |0^n\rangle_{Y_1} |y\rangle_{Y_2} \mapsto \sum_{x,y} |1, x\rangle_X |g_p(x)\rangle_{Y_1} |y \oplus f(x)\rangle_{Y_2}, \tag{13}$$

and then discards the first qubit and the Y_2 register; this yields $\sum_x |x\rangle |g_p(x)\rangle$. The adversary can extract p via period-finding from polynomially-many such states, and then output $(0 \| p, 0^{2n})$. This attack only queries the MAC on messages starting with 1 (e.g., “from Alice”), and then forges at a message which starts with 0 (e.g., “from Bob.”) We emphasize that the forgery was never queried, not even with negligible amplitude. It is thus intuitively clear that this MAC does not provide secure authentication. And yet, despite this obvious and intuitive vulnerability, this MAC is in fact PO-secure.

Theorem 21. *The MAC from Construction 20 is PO-secure.*

The proof of this theorem can be found in the full version [4]. The proof idea is as follows. The superposition oracle technique outlined in Sect. 3 achieves something that naively seems impossible due to the quantum no-cloning theorem: it records on which inputs the adversary has made non-trivial⁴ queries. The information recorded in this way cannot, in general be utilized in its entirety – after all, the premise of the superposition oracle is that the measurement \mathcal{M}_F that samples the random function is delayed until after the algorithm has finished, but it still has to be performed. Any measurement \mathcal{M}' that does not commute

⁴ For the standard unitary oracle for a classical function, a query has no effect when the output register is initialized in the uniform superposition of all strings.

with \mathcal{M}_F and is performed before \mathcal{M}_F , can disturb the outcome of \mathcal{M}_F . If however, \mathcal{M}' only has polynomially many possible outcomes, that disturbance is at most inverse polynomial according to Lemma 1.

Here, we sample the random function f using a superposition oracle, and we chose to use a measurement \mathcal{M}' to determine the *number* of nontrivial queries that the adversary has made to f , which is polynomial by assumption. Random functions are PO-secure [7], so the only way to break PO security is to output $(0\|p, 0^{2n})$ and q other input-output-pairs. Querying messages that start with 0 clearly only yields a negligible advantage in guessing p by the Grover lower bound, so we consider an adversary querying only on strings starting with 1. We distinguish two cases, either the adversary makes or exactly q non-trivial queries to f , or less than that. In the latter case, the success probability is negligible by the PO-security of f and h . In the former case, we have to analyze the probability that the adversary guesses p correctly. f is not needed for that, so the superposition oracle register can be used to measure the set of q queries that the adversary made. Using an inductive argument reminiscent of the hybrid method [5] we show that this set is almost independent of p , and hence the period is equal to the difference of two of the queried inputs only with negligible probability. But if that is not the case, the periodic version of g is indistinguishable from a random function for that adversary which is independent of p .

It's not hard to see that the MAC from Construction 20 is not GYZ-secure. Indeed, observe that the forging adversary described above queries on messages starting with 0 only, and then forges successfully on a message starting with 1. If the scheme was GYZ secure, then in the accepting case, the portion of this adversary between the query and the final output would have a simulator which leaves the computational basis invariant. Such a simulator cannot change the first bit of the message from 0 to 1, a contradiction.

By Theorem 11, this PO-secure MAC is also not BU-secure.

Corollary 4. *The MAC from Construction 20 is BU-insecure.*

Acknowledgements. CM thanks Ronald de Wolf for helpful discussions on query complexity. GA acknowledges support from NSF grant CCF-1763736. CM was funded by a NWO VIDI grant (Project No. 639.022.519) and a NWO VENI grant (Project No. VI.Veni.192.159). FS acknowledges support from NSF grant CCF-1901624. AR acknowledges support from NSF grant CCF-1763773.

A Technical Proofs

The Fourier Oracle number operator. We now restate and prove Lemma 2.

Lemma 3. *The number operator satisfies $\| [N_F, U_{XYF}^{FO}] \|_\infty = 1$. In particular, the joint state of a quantum query algorithm and the oracle after the q -th query is in the kernel of P_l for all $l > q$.*

Proof. Let $|\psi\rangle_{XYEF}$ be an arbitrary query state, where X and Y are the query input and output registers, E is the algorithm's internal register and F is the FO register. We expand the state in the computational basis of X ,

$$|\psi\rangle_{XYEF} = \sum_{x \in \{0,1\}^n} p(x) |x\rangle_X |\psi_x\rangle_{YEF}. \quad (14)$$

Now observe that $U_{XYF}^{\text{FO}} |x\rangle_X |\psi_x\rangle_{YEF} = |x\rangle_X \left(\widetilde{\text{CNOT}}^{\otimes m} \right)_{Y:F_x} |\psi_x\rangle_{YEF}$ with $\widetilde{\text{CNOT}}_{A:B} = H_A \text{CNOT}_{A:B} H_A$, and therefore

$$\begin{aligned} [N_F, U_{XYF}] |x\rangle_X |\psi_x\rangle_{YEF} &= |x\rangle_X \left[N_F, \left(\widetilde{\text{CNOT}}^{\otimes m} \right)_{Y:F_x} \right] |\psi_x\rangle_{YEF} \\ &= |x\rangle_X \left[(\mathbb{1} - |0\rangle\langle 0|)_{F_x}, \left(\widetilde{\text{CNOT}}^{\otimes m} \right)_{Y:F_x} \right] |\psi_x\rangle_{YEF}. \end{aligned}$$

It follows that

$$\begin{aligned} &\left\| [N_F, U_{XYF}] |\psi\rangle_{XYEF} \right\|_2 \quad (15) \\ &= \sum_{x \in \{0,1\}^n} p(x) \left\| [N_F, U_{XYF}] |\psi_x\rangle_{YEF} \right\|_2 \\ &= \sum_{x \in \{0,1\}^n} p(x) \left\| \left[(\mathbb{1} - |0\rangle\langle 0|)_{F_x}, \left(\widetilde{\text{CNOT}}^{\otimes m} \right)_{Y:F_x} \right] |\psi_x\rangle_{YEF} \right\|_2 \\ &\leq \left\| \left[(\mathbb{1} - |0\rangle\langle 0|)_{F_0^n}, \left(\widetilde{\text{CNOT}}^{\otimes m} \right)_{Y:F_0^n} \right] \right\|_{\infty}, \quad (16) \end{aligned}$$

where we have used the definition of the operator norm and the normalization of $|\psi\rangle_{XYEF}$ in the last line. For a unitary U and a projector P , it is easy to see that $\|[U, P]\|_{\infty} \leq 1$, as $[U, P] = PU(\mathbb{1} - P) - (\mathbb{1} - P)UP$ is a sum of two operators that have orthogonal support and singular values smaller or equal to 1. We therefore get $\|[N_F, U_{XYF}] |\psi\rangle_{XYEF}\|_2 \leq 1$, and as the state $|\psi\rangle$ was arbitrary, this implies $\|[N_F, U_{XYF}]\|_{\infty} \leq 1$. The example from Eq. (2) shows that the above is actually an equality. The observation that $P_l \eta_F = 0$ for all $l > 0$ and an induction argument proves the second statement of the lemma.

BU-security of Lamport. In this section, we provide the full proof of Theorem 17, showing that the Lamport construction is BU-secure in the QROM.

Proof. We implement the random oracle h as a superposition oracle with register F . In the 1-BlindForge experiment we execute the sampling part of the key generation by preparing a superposition as well. More precisely, we can just prepare $2n$ n -qubit registers S_i^j in a uniform superposition, with the intention of measuring

them to sample s_i^j in mind. We are talking about a classical one-time signature scheme, and all computation that uses the secret key is done by an honest party, and is therefore classical. It follows that the measurement that samples s_i^j commutes with all other operations which are implemented as quantum-controlled operations controlled on the secret key registers, i.e., we can postpone it to the very end of the 1-BlindForge experiment, just like the measurement that samples an actual random oracle using a superposition oracle. The joint state $|\psi_0\rangle$ with oracle register F and secret key register $SK = (S_i^j)_{i \in \{1, \dots, n\}, j=0,1}$ is now in a uniform superposition, i.e.,

$$|\psi_0\rangle_{SKF} = |\phi_0\rangle_{SK}^{\otimes 2n} \otimes |\phi_0\rangle_F^{\otimes 2^n}. \tag{17}$$

To subsequently generate the public key, the superposition oracle for h is queried on each of the S_i^j with an empty output register P_i^j , producing the state $|\psi_1\rangle_{SKPKF}$ equal to

$$2^{-2n^2} \sum_{\substack{s_i^j \in \{0,1\}^n \\ p_i^j \in \{0,1\}^n \\ i \in \{1, \dots, n\}, j=0,1}} \left(\bigotimes_{i \in \{1, \dots, n\}} |s_i^j\rangle_{S_i^j} \right) \otimes \left(\bigotimes_{i \in \{1, \dots, n\}} |p_i^j\rangle_{P_i^j} \right) \otimes |f_{sk,pk}\rangle_F,$$

where $|f_{sk,pk}\rangle_F$ is the superposition oracle state where $F_{s_i^j}$ is in state $|p_i^j\rangle$ and all other registers are still in state $|\phi_0\rangle$. Then the registers P_i^j are measured to produce an actual, classical, public key that can be handed to the adversary. Note that there is no hash function key k now, as it has been replaced by the random oracle. Treating the public key as classical information from now on and removing the registers PK , the state takes the form

$$|\psi_2(pk)\rangle_{SKF} = 2^{-n^2} \sum_{\substack{s_i^j \in \{0,1\}^n \\ i \in \{1, \dots, n\}, j=0,1}} \left(\bigotimes_{i \in \{1, \dots, n\}} |s_i^j\rangle_{S_i^j} \right) \otimes |f_{sk,pk}\rangle_F, \tag{18}$$

Now the interactive phase of the 1-BlindForge experiment can begin, and we provide both the random oracle h and the signing oracle (that can be called exactly once) as superposition oracles using the joint oracle state $|\psi_2(pk)\rangle$ above. The random oracle answers queries as described in Sect. 3. The signing oracle, when queried with registers XZ with $Z = Z_1 \dots Z_n$, applies $\text{CNOT}_{S_i^{x_i}:Z_i}^{\otimes n}$, $i = 1, \dots, n$ controlled on X being in the state $x \notin B_\epsilon$.

Now suppose \mathcal{A} , after making at most one query to Sign and an arbitrary polynomial number of queries to h , outputs a candidate message signature pair (x^0, z^0) with $z^0 = z_1^0 \parallel \dots \parallel z_n^0$. If $x^0 \notin B_\epsilon$, \mathcal{A} has lost. Suppose therefore that $x^0 \in B_\epsilon$. We will now make a measurement on the oracle register to find an index i such that $S_i^{x_i^0}$ has not been queried. To this end we first need to decorrelate SK and F . This is easily done, as the success test only needs computational basis

measurement results from the register SK , allowing us to perform any controlled operation on F controlled on SK . Therefore we can apply the operation $\oplus p_i^j$ followed by $H^{\otimes n}$ to the register $F_{s_i^j}$ controlled on S_i^j being in state $|s_i^j\rangle$, for all $i = 1, \dots, n$ and $j = 0, 1$. For an adversary that does not make any queries to h , this has the effect that all F -registers are in state $|\phi_0\rangle$ again now.

We can equivalently perform this restoring procedure before the adversary starts interaction, and answer the adversary's h -queries as follows. Controlled on the adversary's input being equal to one of the parts s_i^j of the secret key, answer with the corresponding public key, otherwise use the superposition oracle for h .

For any fixed secret key register S_i^j , the unitary that is applied upon an h -query has hence the form

$$U'_h = U_\perp + \sum_{x \in \{0,1\}^n} (U_x - U_\perp) |x\rangle\langle x|_X |x\rangle\langle x|_{S_i^j} \tag{19}$$

$$= U_\perp + \sum_{x \in \{0,1\}^n} |x\rangle\langle x|_X |x\rangle\langle x|_{S_i^j} (U_x - U_\perp), \tag{20}$$

where the second equality follows because the unitaries U_\perp and U_x are controlled unitaries with X and S_i^j part of the control register. Using the above equation we derive a bound on the operator norm of the commutator of this unitary and the projector onto $|\phi_0\rangle$,

$$\begin{aligned} & \| [U'_h, |\phi_0\rangle\langle\phi_0|] \|_\infty \\ &= 2^{-n/2} \left\| \sum_{x \in \{0,1\}^n} \left((U_x - U_\perp) |x\rangle\langle x|_X |x\rangle\langle\phi_0|_{S_i^j} - |x\rangle\langle x|_X |\phi_0\rangle\langle x|_{S_i^j} (U_x - U_\perp) \right) \right\|_\infty \\ &= 2^{-n/2} \max_{x \in \{0,1\}^n} \left\| \left((U_x - U_\perp) |x\rangle\langle x|_X |x\rangle\langle\phi_0|_{S_i^j} - |x\rangle\langle x|_X |\phi_0\rangle\langle x|_{S_i^j} (U_x - U_\perp) \right) \right\|_\infty \\ &\leq 2 \cdot 2^{-n/2}, \end{aligned}$$

where the second equality follows again because U_\perp and U_x are controlled unitaries with X and S_i^j part of the control register.

It follows that a query to h does not decrease the number of registers S_i^j that are in state $|\phi_0\rangle$, except with probability $8n \cdot 2^{-n}$.

As we assume that x^0 is blinded, we have that for any message $x \notin B_\varepsilon$, there exists an $i \in \{1, \dots, n\}$ such that $x_i \neq x_i^0$. But \mathcal{A} interacts with a blinded signing oracle, i.e., controlled on his input being not blinded, it is forwarded to the signing oracle, otherwise \perp is XORed into his output register. Therefore only non-blinded queries have been forwarded to the actual signing oracle, so the final state is a superposition of states in which the register SK has at least n subregisters S_i^j are in state $|\phi_0\rangle$, and at least one of them is such that $x_i^0 = j$. We can therefore apply an n -outcome measurement to the oracle register to obtain an index i_0 such that $S_{i_0}^{x_{i_0}^0}$ is in state $|\phi_0\rangle$. By Lemma 1, this implies that \mathcal{A} 's forgery is independent of s_{i_0} , so \mathcal{A} 's probability of succeeding in BlindForge is negligible. \square

B More on Bernoulli-Preserving Hash

In this section, we prove several results about Bernoulli-preserving hash functions. Recalling Definition 8, we refer to blinding according to \mathcal{B}_ϵ as “uniform blinding,” and blinding according to $\mathcal{B}_\epsilon^{\mathcal{H}}$ as “hash blinding.” First, we show that random and pseudorandom functions are Bernoulli-preserving, and that this property is equivalent to collision-resistance against classical queries.

Lemma 4. *Let $H : X \rightarrow Y$ be a function such that $1/|Y|$ is negligible. Then*

1. *If H is a random oracle or a qPRF, then it is a Bernoulli-preserving hash.*
2. *If H is $4q$ -wise independent, then it is a Bernoulli-preserving hash against q -query adversaries.*

Proof. The claim for random oracles is obvious: by statistical collision-resistance, uniform blinding is statistically indistinguishable from hash-blinding. The remaining claims follow from the observation that one can simulate one quantum query to $\chi_{B_\epsilon^h}$ using two quantum queries to h (see, e.g., the proof of Corollary 3). \square

Theorem 22. *A function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ is Bernoulli-preserving against classical-query adversaries if and only if it is collision-resistant.*

Proof. First, the Bernoulli-preserving hash property implies collision-resistance: testing whether two colliding inputs are either (i) both not blinded or both blinded, or (ii) exactly one of them is blinded, yields always outcome (i) when dealing with a hash-blinded oracle and a uniformly random outcome for a blinded oracle and $\epsilon = 1/2$. On the other hand, consider an adversary \mathcal{A} that has inverse polynomial distinguishing advantage between blinding and hash-blinding, and let x_1, \dots, x_q be its queries. Assume for contradiction that with overwhelming probability $h(x_i) \neq h(x_j)$ for all $x_i \neq x_j$. Then with that same overwhelming probability the blinded and hash blinded oracles are both blinded independently with probability ϵ on each x_i and are hence statistically indistinguishable, a contradiction. It follows that with non-negligible probability there exist two queries $x_i \neq x_j$ such that $h(x_i) = h(x_j)$, i.e., \mathcal{A} has found a collision. \square

Bernoulli-preserving hash from LWE. We have observed that any qPRF is a Bernoulli-preserving hash function, which can be constructed from various quantum-safe computational assumption (e.g., LWE). Nonetheless, qPRF typically does not give short digest, which would result in long tags, and it requires a secret key.⁵

Here we point out an alternative construction of a public Bernoulli-preserving hash function based on the quantum security of LWE. In fact, we show that the collapsing hash function in [26] is also Bernoulli-preserving hash. This construction relies on a lossy function family $F : X \rightarrow Y$ and a universal hash function

⁵ In practice, it is probably more convenient (and more reliable) to instantiate a qPRF from block ciphers, which may not be ideal for message authentication.

$G = \{g_k : Y \rightarrow Z\}_{k \in \mathcal{K}}$. A lossy function family admits two types of keys: a lossy key $s \leftarrow \mathcal{D}_{\text{los}}$ and an injective key $s \leftarrow \mathcal{D}_{\text{inj}}$, which are computationally indistinguishable. $F_s : X \rightarrow Y$ under a lossy key s is compressing, i.e., $|\text{im}(F_s)| \ll |Y|$; whereas under an injective key s , F_s is injective. We refer a formal definition to [26, Definition 2], and an explicit construction based on LWE to [20]. There exist efficient constructions for universal hash families by various means [28]. Then one constructs a hash function family $H = \{h_{s,k}\}$ by $h_{s,k} := g_k \circ F_s$ with public parameters generated by $s \leftarrow \mathcal{D}_{\text{los}}, k \leftarrow \mathcal{K}$. The proof of Bernoulli-preserving for this hash function is similar to Unruh’s proof that H is collapsing; see the full version [4].

Relationship to collapsing. Finally, we relate Bernoulli-preserving hash to another quantum generalization of classical collision-resistance: the collapsing property. Collapsing hash functions are particularly relevant to post-quantum signatures. We first define the collapsing property (slightly rephrasing Unruh’s original definition [27]) as follows. Let $h : X \rightarrow Y$ be a hash function, and let \mathcal{S}_X and \mathcal{S}_{XY} be the set of quantum states (i.e., density operators) on registers corresponding to the sets X and $X \times Y$, respectively. We define two channels from \mathcal{S}_X to \mathcal{S}_{XY} . First, \mathcal{O}_h receives X , prepares $|0\rangle$ on Y , applies $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus h(x)\rangle$, and then measures Y fully in the computational basis. Second, \mathcal{O}'_h first applies \mathcal{O}_h and then also measures X fully in the computational basis.

$$\begin{aligned} \mathcal{O}_h : |x\rangle_X &\xrightarrow{h} |x, h(x)\rangle_{X,Y} \xrightarrow{\text{measure } Y} (\rho_{X,Y}^y, y), \\ \mathcal{O}'_h : |x\rangle_X &\xrightarrow{h} |x, h(x)\rangle_{X,Y} \xrightarrow{\text{measure } X \& Y} (x, y). \end{aligned}$$

If the input is a pure state on X , then the output is either a superposition over a fiber $h^{-1}(s) \times \{s\}$ of h (for \mathcal{O}_h) or a classical pair $(x, h(x))$ (for \mathcal{O}'_h).

Definition 9 (Collapsing). *A hash function h is collapsing if for any single-query QPT \mathcal{A} , it holds that $|\Pr[\mathcal{A}^{\mathcal{O}_h}(1^n) = 1] - \Pr[\mathcal{A}^{\mathcal{O}'_h}(1^n) = 1]| \leq \text{negl}(n)$.*

To prove that Bernoulli-preserving hash implies collapsing, we need a technical fact. Recall that any subset $S \subseteq \{0, 1\}^n$ is associated with a two-outcome projective measurement $\{\Pi_S, \mathbb{1} - \Pi_S\}$ on n qubits defined by $\Pi_S = \sum_{x \in S} |x\rangle\langle x|$. We will write Ξ_S for the channel (on n qubits) which applies this measurement.

Lemma 5. *Let S_1, S_2, \dots, S_{c_n} be subsets of $\{0, 1\}^n$, each of size 2^{n-1} , chosen independently and uniformly at random. Let Ξ_{S_j} denote the two-outcome measurement defined by S_j , and denote their composition $\tilde{\Xi} := \Xi_{S_{c_n}} \circ \Xi_{S_{c_n-1}} \circ \dots \circ \Xi_{S_1}$. Let Ξ denote the full measurement in the computational basis. Then $\Pr[\tilde{\Xi} = \Xi] \geq 1 - 2^{-\varepsilon n}$, whenever $c \geq 2 + \varepsilon$ with $\varepsilon > 0$,*

A proof is given in the full version [4]. We remark that, to efficiently implement each Ξ_S with a random subset S , we can sample $h_i : [M] \rightarrow [N]$ from a pairwise-independent hash family (sampling an independent h_i for each i), and then define $x \in S$ iff. $h(x) \leq N/2$. For any input state $\sum_{x,z} \alpha_{x,z} |x, z\rangle$, we can compute

$$\sum_{x,z} \alpha_{x,z} |x, z\rangle \mapsto \sum_{x,z} |x, z\rangle |b(x)\rangle, \quad \text{where } b(x) := h(x) \stackrel{?}{\leq} N/2,$$

and then measure $|b(x)\rangle$. Pairwise independence is sufficient by Theorem 10 because only one quantum query is made.

Theorem 23. *If $h : X \rightarrow Y$ is Bernoulli-preserving, then it is collapsing.*

Proof. Let \mathcal{A} be an adversary with inverse-polynomial distinguishing power in the collapsing game. Choose n such that $X = \{0, 1\}^n$. We define $k = cn$ hybrid oracles H_0, H_1, \dots, H_k , where hybrid H_j is a channel from \mathcal{S}_X to \mathcal{S}_{XY} which acts as follows: (1) adjoin $|0\rangle_Y$ and apply the unitary $|x\rangle_X |y\rangle_Y \mapsto |x\rangle_X |y \oplus h(x)\rangle_Y$; (2) measure the Y register in the computational basis; (3) repeat j times: (i) select a uniformly random subset $S \subseteq X$ of size 2^{n-1} ; (ii) apply the two-outcome measurement Ξ_S to the X register; (4) output registers X and Y .

Clearly, H_0 is identical to the \mathcal{O}_h channel in the collapsing game. By Lemma 5, H_k is indistinguishable from the \mathcal{O}'_h . By our initial assumption and the triangle inequality, there exists a j such that

$$|\Pr[\mathcal{A}^{H_j}(1^n) = 1] - \Pr[\mathcal{A}^{H_{j+1}}(1^n) = 1]| \geq 1/\text{poly}(n). \tag{21}$$

We now build a distinguisher \mathcal{D} against the Bernoulli-preserving property (with $\epsilon = 1/2$) of h . It proceeds as follows: (1) run $\mathcal{A}(1^n)$ and place its query state in register X ; (2) simulate oracle H_j on XY (use 2-wise independent hash to select sets S); (3) prepare an extra qubit in the $|0\rangle$ state in register W , and invoke the oracle for χ_B on registers X and W ; (4) measure and discard register W ; (5) return XY to \mathcal{A} , and output what it outputs.

We now analyze \mathcal{D} . After the first two steps of H_j (compute h , measure output register) the state of \mathcal{A} (running as a subroutine of \mathcal{D}) is given by

$$\sum_z \sum_{x \in h^{-1}(s)} \alpha_{xz} |x\rangle_X |s\rangle_Y |z\rangle_Z.$$

Here Z is a side information register private to \mathcal{A} . Applying the j measurements (third step of H_j) results in a state of the form $\sum_z \sum_{x \in M} \beta_{xz} |x\rangle |s\rangle |z\rangle$, where M is a subset of $h^{-1}(s)$. Applying the oracle for χ_B into an extra register now yields

$$\sum_z \sum_{x \in M} \beta_{xz} |x\rangle |s\rangle |z\rangle |\chi_B(x)\rangle_W.$$

Now consider the two cases of the Bernoulli-preserving game.

First, in the “hash-blinded” case, $B = h^{-1}(C)$ for some set $C \subseteq Y$. This implies that $\chi_B(x) = \chi_C(h(x)) = \chi_C(s)$ for all $x \in M$. It follows that W simply contains the classical bit $\chi_C(s)$; computing this bit, measuring it, and discarding it will thus have no effect. The state returned to \mathcal{A} will then be identical to the output of the oracle H_j . Second, in the “uniform blinding” case, B is a random subset of X of size 2^{n-1} , selected uniformly and independently of everything else in the algorithm thus far. Computing the characteristic function of B into an

extra qubit and then measuring and discarding that qubit implements the channel Ξ_B , i.e., the measurement $\{\Pi_B, \mathbb{1} - \Pi_B\}$. It follows that the state returned to \mathcal{A} will be identical to the output of oracle H_{j+1} .

By (21), it now follows that \mathcal{D} is a successful distinguisher in the Bernoulli-preserving hash game for h , and that h is thus not a Bernoulli-preserving hash. \square

References

1. Aaronson, S.: Quantum lower bound for recursive Fourier sampling. *Quantum Inf. Comput.* **3**(2), 165–174 (2003)
2. Alagic, G., Gagliardini, T., Majenz, C.: Can you sign a quantum state? arXiv preprint [arXiv:1811.11858](https://arxiv.org/abs/1811.11858) (2018)
3. Alagic, G., Gagliardini, T., Majenz, C.: Unforgeable quantum encryption. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10822, pp. 489–519. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_16
4. Alagic, G., Majenz, C., Russell, A., Song, F.: Quantum-secure message authentication via blind-unforgeability. arXiv preprint [arXiv:1803.03761](https://arxiv.org/abs/1803.03761) (2020)
5. Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.V.: Strengths and weaknesses of quantum computing. *SIAM J. Comput.* **26**(5), 1510–1523 (1997)
6. Biasse, J.-F., Song, F.: Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2016, Philadelphia, PA, USA, pp. 893–902. Society for Industrial and Applied Mathematics (2016)
7. Boneh, D., Zhandry, M.: Quantum-secure message authentication codes. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 592–608. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_35
8. Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 361–379. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_21
9. Chen, L., et al.: Report on post-quantum cryptography. Technical report, National Institute of Standards and Technology (2016)
10. Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 559–585. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_20
11. Czajkowski, J., Majenz, C., Schaffner, C., Zur, S.: Quantum lazy sampling and game-playing proofs for quantum indistinguishability. *Cryptology ePrint Archive, Report 2019/428* (2019). <https://eprint.iacr.org/2019/428>
12. Eisenträger, K., Hallgren, S., Kitaev, A., Song, F.: A quantum algorithm for computing the unit group of an arbitrary degree number field. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing. STOC 2014, pp. 293–302. ACM, New York (2014)
13. Gagliardini, T., Hülsing, A., Schaffner, C.: Semantic security and indistinguishability in the quantum world. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 60–89. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53015-3_3

14. Garg, S., Yuen, H., Zhandry, M.: New security notions and feasibility results for authentication of quantum data. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10402, pp. 342–371. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_12
15. Hayashi, M.: Optimal sequence of quantum measurements in the sense of Stein’s lemma in quantum hypothesis testing. *J. Phys. A: Math. Gen.* **35**(50), 10759 (2002)
16. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Breaking symmetric cryptosystems using quantum period finding. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9815, pp. 207–237. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_8
17. Kuwakado, H., Morii, M.: Quantum distinguisher between the 3-round Feistel cipher and the random permutation. In: Proceedings of IEEE International Symposium on Information Theory, pp. 2682–2685, June 2010
18. Kuwakado, H., Morii, M.: Security on the quantum-type Even-Mansour cipher. In: Proceedings of the International Symposium on Information Theory and Its Applications, pp. 312–316. IEEE Computer Society (2012)
19. Lamport, L.: Constructing digital signatures from a one way function. Technical report SRI-CSL-98, SRI International Computer Science Laboratory (1979)
20. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC 2008, pp. 187–196. ACM, New York (2008)
21. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC 2014, pp. 475–484. ACM (2014)
22. Santoli, T., Schaffner, C.: Using Simon’s algorithm to attack symmetric-key cryptographic primitives. *Quantum Inf. Comput.* **17**(1&2), 65–78 (2017)
23. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**(5), 1484–1509 (1997)
24. Song, F., Yun, A.: Quantum security of NMAC and related constructions. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10402, pp. 283–309. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_10
25. Forrest Stinespring, W.: Positive functions on c^* -algebras. *Proc. Am. Math. Soc.* **6**(2), 211–216 (1955)
26. Unruh, D.: Collapse-binding quantum commitments without random oracles. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 166–195. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_6
27. Unruh, D.: Computationally binding quantum commitments. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 497–527. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_18
28. Vadhan, S.P.: Pseudo randomness. *Found. Trends[®] Theor. Comput. Sci.* **7**(1–3), 1–336 (2012)
29. Zhandry, M.: How to construct quantum random functions. In: Proceedings of the 53rd Annual Symposium on Foundations of Computer Science, FOCS 2012, pp. 679–687. IEEE Computer Society, Washington, DC (2012)
30. Zhandry, M.: How to record quantum queries, and applications to quantum indistinguishability. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11693, pp. 239–268. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26951-7_9
31. Zhandry, M.: Quantum lightning never strikes the same state twice. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11478, pp. 408–438. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17659-4_14