On the Power of Nonuniformity in Proofs of Security

Kai-Min Chung[*], Huijia Lin[†], Mohammad Mahmoody[‡], Rafael Pass[§]

January 18, 2013

Abstract

Nonuniform proofs of security are common in cryptography, but traditional black-box separations consider only uniform security reductions. In this paper, we initiate a formal study of the power and limits of nonuniform black-box proofs of security. We first show that a known protocol with a nonuniform security reduction to one-way permutations *cannot* be proven secure through a uniform security reduction based on one-way permutations. Therefore, nonuniform proofs of security are indeed provably more powerful than uniform ones.

We complement this result by showing that many known black-box separations in the uniform regime actually do extend to the nonuniform regime. We prove our results by providing general techniques for extending certain types of black-box separations to handle nonuniformity.

Keywords:. Black-Box Separations, Proofs of Security, Nonuniformity.

[*]Cornell, chung@cs.cornell.edu.

[†]MIT and Boston University, huijia@csail.mit.edu.

[‡]Cornell, mohammad@cs.cornell.edu.

[§]Cornell, rafael@cs.cornell.edu.

Contents

1 Introduction

Most of cryptography relies on unproved hardness assumptions, such as the existence of one-way functions and one-way permutations, the hardness of factoring, etc. Understanding the minimal assumptions needed for proving the security of cryptographic tasks is thus of fundamental importance. Understanding *barriers* to such proofs of security has been an active line of research in the last decades since the seminal work of Impagliazzo and Rudich [IR89]. Their work demonstrates barriers to providing a "black-box construction" of key-agreement from one-way permutations—a black-box construction of a primitive $\mathcal{Q}$ (e.g., key-agreement) uses a primitive $\mathcal{P}$ (e.g., one-way permutations) as an oracle, while the specific details of the implementation of $\mathcal{P}$ is ignored by the construction; additionally, the proof of security is also black-box in the sense that the security reduction only uses the presumed attacker to $\mathcal{Q}$ as a black-box in order to violate the security of $\mathcal{P}$. Subsequently, many other black-box separations between cryptographic primitives have been established (e.g., [Sim98, GKM+00, GKM+00, GMR01, BPR+08, Vah10, KSY11, M11]); this paradigm has also been used to demonstrate lower bounds on the *efficiency* of black-box constructions (e.g., [KST99, GGKT05, LTW05, HHRS07, BMG07, BMM09, DLMM11]). Very recently, several works [PTV11, Pas11, GW11] demonstrated barriers to proofs of security that apply even when the construction is non-black-box (that is, the implementation of $\mathcal{Q}$ may use the code of $\mathcal{P}$ instead of just treating it as an oracle) as long as just the proof of security is black-box.

The result of [IR89] and its follow-ups, however, suffer from the following restriction: They rule out only constructions with *uniform* proofs of security—that is, the "security reduction" is a uniform polynomial-time algorithm. In contrast, some quite commonly used techniques in cryptography make use of *nonuniformity* in the proof of the security (see e.g., [GMR89, GMW91, GO94, ILL89, LPV08]). For example, a common technique is to use a hybrid argument that involves nonuniformly fixing some "prefix" of the experiment to be the "best possible" value for the reduction. This may be thought of as a mild form of non-black-box access to the code (of the primitive and the adversary) by the security reduction. In some cases, initial nonuniform proofs of security were eventually made uniform. A celebrated example is the result of of [Hås90], making the security reduction of the pseudorandom generator construction of [ILL89] uniform (see [Gol93] for more examples), but in general it is not clear whether nonuniformity makes proofs of security more powerful or not, leaving open the following question:

> *Are nonuniform proofs of security inherently more powerful than uniform ones, or can any nonuniform proof of security be made uniform?*

A closely related question is whether (known) barriers for uniform security extend to the nonuniform regime:

> *Do (known) black-box separations handling uniform security reductions extend to handle also nonuniform security reductions?*

In this work we address the above two questions. We answer the first question affirmatively—showing that a known protocol (specifically, the non-malleable commitment of [LPV08], which is based on the existence of one-way permutations) that uses a nonuniform black-box proof of security cannot be proven secure using a uniform black-box proof of security. Thus, assuming the existence of one-way permutations, nonuniform proofs of security are provably more powerful than uniform ones.

Regarding the second question, we show that, although in general black-box separations for the uniform regime do not extend to uniform regime, many *known* black-box separations for the uniform regime actually do extend to the nonuniform regime. We prove our results by providing general techniques for extending certain types of black-box separations to handle nonuniformity. But before further explaining our result, we need to formally define nonuniform proofs of security.

1.1 Formalizing Nonuniform Proofs of Security

Let us start by recalling the formalization of a black-box proofs of security from [RTV04]. In a black-box proof of security for a cryptographic scheme Q based on some cryptographic assumption P, a security proof S^A is an efficient oracle Turing machine that uses any adversary A who breaks Q as an oracle and breaks the security of P; we use the terms "security proof" or "security reduction" interchangeably. In a *black-box construction* [RTV04] the implementation of the new primitive Q uses a implementation of the primitive P as an oracle, and the security reduction also may only access P as an oracle.

In order to allow the security reduction to be nonuniform, an initial approach would be to allow the security reduction S to be a nonuniform circuit (rather than an efficient Turing machine). More formally, we may extend the definition of a black-box security reduction to the nonuniform setting by requiring the existence of an efficient *circuit* S that for every adversary A breaking the security of Q, S^A breaks P. This "naive" way of defining a nonuniform proof of security, however, does not capture known nonuniform proof techniques where, for instance, the security reduction fixes some "prefix" of a computation to its "best value". The reason is that for such nonuniform proofs of security, the nonuniform advice (i.e., the "best value" fixed by the reduction) can depend on the adversary A. Thus, such use of nonuniform advice may be viewed as a limited form of non-black-box use of the adversary. To also capture such techniques, we instead allow the nonuniform advice to be selected as a function of the adversary A and P.

Definition 1.1 (General Constructions with Nonuniform Security Reductions — Informally Stated)**.** A (general) *construction* of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ is consists of the following:

- **Construction Mapping Q.** There exists a mapping $Q(\cdot)$ such that: if P is an efficient implementation for $\mathcal{P}$, $Q(P)$ is an efficient implementation for $\mathcal{Q}$.
- **Nonuniform Proof of Security.** For every (computationally unbounded) adversary A breaking the security of the efficient construction $Q(P)$, there is a polynomial-size circuit S (which may depend on A and Q) such that S^A breaks the security of P.

Definition 1.2 (Fully Black-box Constructions with Nonuniform Security Reductions — Informally Stated)**.** A *fully black-box construction* of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ consists of:

- **Black-box Construction Q.** For every (computationally unbounded) oracle P implementing $\mathcal{P}$, Q^P implements $\mathcal{Q}$.
- **Nonuniform Proof of Security.** For every (computationally unbounded) oracle P implementing $\mathcal{P}$ and every (computationally unbounded) adversary A breaking the security of Q^P (as an implementation of $\mathcal{Q}$), there exists a polynomial-size security reduction S (which may depending on A and P) such that $S^{P,A}$ breaks the security of P as an implementation of $\mathcal{P}$.

The above way of treating nonuniform proofs of security is closely related to a notion considered in [Lu09, SV10, AS11] in the context of hardness amplification, and Lu [Lu09] in the context of constructions of pseudorandom generators from one-way functions; however, as far as we know, no general treatment of nonuniformity in the context of black-box separations exists.

Before describing our results, let us briefly mention the work of Koblitz an Menezes [KM12] which argues that it is "unnatural and undesirable" to use nonuniform proof of security for "practice-oriented" cryptography. We do not take a stand on this claim; however, in our opinion, due to the extensive use of nonuniform proofs of security in the current cryptographic literature, we believe it is of crucial importance to understand the power and limitations of such techniques in a precise and formal way.

1.2 Our Results

1.2.1 Separating the power of uniform and nonuniform security reductions

Our first results shows a separation between the power of uniform and nonuniform proofs of security.

Theorem 1.3 (Informally Stated)**.** *There exists a construction C of commitments from one-way permutations with a nonuniform security reduction such that: the existence of a uniform security reduction for C means that the one-way permutation in use is efficiently invertible. Thus, nonuniform proofs of security are provably more powerful than uniform ones.*

As mentioned earlier, the commitment scheme we consider is not an artificial one—it is the actual non-malleable commitment scheme constructed in [LPV08].[1] To prove the above theorem, we show that the recent framework of [Pas11]—which proves barriers to cryptographic constructions using uniform security reductions—can be extended to rule security proofs for commitment schemes of the above type. Interestingly, as we shall shortly see, the actual result proven in [Pas11] indeed extends to the nonuniform regime, but we can still use this framework for our separation. See Section 4 for more details.

Related Work. The work of Backes and Unruh [BU08] compares the power of uniform and nonuniform security reductions (the former are called *constructive* proofs in [BU08]). They present a cryptographic protocol Π, assuming some new (and very strong) complexity-theoretic objects, such that security of Π cannot be based on the collision-resistance of a class of hash functions using a uniform security reduction, but can be proven secure using a nonuniform security reductions; the authors also argue that using non-standard assumption may be needed to establish such a separation result. In contrast, our Theorem 1.3 is unconditional (and is additionally demonstrated for a natural scheme already appearing the literature).

1.2.2 Handling nonuniform reductions in fully black-box constructions

As mentioned, following the seminal work of Impagliazzo and Rudich [IR89] there has been a large body of work proving barriers to providing a "black-box construction" of various tasks from various primitives. Ten years after the work of [IR89], using a "reconstruction technique," Gennaro and Trevisan [GT00] (in the context of studying the efficiency of the black-box constructions) showed

[1] [LPV08] also provide a construction of a non-malleable commitment based on one-way functions; our separation does not seem to apply to this protocol.

that a random permutation $P\colon \{0,1\}^n \mapsto \{0,1\}^n$ with overwhelming probability is *nonuniformly-hard* to invert. We first observe that the result of [GT00] can be used to directly extend the original result of [IR89] to rule out black-box constructions of key-agreement from one-way permutations also with respect to nonuniform proofs of security. This follows since: **(1)** [IR89] show how to break all key-agreement protocols through a *single* inefficient but $\mathrm{poly}(n)$-query attack, and **(2)** [GT00] show that any fixed computationally unbounded algorithm that gets $\mathrm{poly}(n)$ bits of advice about the random permutation P and may ask $\mathrm{poly}(n)$ queries to P, has negligible probability of inverting P; thus, if a black-box construction with a nonuniform security proof had existed, we could use the $\mathrm{poly}(n)$-query attacker A^P of [IR89] in conjunction with the nonuniform security reduction S to obtain a fixed computationally unbounded algorithm S' that inverts a random permutation P by getting $\mathrm{poly}(n)$ bits of advice about P and asking only $\mathrm{poly}(n)$-queries to it (contradicting [GT00]).

The "reconstruction technique" of [GT00] has subsequently been employed in several other black-box separation results [GGK03, HHRS07, HH09] and by the same argument these results also extend to the nonuniform setting.

In this work, we also establish new black-box separations results in the nonuniform setting, and a primitive that we focus on is the families of collision-resistant hash functions. We note that although nonuniform hardness results have been proved also for other (idealized forms of) cryptographic primitives [DTT10], as far as we know, no nonuniform hardness result have been proved for families of collision-resistant hash functions, leaving open the question of whether there exists a black-box construction of a key-agreement protocol from families of collision-resistant hash functions with a nonuniform proof of security. (It is well-known that in the uniform setting, the separation of [IR89] easily extends to hash functions.) We start by proving such a separation also in the nonuniform setting.

Theorem 1.4. *There is no fully black-box construction of key-agreement protocols from families of collision-resistant hash functions even with a nonuniform proof of security.*

Since public-key encryption and oblivious transfer both imply key-agreement in a black-box way [GKM+00], Theorem 1.4 extends to separate families of collision-resistant hash functions from those primitives as well. Our proof proceeds by proving a nonuniform hardness lower bound for families of collision-resistant hash functions, and next relying on the above proof template sketched for the case of one-way permutations.

Theorem 1.5 (Nonuniform Collision Resistance)**.** *Let the function $h\colon \{0,1\}^k \times \{0,1\}^m \mapsto \{0,1\}^n$ be chosen uniformly at random and $k \geq 4n, m > n$. Then with $1 - \mathrm{negl}(n)$ probability h will be $2^{n/10}$-secure as a collision-resistant hash function family $\{h_K : K \in \{0,1\}^k\}$. Namely, any circuit of size $2^{n/10}$ with h gates can find collision in h for at most $2^{n/10}$ fraction of the keys $K \in \{0,1\}^k$.*

The proof of Theorem 1.5 follows by a simple counting argument and an application of a lemma due to Unruh [Unr07] (see Lemma 4.2). The proof extends to any *family* version of natural cryptographic primitives.

By applying Theorem 1.5 we also extend some earlier lower bounds [GGKT05, BMG07] on the *efficiency* of black-box constructions from families of collision-resistant hash functions to the nonuniform regime.

Theorem 1.6. (Efficiency of constructions using FCRHs.)

1. *Any fully black-box construction $G\colon \{0,1\}^{\ell} \mapsto \{0,1\}^{\ell+k}$ of PRGs from families of collision-resistant hash functions with a nonuniform security reduction needs $\Omega(k/\log n)$ oracle calls to FCRHs.*

2. *Any fully black-box construction of digital signatures from FCRHs with a nonuniform security reduction and for messages of length at least n bits and needs $\Omega(n)$ oracle calls to FCRHs.*

1.2.3 Handling nonuniform reductions in general constructions

In recent years, new types of black-box separations have emerged. These types of separation apply even to non-black-box constructions, but still only rule out black-box proofs of security: Following the works of Brassard [Bra79] and Akavia et al [AGGM06], demonstrating limitations of "NP-hard Cryptography",[2] Pass [Pas06] and Pass, Tseng and Venkitasubramaniam [PTV11] demonstrate that under certain (new) complexity theoretic assumptions, various cryptographic tasks cannot be based on *one-way functions* using a black-box security reduction, even if the protocol uses the one-way function in a non-black-box way. Very recently, two independent works demonstrate similar types of lower bounds, but this time ruling our security reductions to a *general* set of intractability assumptions: Pass [Pas11] demonstrates unconditional lower bounds on the possibility of using black-box reductions to prove the security of several primitives (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening, Chaum Blind signatures, etc) based on any "bounded-round" intractability assumption (where the challenger uses an a-priori bounded number of rounds, but is otherwise unbounded). Gentry and Wichs [GW11] provides a lower bound (assuming the existence of strong pseudorandom generators) on the possibility of using black-box security reductions to prove soundness of "succinct non-interactive arguments" (SNARGs) based on any "falsifiable" assumption (where the challenger is computationally bounded). Both of the above-mentioned work fall into the "meta-reduction" paradigm of Boneh and Venkatesan [BV98], which was previously used to prove separations for restricted types of reductions (see e.g., [BMV08, HRS09, FS10]).

As with the literature on fully black-box separations, the results of [Pas11, GW11] are only proved in the context of uniform security reductions. Here, we show that these results actually do extend to the nonuniform regime as well.

Theorem 1.7 (Informally Stated)**.** *The separations results of [Pas11, GW11] hold also when considering nonuniform black-box proofs of security.*

2 Preliminaries

2.1 Notation

For any Boolean string x, by $|x|$ we denote the length of x. By $[k]$ we denote the set $\{1,\ldots,k\}$. We use bold letters (e.g., $\mathbf{x}$) when referring to random variables. By $x \xleftarrow{\$} \mathbf{x}$ we mean that x is sampled according to the distribution of the random variable $\mathbf{x}$. We use calligraphic letters (e.g., $\mathcal{S}$) to denote sets (e.g., events over random variables) and cryptographic primitives (e.g., one-way function). We use sans-serif letters (e.g., NP) to denote complexity classes. For a set $\mathcal{S}$, by $\mathbf{U}_{\mathcal{S}}$ we

[2] See also the results of Feigenbaum and Fortnow [FF93] and the result of Bogdanov and Trevisan [BT06] that demonstrate limitations of NP-hard cryptography for *restricted* types of reductions.

mean the random variable with uniform distribution over $\mathcal{S}$, and by $x \xleftarrow{\$} \mathcal{S}$ we mean $x \xleftarrow{\$} \mathbf{U}_S$. By the *support* of the random variable $\mathbf{y}$, represented by $\mathrm{Supp}(\mathbf{y})$, we mean $\{y \mid \Pr[\mathbf{y} = y] > 0\}$. For jointly distributed random variables $(\mathbf{x}, \mathbf{y})$, and for any $y \in \mathrm{Supp}(\mathbf{y})$, the conditional distribution $(\mathbf{x} \mid y)$ is the random variable $\mathbf{x}$ conditioned on $\mathbf{y} = y$.

When we say that an event parameterized by n occurs with negligible probability, denoted by $\mathrm{negl}(n)$, we mean that it occurs with probability $n^{-\omega(1)}$, and we say it happens with overwhelming probability if it happens with probability $1 - \mathrm{negl}(n)$. We call two random variables $\mathbf{x}, \mathbf{y}$ (or their corresponding distributions) over the support set $\mathcal{S}$ ε-close if their statistical distance, defined as $\Delta(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \cdot \sum_{s \in \mathcal{S}} |\Pr[\mathbf{x} = s] - \Pr[\mathbf{y} = s]|$, is at most ε. We call an algorithm D an ε-distinguisher between the random variables $\mathbf{x}$ and $\mathbf{y}$ if $|\Pr[D(\mathbf{x}) = 1] - \Pr[D(\mathbf{y}) = 1]| \geq \varepsilon$. It is easy to see that if there is any D that ε-distinguishes between $\mathbf{x}$ and $\mathbf{y}$, then $\Delta(\mathbf{x}, \mathbf{y}) \geq \varepsilon$.

We use the term *efficient* for any *probabilistic* algorithm that runs in polynomial time over its input length. By the *size* of a circuit we refer to the number of bits that is required to describe it. So the number of circuits of size k will be at most 2^k.

2.2 Intractability Assumptions and Security Reductions

We recall the definition of an intractability assumption from [Pas11] (see also [Nao03,DOP05,HH09, RV10,GW11]).

Definition 2.1 (Intractability Assumptions)**.** A *intractability assumption* $\mathcal{C}$ is a two party game between a *challenger* Chal and an adversary Adv where both parties get 1^n as common input and Chal at the end outputs accept or reject. Any intractability assumption $\mathcal{C}$ has a *security threshold* $\tau_{\mathcal{C}}$ assigned to it which is a constant in the interval $[0, 1)$. We say that an interactive algorithm Adv *breaks* $\mathcal{C}$, if Adv (over the common input 1^n) can make Chal accept with probability $\tau_{\mathcal{C}} + \varepsilon(n)$ for a nonnegligible function $\varepsilon(n)$. When the adversary wins with probability $\tau_{\mathcal{C}} + \varepsilon$ we say that he has won the game with *advantage* ε. We say that $\mathcal{C}$ can be *uniformly (resp. nonuniformly) broken* if there is a PPT (resp. $\mathrm{poly}(n)$-sized circuit) Adv that breaks $\mathcal{C}$. A *falsifiable* assumption is an intractability assumption where Chal is polynomial-time in the length of the messages it receives. A *bounded-round* assumption is an intractability assumption where the game between $\mathcal{C}$ and Adv has a fixed $\mathrm{poly}(n)$ number of rounds.

Cryptographic Primitives. A cryptographic primitive is a syntactical requirement over a set of algorithms performing some cryptographic task. For example if $\mathcal{P}$ denotes the primitive one-way permutation, then it simply requires some algorithm P that computes some (hopefully one-way) permutation. We call (a computationally unbounded) oracle P an *implementation* of the primitive $\mathcal{P}$, if P satisfies the syntactical requirements of $\mathcal{P}$ (when all composed in one algorithm). We say an algorithm P *efficiently* implements $\mathcal{P}$ if it implements $\mathcal{P}$ and runs in polynomial time. We always assume that the algorithm P implementing the primitive $\mathcal{P}$ takes as its first input 1^n (and if it is efficient, it will run in time $\mathrm{poly}(n, |x|)$ where x is the main input). In most primitives the security parameter n can be related to the input length. For example in case of one-way function, n is usually taken to be equal to the input length, or for the case of PKE, it could be the length of the random seed used in the key-generation.

The security of (the efficient implementations of) almost all cryptographic primitives can be modeled as an intractability assumption. For instance, the security game of signature schemes is falsifiable, but it is not bounded-round, since the adversary is allowed to choose the number of

received signatures before trying to forge one; soundness of a constant-round interactive argument, on the other hand, is a bounded-round assumption, but not falsifiable since checking whether an attacker succeeds is not efficient. The security threshold is usually either 0 (e.g., inverting a one-way function) or 1/2 (e.g., distinguishing a PRG from a uniform string).

Definition 2.2 (Cryptographic Primitives)**.** A cryptographic primitive $\mathcal{P}$ is a tuple $(\mathcal{F}_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}})$ where every $P \in \mathcal{F}_{\mathcal{P}}$ is a function implementing P, and $\mathcal{R}_{\mathcal{P}}$ is a relation whose first component is always in $\mathcal{F}_{\mathcal{P}}$. When $(P, A) \in \mathcal{R}_{\mathcal{P}}$, we say that the "adversary" A breaks P (as an implementation of $\mathcal{P}$). By P_n we denote the implementation P restricted to the security parameter n. We restrict ourself to "natural" primitives where the security of any implementation P of $\mathcal{P}$ is defined through an intractability game with security threshold $\tau_{\mathcal{P}}$ that only depends on $\mathcal{P}$.

In the following, we formalize the definition of cryptographic constructions (fully black-box and general) and security reductions (uniform and nonuniform) separately. Formal variants of Definitions 1.2 and 1.1 can be obtained directly from these definitions.

Definition 2.3 (Cryptographic Constructions)**.** A (general) *construction* of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ is a mapping $Q(\cdot)$ such that: if P is an efficient implementation for $\mathcal{P}$, $Q(P)$ is an efficient implementation for $\mathcal{Q}$. A *black-box construction* is a particular form of construction such that: $Q(P)$ is defined through an efficient (uniform) oracle machine Q accessing P only as an oracle, and Q^P is an implementation for $\mathcal{Q}$ for *any* (possibly inefficient) implementation P of $\mathcal{P}$.

We might use the terms "non-black-box" and "general" (constructions) interchangeably.

Definition 2.4 (Uniform and Nonuniform Security Reductions)**.** We say that a (general or black-box) construction Q of a primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ has a *nonuniform* (black-box) security reduction, if for every implementations P for $\mathcal{P}$ and $Q(P)$ for $\mathcal{Q}$, and every adversary A such that A_n breaks $Q_n(P)$ (over security parameter n) with advantage $\varepsilon \geq 1/\operatorname{poly}(n)$, there is a $\operatorname{poly}(n/\varepsilon)$-sized *oracle circuit* S (whose code might depend on (P, A)) such that S^{P,A_n} breaks the security of P_m for some polynomially related security parameter $m = n^{\Theta(1)}$. A *uniform* security reduction is defined similarly by requiring the code of $S^{P,A_n}(1^n, 1^{1/\varepsilon})$ to be uniform and independent of the choices of the oracles (P, A).

We might use the terms "security reduction" and "proof of security" interchangeably.

One can potentially define non-black-box proofs of security as well, but throughout this paper security reductions are always black-box.

Remark 2.5. The reason that we allow the security reduction S to call A only over a single security parameter n (which is polynomially related to $m = n^{\Theta(1)}$) is that here we work with the "almost-everywhere" notion of security (where any infinite–but arbitrarily sparse–sequence of security parameters $\{n_1 < n_2 \ldots\}$ over which A "wins" shall be transformed into a sequence of security parameters $\{m_1 < m_2 \ldots\}$ over which P is broken).

The formal definition of a general construction of a primitive $\mathcal{Q}$ from another primitive together $\mathcal{P}$ with a nonuniform security reduction (i.e. the formalized form of Definition 1.1) can be obtained directly from Definitions 2.3 and 2.4. For the case of black-box constructions, we employ the following terminology whose uniform variant is due to [RTV04].

Definition 2.6 (Fully Black-Box Constructions)**.** A *fully* black-box construction of a primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ consists of a black-box construction Q of $\mathcal{Q}$ from $\mathcal{P}$ together with a (uniform or nonuniform) black-box security reduction for this construction.

For every fixed efficient implementation P of a primitive $\mathcal{P}$, and every (black-box or non-black-box) construction Q of primitive $\mathcal{Q}$ from $\mathcal{P}$, a black-box reduction of the security of $Q(P)$ to that of P can usually be modeled as an intractability assumption. Thus, intractability assumptions allow us to model the (uniform or nonuniform) proofs of security for cryptographic constructions directly (regardless of whether the construction is black-box or not).

Note that any intractably assumption $\mathcal{C} = (\mathsf{Chal}, \mathsf{Adv})$ can be considered as a cryptographic primitive $\widetilde{\mathcal{C}}$ such that: the set of implementations of $\widetilde{\mathcal{C}}_n$ (over security parameter n) only contains the empty function, and the security of $\widetilde{\mathcal{C}}_n$ is defined based on the interactive game $(\mathsf{Chal}, \mathsf{Adv})$ over security parameter n. Based on this perspective, the following definition formalizes what it means to base a cryptographic primitive on such a primitive.

Definition 2.7 (Black-Box Reductions to Intractability Assumptions)**.** We say that a cryptographic primitive $\mathcal{Q}$ can be based on a intractability assumption $\mathcal{C} = (\mathsf{Chal}, \mathsf{Adv})$ through a uniform (resp. nonuniform) black-box reduction iff there exists a construction of $\mathcal{Q}$ from $\mathcal{P} = \widetilde{\mathcal{C}}$ with a uniform (resp. nonuniform) security reduction.

Refuting the possibility of basing a cryptographic primitive $\mathcal{Q}$ on any (falsifiable/bounded round/general) intractability assumption through a uniform (resp. nonuniform) black-box reduction immediately rules out the possibility of basing $\mathcal{Q}$ on a large class of natural cryptographic primitives (whose security for efficient implementations are of the form of intractability assumptions) for uniform (resp. nonuniform) black-box security reductions.

2.3 Special Soundness and Witness Hiding

We assume the reader is familiar with the notions Witness Indistinguishability and Commitment schemes. We refer the reader to [Gol04] for formal definitions.

Special Soundness. Recall that a three-round public-coin interactive proof is said to be *special-sound*, if a valid witness to the statement x can be efficiently computed from any two accepting proof-transcripts of x which have the same first message but different second messages. [Pas11] considers a relaxation of this notion—referred to as *computational special-soundness*—where **(a)** the number of communication rounds is any constant (instead of just three), **(b)** the extractor may need a polynomial number of accepting transcripts (instead of just two), and **(c)** extraction need only succeed if the transcripts are generated by communicating with a computationally-bounded prover (see Section 4 for a formal definition). All traditional constant-round public-coin proofs of knowledge protocols (such as [GMR89, GMW91, Blu87, Sch91], as well as instantiations of [GMW91, Blu87] using statistically-hiding commitments) satisfy this property, and continue to do so also under parallel repetition. We say that a computationally special-sound protocol has a *large challenge space* if the length of the verifier challenge is $\omega(\log n)$ on common inputs of length n.

Definition 2.8 (Computational Special-Soundness – Definition 6 in [Pas11])**.** Let (P, V) be a k-round (where k is a constant) public-coin interactive argument for the language $L \in \mathsf{NP}$ with witness

relation R_L. (P, V) is said to be *computationally special-sound* if there exists a polynomial $m(\cdot)$, and a polynomial-time extractor machine X, such that for every polynomial-time deterministic machine P^*, and every polynomial $p(\cdot)$, there exists a negligible function $\mu(\cdot)$ such that the following holds for every $x \in L$ and every auxiliary input z for P^*. Let $\vec{T} = (T_1, T_2, \dots T_{p(|x|)})$ denote transcripts in $p(|x|)$ random executions between $P^*(x, z)$ and $V(x)$ where V uses the same randomness for the first $k-2$ messages (thus, the first $k-1$ messages are the same in all transcripts). Then, the probability (over the randomness used to generate $\vec{T}$) that:

1. $\vec{T}$ contains a set of $m(|x|)$ accepting transcripts with different round $k-1$ messages; and
2. $X(\vec{T})$ does not output a witness $w \in R_L(x)$

is smaller than $\mu(|x|)$. We say that a computationally special-sound protocol has a *large challenge space* if the length of the verifier challenge is $\omega(\log n)$ on common inputs of length n.

Witness Hiding. A desirable property of interactive proofs is that they "hide" the witness used by the prover. We will consider a very weak notion of sequential witness hiding: roughly speaking, a protocol is said to be weakly sequential witness hiding if no polynomial time attacker can always recover the witness for any statement that it hear polynomially many sequential proofs of.

Definition 2.9 (Breaking Weak Witness Hiding)**.** Let (P, V) be an argument for the language L with witness relation R_L. We say that (a potentially unbounded) A *breaks weak $\ell(\cdot)$-sequential witness hiding of* (P, V) with respect to R_L if for every $n \in \mathbb{N}$, $x \in L \cap \{0,1\}^n$, and $w \in R_L(x)$, A wins in the following experiment with probability 1: Let $A(x)$ sequentially communicate with $P(x, w)$ $\ell(n)$ times; A is said to win if it outputs a witness w' such that $w' \in R_L(x)$. (P, V) is called *weakly $\ell(\cdot)$-sequentially witness hiding* w.r.t R_L if no polynomial time algorithm A breaks weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L.

Let us now state the uniform separation result of [Pas11]:

Theorem 2.10 (Main Result of [Pas11])**.** *Let (P, V) be a computationally-special-sound argument with large challenge space for the language L with a unique witness relation R_L, and let $\mathcal{C}$ be an $r(\cdot)$-round assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists a black-box security reduction S for basing weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L on intractability assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an efficient algorithm B such that $B(1^n)$ breaks $\mathcal{C}$ w.r.t. $\tau_{\mathcal{C}}$ with advantage $\varepsilon(n)$ for a nonnegligible $\varepsilon(n)$.*

Pass [Pas11] shows how Theorem 2.10 can be used to separate several well-known cryptographic protocols/primitive (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening) from any intractability assumption w.r.t. uniform black-box reductions. It directly follows that if Theorem 2.10 is extended to the nonuniform setting, then these corollaries also extend to the nonuniform setting. We refer the reader to [Pas11] for more details on these corollaries.

3 Separating Uniform and Nonuniform Security Reductions

In this section, we demonstrate that a commitment scheme from [LPV08] (LPV), which is proven secure based on the existence of one-way permutations through a *nonuniform* black-box proof of

security, cannot be based on any bounded-round falsifiable assumptions through a uniform black-box proof of security; in particular, it cannot be based on the existence of one-way permutations using a uniform black-box proof of security. As such, we separate the power of nonuniform and uniform black-box proofs of security (assuming the existence of one-way permutations). Let us start by reviewing the LPV protocol.

The LPV protocol is a non-malleable commitment scheme [DDN00a], where the committer and the receiver receive as common input a security parameter n and an identity id of some length $\ell(n)$. To commit to a value v, committer and the receiver proceeds in three stages. In Stage 1, the receiver sends a random string $s \in \{0,1\}^n$. Then in Stage 2, the committer commits to v by sending string c using a non-interactive perfectly binding commitment scheme com. Finally, in Stage 3, the committer proves that either c is a valid commitment to v or it knows the pre-image of s through a one-way permutation f (that is, that it knows a string r such that $f(r) = s$). More specifically, the witness relation used is defined as follows: witness relation $R_L = \{((s,c),y) \mid f(y) = s \text{ or } c = \mathsf{com}(y;r) \text{ for some } r\}$. This is proved using $4\ell(n)$ invocations of a 3-round public-coin WI special-sound (WISSP) argument. Messages in these WISSP arguments are scheduled according to a special scheduling based on id that guarantees that there exist at least $2\ell(n)$ sequential WISSP arguments in the protocol. Our result will only rely on this fact, and thus for simplicity of exposition, below we outline our result w.r.t. a simplified protocol consisting of 2ℓ sequential WISSP arguments in Stage 3.[3] We refer to this protocol as $(C,R)_\ell$.

The following lemma is shown in [LPV08].

Lemma 3.1 ([LPV08]). *Let ℓ be any polynomial. Assuming the existence of one-way permutations, $(C,R)_\ell$ is perfectly binding and computationally hiding with a nonuniform black-box proof of security.*

To highlight the power of nonuniformity in security reductions, let us briefly sketch a proof of this lemma. Fix a polynomial ℓ. Assume that there is an adversary A (w.l.o.g. deterministic) that breaks the hiding property of the LPV protocol $(C,R)_\ell$; that is, for two values v_0, v_1, the adversary A after receiving a commitment to one of the values chosen at random, can guess with inverse polynomial probability which value it has received a commitment to. We now demonstrate the existence of a nonuniform reduction R that with black-box access to A breaks the computational hiding property of the perfectly binding commitment com. Since A is deterministic, the first message s that it sends is fixed. Thus the reduction R can receive as a nonuniform advice the pre-image r of s through the OWP f (i.e, a string r such that $f(r) = s$). Then, after receiving a com commitment c to one of the two values v_0, v_1 chosen at random, it can use the adversary A to guess which value it receives a commitment to by forwarding c to A and simulating all the WISSP arguments using r as a fake witness; finally, it outputs A's guess. It follows from the witness indistinguishability property of the WISSP argument that R^A has inverse polynomial advantage in guessing whether it received a commitment to v_0 or v_1.

Let us now turn to showing the impossibility of basing the hiding property of the LPV protocol on any bounded-round falsifiable assumptions through a *uniform* black-box proof of security.

Theorem 3.2. *Let $\mathcal{C}$ be an $r(\cdot)$-round falsifiable assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists an efficient black-box security reduction S for basing the hiding property of $\langle C,R\rangle_\ell$ on assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an efficient algorithm B and a polynomial $p(\cdot)$ such that for infinitely many $n \in N$, $B(1^n)$ breaks $\mathcal{C}$ with advantage $1/p(n)$.*

[3] It is easy to see that the same argument applies also to the original LPV protocol.

We first provide an outline of the proof and then will describe the full proof.

3.1 Outline of the Proof of Theorem 3.2

At first sight, it would seem that Theorem 3.2 directly follows from Theorem 2.10: At a very high-level, the LPV protocol simply consists of many sequentially repeated special-sound arguments (that are also constant-round and public coin) for the statement (s, c) defined by the messages in Stage 1 and 2. Additionally, demonstrating hiding, at the very least implies that these protocols are weakly $l(n)$-sequentially witness hiding (or else, the committed value can be completely recovered!).

However, this approach does not go through as Theorem 2.10 only provides a separation in the case of *unique witness languages*, whereas R_L admits two witnesses for every statement. To get around this problem, we consider a unique witness relation $R_{L'}$ for which every statement (s, c) only has one unique witness that is the unique committed value in c.

But if we change the language, the proofs in Stage 3 of the protocol are no longer special-sound for the witness relation $R_{L'}$, since for a specific instance (s, c), it might be easy to invert s and therefore the value extracted from a WISSP argument (in Stage 3 of the LPV protocol) might be the pre-image r instead of the committed value, violating the (computational) special soundness property for $R_{L'}$. We resolve the problem by observing that, in fact, the proof of the Theorem 2.10, when restricted to falsifiable bounded-round assumption (as opposed to general bounded-round assumption) in [Pas11] itself is black-box and uniform[4]. Roughly speaking, theorem 2.10 states that if (P, V) is computationally special-sound for a unique witness language L, then it is impossible to base the sequential witness hiding property on any bound-round falsifiable assumptions through uniform black-box security reduction. This is proven by showing that for any (public-coin, constant-round) interactive protocol (P, V), if there is a uniform black-box security reduction S for basing its sequential witness hiding property for $R_{L'}$ on any bounded-round falsifiable assumption, then there is a *uniform meta reduction* M that with black-box access to S can violate the computational special-soundness of (P, V) w.r.t. $R_{L'}$. In our context, since (P, V) is a special-sound proof for R_L (as opposed to $R_{L'}$), this meta reduction may either violate computational special-soundness of (P, V) w.r.t. R_L, or may output a pre-image r to s. However, as long as s is chosen at random by M, we can thus use such an M to invert a random string s through f (assuming that computational special-soundness of (P, V) for R_L holds). See Section 3.2 for the detailed proof.

Let us briefly comment on the why the above proof sketch does not extend to nonuniform proofs of security (whereas as we shall see in Section 4.1.2, the main theorem of [Pas11] does). The problem is that if the reduction S is nonuniform, it may get as nonuniform advice the string s (or some function of it); this means that in the above proof, M can no longer choose the string s uniformly at random (since S would notice this). Indeed, in the actual nonuniform proof of security of LPV, the reduction does get the pre-image of s as nonuniform advice.

3.2 Formal Proof of Theorem 3.2

In this section, we describe a protocol that can be proven secure with a (black-box) nonuniform proof of security, but cannot be proven secure with a (black-box) uniform proof of security. The protocol is the concurrent non-malleable commitment scheme introduced in [LPV08] (LPV). Below we recall the protocol.

[4]In fact, the proof of Theorem 2.10 as stated in [Pas11] in actually *nonuniform*, but for the special case that the intractability assumption is falsifiable, the proof is uniform.

The LPV protocol is based on Feige-Shamir's zero-knowledge protocol [FS87] while relying on the *message scheduling technique* of Dolev, Dwork and Naor [DDN00b]. Their protocol relies on the existence of one-way functions with efficiently recognizable range, and can be modified to work with arbitrary one-way functions by additionally providing a witness hiding proof that an element is in the range; here we directly instantiate their protocol with one-way permutations. Their protocol also uses a two round statistically binding commitment scheme and we instantiate it with an non-interactive perfectly binding commitment scheme. Let ℓ be a polynomial; the LPV commitment protocol for identities of length ℓ—denoted as $\langle C, R\rangle_\ell$, proceeds in the following three stages on common input a security parameter n and an identity $\mathsf{id} \in \{0,1\}^{\ell(n)}$.

1. In Stage 1, the Receiver picks a random string $s \in \{0,1\}^n$, and sends it to the Committer. Let r be the pre-image of s through a one-way permutation f.

2. In Stage 2, the Committer sends a commitment c to v, using a non-interactive perfectly binding commitment scheme com.

3. In Stage 3, the Committer proves that c is a valid commitment to v or it knows the pre-image r of s through the one-way permutation f. This is proved by $4\ell(n)$ invocations of a special-sound WI proof where the messages are scheduled based on the id. More precisely, there are $\ell(n)$ rounds, where in round i, the schedule $\mathsf{design}_{\mathsf{id}_i}$ is followed by $\mathsf{design}_{1-\mathsf{id}_i}$ (See Figure 1).

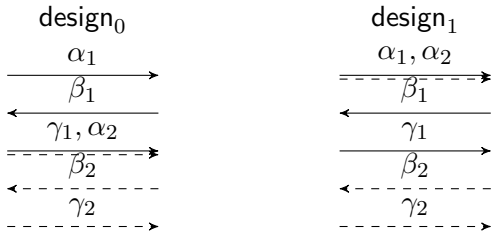

Figure 1: Description of the schedules used in Stage 3 of the protocol. $(\alpha_1, \beta_1, \gamma_1)$ and $(\alpha_2, \beta_2, \gamma_2)$ are respectively the transcripts of a pair of 3-round special-sound proofs.

In the LPV protocol, the special scheduling of the WISSP proofs in Stage 3 is crucial for proving the concurrent non-malleability property of the protocol. In this work, our separation result only relies on the fact that there are many sequential WISSP proofs in Stage 3 (it is easy to see that this is implied by the special scheduling, as all the design's are arranged sequentially). The number of WISSP proofs is decided by the length $\ell(n)$ of the identity. Thus, the protocol contains at least $\ell(n)$ sequential WISSP proofs for some constant c. It is shown in [LPV08] that the above protocol is hiding.

Lemma 3.3 ([LPV08]). *Let ℓ be any polynomial. $\langle C, R\rangle_\ell$ is perfectly binding and computationally hiding with a black-box nonuniform proof of security.*

We show that in contrast, it is impossible to base the hiding property of the above protocol on any falsifiable assumptions through a black-box uniform proof of security; namely we prove Theorem 3.2.

In the following, w.l.o.g we assume that the security reduction S only asks distinct oracle queries. A query of S is of the form $(x, p_1, p_2, \ldots, p_i)$ for some $i \leq k \cdot \ell$ (recall that k is the number of rounds

of (P, V)) where x is a common input to (P, V) and $p_1, \ldots, p_i$ are the prover's messages. If $i < k \cdot \ell$, the oracle answer will be of the form $(v_1, v_2, \ldots, v_{i+1})$ denoting the messages of the verifier, and if $i = k \cdot \ell$, the oracle answer will be of the form $(v_1, v_2, \ldots, v_{k\ell}, w)$ where $(x, w) \in R_L$. W.l.o.g we assume that if the reduction S it making the query $(x, p_1, \ldots, p_i)$, it has previously queried $(x, p_1, \ldots, p_{i-1})$. When it is clear from the context, we represent the query $(x, p_1, \ldots, p_i)$ by $q = p_i$ and represent the answer by $a = v_{i+1}$ (if $i < k \cdot \ell$) or $a = w$ (if $i = k \cdot \ell$).

Ideal Adversary A. In a straight-line interaction, the ideal adversary A behaves like the honest verifier and sends fresh random coins in response to any message from the "prover" (whose role is played by the reduction S), and at the end of $\ell(|x|)$ sequential interactions over the common input x, A returns the unique witness w such that $(x, w) \in R_L$. A also uses fresh randomness in answering any new query *even when rewound*. More formally, A will internally access its own oracle R, and upon any query q as the last message of the partial transcript t, if q is a full transcript of ℓ sequential executions, it returns w such that $(x, w) \in R_L$, and if t is not a full transcript of ℓ executions, A applies R to t and uses $R(t)$ as the randomness used in response to q.

Lemma 3.4 ([Pas11])**.** *Suppose S is a (potentially unbounded) algorithm calling the adversary A described above as a black-box and wins against a challenger* Chal *with probability at least ε. There is an efficient algorithm B that uses S internally* as a black-box *(but could rewind S), interacts with* Chal *externally, and wins with probability at least $\varepsilon - \mathrm{negl}(n)$.*

In order to use the theorem in a modular way, we first observe that in [Pas11], Lemma 3.4 is proven through a *black-box uniform* proof reduction to the computationally-special-soundness of the protocol (P, V) whose sequential witness hiding property is under analysis, when the assumption to reduce to is a *falsifiable assumption*. More precisely, Lemma 3.4 can be restated and extended when considering only falsifiable assumptions as follows: We say that a protocol (P, V) is a canonical interactive protocol if it is public-coin, k-round for a constant k, has large challenge space and the verifier's output is boolean. A canonical interactive protocol is potentially computationally-special-sound for a NP language L, w.r.t. an extractor X that outputs a witness of a statement x on input $m(|x|)$ accepting transcripts with different $k - 1$ messages.

Lemma 3.5 (Uniform Security Proof of [Pas11])**.** *Let L be an* NP *language with unique witness relation R_L, (P, V) a canonical interactive protocol, with potential special-soundness extractor X and polynomial m, and $\mathcal{C}$ be a $r(\cdot)$-round falsifiable assumption with efficient challenger* Chal*, where $r(\cdot)$ is a polynomial. If there exists an efficient algorithm S, for which Lemma 3.4 does not hold w.r.t. R_L, (P, V), X, m and $\mathcal{C}$. There exists an efficient algorithm D that uses S internally* as a black-box *and violates the computational special-soundness property of (P, V) w.r.t. extractor X and polynomial m, with polynomial probability, for a statement $x \in L$ output by S in reply to one of the queries made by D during its execution.*

Now we are ready to prove Theorem 3.2.

of Theorem 3.2. Assume for contradiction that there exists a falsifiable assumption $\mathcal{C}$ with challenger Chal and an efficient black-box uniform security reduction S for basing the hiding property of the LPV protocol on assumption $\mathcal{C}$ with threshold $\tau_{\mathcal{C}}$. Then we derive a contradiction as follows.

Consider an ideal adversary A that breaks the hiding property of the LPV protocol as defined above. Namely, the adversary A acts as an honest receiver of the LPV protocol by sending fresh

random coins using its own random oracle R, and at the end of an accepting commitment execution, it returns the unique committed value defined by the Stage 1 com commitment (if the commitment is invalid, it returns $\perp$). It follows from the soundness of the WISSP arguments in Stage 3 that, for any two messages m_0 and m_1, the probability that A after receiving a commitment to a random message m_b from an honest committer, guesses correctly b with overwhelming probability. That is, A violates the hiding property of the LPV protocol.

Then by our hypothesis, the security reduction S with black-box access to A can win against the challenger Chal with some polynomial probability $1/p(n)$. Next, we view Stage 1 and 2 of the LPV as a process for sampling an instance of a NP language L with the following unique witness relation R_L,

$$R_L = \{(x, w) \ : \ x = (s, c) \text{ and } \exists\ \sigma \text{ s.t. } c = \mathsf{com}(w, \sigma)\},$$

and we view the WISSP arguments in Stage 3 as a canonical interactive protocol (P, V), with potential special-soundness extractor X that is just the extractor of the WISSP proof and polynomial $m(\cdot) = 2$. Using A and S, we construct another ideal adversary $\tilde{A}$ and reduction $\tilde{S}$ for the sequential witness hiding property of the protocol (P, V) w.r.t. R_L.

Ideal Adversary $\tilde{A}$: $\tilde{A}$ on input an instance (s, c) acts as the verifier of $\ell(n)$ sequential executions of (P, V) (i.e., as the verifier of one WISSP arguments in Stage 3) by sending fresh random coins using its own oracle R. At the end of $\ell(n)$ sequential executions of (P, V), if in all executions, the verifier outputs 1, then return the unique witness $w \in R_L((s, c))$ if $s \in L$, or $\perp$ otherwise. Clearly, $\tilde{A}$ breaks the ℓ-sequential witness hiding property of (P, V) for R_L.

Security Reduction $\tilde{S}$: $\tilde{S}$ internally runs S and externally interacts with the challenger Chal. it forwards all the messages between S and Chal, and answers queries from S to the ideal adversary A as follows: $\tilde{S}$ tosses some fresh random coins to obtain a random string s, whenever S expects to receive the first message from A, $\tilde{S}$ returns the random string s; whenever S expects the answer from A for a query q which is a partial transcript of the LPV protocol with Stage 1 and 2 messages s' and c, if $s' = s$, $\tilde{S}$ simply forwards this query q to $\tilde{A}$ using it as a partial transcript of ℓ executions of (P, V) w.r.t. statement $x = (s, c)$; otherwise, if $s' \neq s$ it does not reply.[5]

Since $\tilde{S}$ emulates the view of S perfectly internally, we have that $\tilde{S}$ with black-box access to $\tilde{A}$ wins against the challenger Chal with the same probability as S with black-box access to A does, which is $1/p(n)$. Then we claim that there is an efficient algorithm B with black-box access to $\tilde{S}$ wins against Chal with probability $1/p(n) - \mathrm{negl}(n)$, which gives a contradiction to the security of the falsifiable assumption $\mathcal{C}$.

Assume that there is no efficient algorithm B that satisfy the above property. Then by Lemma 3.5, we have that there exists an efficient algorithm D that with black-box access to $\tilde{S}$ can violate the computational-special-soundness of (P, V) for R_L, $m(\cdot) = 2$ and X, for a statement that is output by $\tilde{S}$ when answering one of the queries from D. By construction, any statement (s, c) output by $\tilde{S}$ has a s sampled at random using the random tape of $\tilde{S}$. Then we can build another machine $\tilde{D}$ that on input a random string s', incorporates D and $\tilde{S}$ internally; it sets the

[5]The LPV protocol contains many sequential design's and each design contains either two sequential WISSP arguments or two interleaved WISSP arguments. Therefore, the queries that $\tilde{S}$ sends to $\tilde{A}$ may contain some interleaved executions of (P, V). This technical issue can be dealt with by having $\tilde{S}$ only forward to $\tilde{A}$ sequential executions of (P, V). For simplicity, we omit the details here.

random tape of $\tilde{S}$ to s', and forwards the statement (s, c) and message of (P, V) from D externally. As argued above, by construction of $\tilde{S}$, s in the statement must be the same as s'. Therefore $\tilde{D}$ breaks the computational special-soundness of (P, V) w.r.t. a statement (s', c) where s' is chosen by an external challenger. This means, with some polynomial probability, either (1) the extractor fails to obtain two accepting transcripts of (P, V) with different second messages, or (2) the extractor X on two accepting transcripts with different second messages outputs a value that is not the unique committed value v in c. It follows from the special-soundness of the WISSP argument, the first scenario happens with negligible probability. On the other hand, by construction of the WISSP arguments in Stage 3 of the LPV protocol and its special-soundness property, except with negligible probability, the value X extracts is either v or a pre-image of s through one-way permutation f. Given that the extracted value is not v, we have that with some polynomial probability, the extracted value is a pre-image of s' which violates that f is one-way. This gives a contradiction. □

4 Refuting Nonuniform Reductions to Intractability Assumptions

In this section we prove Theorem 1.7.

4.1 Extending Theorem 2.10 to Nonuniform Security Reductions

Here we show how to extend Theorem 2.10 to handle nonuniform proofs of security. Pass showed how Theorem 2.10 can be used to prove that certain well known cryptographic protocols/primitive (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening) can not be based on any bounded-round intractability assumption through a black-box proof of security. It follows that all corollaries of [Pas11] also extend to the nonuniform regime.

Theorem 4.1. *Let (P, V) be a computationally-special-sound argument with large challenge space for the language L with a unique witness relation R_L, and let $\mathcal{C}$ be an $r(\cdot)$-round assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists a nonuniform black-box security reduction S for basing weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L on intractability assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an nonuniform polynomial-time algorithm B and a polynomial $p(\cdot)$ such that for infinitely many $n \in N$, $B(1^n)$ breaks $\mathcal{C}$ with advantage $1/p(n)$.*

Note that in Theorem 4.1 we consider also nonuniform security reductions S, but the conclusion is slightly weaker than the conclusion of Theorem 2.10: the machine B breaking the assumption $\mathcal{C}$ is no longer uniform, but it now is a nonuniform polynomial-time algorithms.

We first outline the proof of Theorem 4.1 and then will present a full proof.

4.1.1 Outline of the Proof of Theorem 4.1

Let us start by briefly outlining the high-level approach used in the result of [Pas11] and explain why handling nonuniform proofs of security becomes an issue. Assume there exists a security reduction S (and for now, assume that S is uniform) such that S^A breaks the assumption $\mathcal{C}$ whenever A breaks weak sequential witness hiding of a (computationally) special-sound argument (P, V) for a language with unique witnesses. We want to use S to directly break $\mathcal{C}$ *without the help of* A. So, (following the paradigm of [BV98]), the goal will be to efficiently emulate A for S (i.e., we will construct a "meta-reduction" M which uses the underlying reduction S to break $\mathcal{C}$). We consider a

particular computationally unbounded oracle A that after hearing an appropriate number of proofs using (P, V) (acting as a verifier) simply outputs a witness to the statement proved. The idea is to "extract" out the witness that A is supposed to provide S by "rewinding" S—since (P, V) is computationally special-sound, S, intuitively, must know a witness for all statements x that it proves to A. (There are several obstacles in formalizing this approach. The main one is that the reduction S is not a "stand-alone" prover—it might *rewind and reset* the oracle A, so it is no longer clear that it needs to "know" a witness for x in order to convince A of x. It is here that the proof of [Pas11] relies on the fact that there are multiple proofs being provided by S; this gives the meta-reduction more opportunities to rewind S, which enables extraction even if S "nests" its queries to A in an arbitrary way. We refer the reader to [Pas11] for further details.

Let us point out a crucial component in the above proof: To succeed in its emulation of S, it is imperative that whenever A is acting as a verifier, it chooses fresh random coins to generate its messages, even in case A is rewound (technically, this is achieved by letting A generate its messages by applying a random function to its queries). This is needed to ensure that M can "rewind" S (in order to extract out a witness), sending it new verifier messages, while ensuring that S provides an answer back with the same probability as if S communicated with A. In other words, to ensure that M succeeds in extracting witnesses from S, we require A's verification message to have essentially "full entropy", or else S may be able to notice that it is being rewound, and may abort its computation. In the context of nonuniform reductions, we can no longer guarantee that A's answers have high entropy: S gets a nonuniform advice string as a function of A, and thus the conditional entropy of the answers of A drops. Our approach for getting around this problem is that: although the conditional entropy of A's answers drops (once S gets its nonuniform advice), for all but a polynomial number of "bad" queries to A, the answers to the remaining "good" queries will still have high enough entropy. In fact, by an argument due to Unruh [Unr07], it can be shown that the conditional distribution of answers to "good" queries is statistically close to the original distribution of A.

Lemma 4.2 (Informal Variant of [Unr07])**.** *Suppose A is a randomized oracle, and suppose S is an oracle algorithm that gets as input a nonuniform (polynomial-size) advice z as function of (the description of) A, and then asks polynomially many queries to A. Then there is a "pre-sampling" algorithm* Samp *that given z samples s query-answers of A (according to their true distribution based on A), and the view of $S^A(z(A))$ is $1/\operatorname{poly}(n)$-close in another experiment in which S is given z, then* $\mathsf{Samp}(z)$ *samples a partial domain of A, then A gets* resampled *on every other point other than the outputs of* $\mathsf{Samp}(z)$*, and finally $S(z)$ gets executed with oracle access to the (newly sampled) A.*

If the reduction S had only queried these "good" queries (that are *not* presampled by Samp), we would already be done. But, S may of course ask also "bad" queries (i.e. the ones presampled by $\mathsf{Samp}(\cdot)$). To deal with this, we present a *nonuniform meta reduction* M—M receives as a nonuniform advice the set of "bad" queries (which may depend on the nonuniform advice of S), and for each of these queries, the answer that A actually would provide. M can then perfectly emulate "bad" queries (using its nonuniform advice), and as before emulate good queries (in a statistically close manner) by using rewindings. As a conclusion we get that the existence of a security reduction S can be used to break the intractability assumption $\mathcal{C}$ in *nonuniform polynomial time.*

4.1.2 Formal Proof of Theorem 4.1

Here we prove Theorem 4.1 formally.

Lemma 3.4 is indeed the result proved in [Pas11] which is more general than Theorem 2.10. We use this lemma to prove our Theorem 4.1. We describe four experiments and prove Theorem 4.1 using Lemmas 3.4 and the following lemma due to Unruh 4.3.

Lemma 4.3 ([Unr07]). *There is an (inefficient) algorithm* Samp *that gets as input some* (z, s) *for* $z \in \{0,1\}^*, s \in \mathbb{N}$ *and outputs a partial function* F *with* s *defined points such that the following holds. If* D *is a (computationally unbounded) oracle algorithm that receives an auxiliary input* z *of length* $|z| = d$ *and asks* t *queries to its oracle, then for any function* $Z\colon \{0,1\}^* \mapsto \{0,1\}^d$ *the view of* D *in the following two experiments is* $\sqrt{dt/2s}$*-close in statistical distance:*

1. **(1)** $RO \xleftarrow{\$} \mathbf{RO}$. **(2)** $z = z(RO)$. **(3)** *Execute* $D^{RO}(z)$.
2. **(1)** $RO \xleftarrow{\$} \mathbf{RO}$. **(2)** $z = z(RO)$. **(3)** $F = \mathsf{Samp}(z)$. **(4)** $RO' \xleftarrow{\$} \mathbf{RO}[F]$ **(5)** *Execute* $D^{RO'}(z)$.

Experiment Inef**.** Similarly to experiment Inef in the proof of Theorem 4.10, in this experiment $S^A(z(A))$ interacts with Chal and breaks the assumption $\mathcal{C}$ by winning against Chal with probability $\varepsilon \geq 1/\operatorname{poly}(n)$ while $z(A)$ is a nonuniform advice about the ideal adversary oracle A.

To describe the next experiment we need to fix a few definitions and notations. Suppose S asks $t = t(n)$ oracle queries, and let s be such that $\varepsilon/2 = \sqrt{dt/2s}$ where d is the length of the advice given to S. Suppose A accesses the random oracle $RO \xleftarrow{\$} \mathbf{RO}$ as its random tape and uses $RO(q)$ in order to answer query q. Suppose D is an oracle algorithm with auxiliary input accessing RO as follows: $D(z(A))$ emulates the interaction of $S^A(z(A))$ with Chal and for every oracle query q, it asks $RO(q)$ to obtain the answer $A(q)$ and continues the emulation. For any sampled random oracle RO let $z(RO) = z(A(RO))$ which denotes the nonuniform advice being computed directly over the random oracle. Suppose $\mathsf{Samp}(\cdot)$ is the pre-sampling algorithm of Lemma 4.3 when applied to the advice function z (computed over the random oracle) and the oracle algorithm D.

Experiment Hyb_1**.** This experiment is defined similarly to the experiment Hyb of the proof of Theorem 4.10. More formally we do the following:

1. Get $RO \xleftarrow{\$} \mathbf{RO}$ and obtain $z = z(RO)$.
2. Run $\mathsf{Samp}(z)$ to get the pre-sampled part F of the oracle which includes s query-answer pairs.
3. Get $RO' \xleftarrow{\$} \mathbf{RO}[F]$.
4. Run the interaction between Chal and $S^{A[RO']}(z)$.

Experiment Hyb_2**.** In this experiment the algorithm S receives a randomized advice $y = (z, F, W)$ as follows: The components (z, F) are sampled from the same exact distribution as that of experiment Hyb_1. The set W contains all pairs (x, w) where w is the unique witness for x (i.e. $(x, w) \in R_L$), and there exists a partial transcript $(x, q_1, \ldots, q_i)$ answered in F. After receiving the advice $y = (z, F, W)$, the experiment Hyb_2 continues as follows.

- For any oracle query q asked by S with respect to the input x:

1. If the partial transcript leading to the query q is already answered in F or W use this answer (note that the answer could be a witness or some random coins).
2. Otherwise, if there is any partial transcript $(x, q_1, \ldots, q_i)$ answered in F (with the same x), then: if $i < k\ell$, toss fresh random coins and use it as the answer, and if $i = k\ell$ use the answer w specified in W.
3. Otherwise, call the oracle A for the answer. Note that the internal random oracle $RO \xleftarrow{\$} \mathbf{RO}$ of A is sampled independently of the randomized advice y.

Note that in Hyb_2, if the reduction S asks A for any query q w.r.t. some x, then it does the same for *all* other queries q' w.r.t. the same x.

Experiment Eff. In this experiment, first the algorithm S receives a randomized advice $y = (z, F, W)$ distributed the same way as that of experiment Hyb_2. Then the efficient algorithm B of Lemma 3.4 interacts with $S(y)$ (as specified in Hyb_2) internally, and with Chal externally.

Now we analyze the probability of winning he challenge in the experiments above. Let $\mathcal{W}_X$ be the event that Chal accepts its interaction in experiment X. Note that $\Pr[\mathcal{W}_{\mathsf{Inef}}] \geq \varepsilon$ holds for a nonnegligible $\varepsilon \geq 1/\operatorname{poly}(n)$ by the the definition of A and the black-box security reduction S. Also, note that to prove Theorem 4.1 it is sufficient to prove that $\Pr[\mathcal{W}_{\mathsf{Eff}}] \geq \delta$ for a nonnegligible δ, because we can always hardwire a fixed advice y_0 that maximizes the probability of $\mathcal{W}_{\mathsf{Eff}}$ and get an efficient circuit that wins against Chal with probability at least δ. The bound $\Pr[\mathcal{W}_{\mathsf{Eff}}] \geq 1/\operatorname{poly}(n)$ follows from $\Pr[\mathcal{W}_{\mathsf{Inef}}] \geq \varepsilon$ and the following three claims.

Claim 4.4. $\Pr[\mathcal{W}_{\mathsf{Hyb}_1}] \geq \Pr[\mathcal{W}_{\mathsf{Inef}}] - \varepsilon/2$.

Proof. Claim 4.11 follows directly from Lemma 4.3 and the way the distinguishing algorithm D is defined as a combination of the interactive algorithms Chal and S. In fact the joint view of the parties (Chal, S) in experiments Inef and Hyb are $(\varepsilon/2)$-close. □

Claim 4.5. $\Pr[\mathcal{W}_{\mathsf{Hyb}_2}] = \Pr[\mathcal{W}_{\mathsf{Hyb}_1}]$.

Proof. For any query oracle q by the reduction S either of the following holds:

- q is a query not determined by the presampled part F of the oracle. In this case, either q is answered randomly in both experiment, or the answer is a unique witness for the corresponding x (which will be the same in both experiments).
- q is a query determined by the presampled part F. But this part is sampled identically in both experiments, and so to bound the statistical distance we can assume that identical F is sampled in both experiments.

□

Claim 4.6. $\Pr[\mathcal{W}_{\mathsf{Eff}}] \geq \Pr[\mathcal{W}_{\mathsf{Hyb}_2}] - \operatorname{negl}(n)$.

Proof. This claim follows directly from Lemma 3.4 and the fact that the algorithm B uses the security reduction S as a black-box. □

4.2 Succinct Non-Interactive Arguments

In this section we show how to extend a result of Gentry and Wichs [GW11] on the impossibility of basing succinct non-interactive arguments (SNARGs) for NP on any falsifiable assumption to the nonuniform regime. We start by recalling the formal definition of SNARGs.

A succinct non-interactive argument system Π consists of three efficient algorithms (G, P, V): The generation algorithm G on input security parameter 1^n outputs a common reference string crs and a private verification state priv. The prover algorithm P on input crs, a statement $x \in \{0,1\}^n$ and a witness w outputs a proof π. The verifier algorithm V on input priv, x, and π outputs a bit $b \in \{0,1\}$ represents whether V accepts or rejects the proof π for x.

Definition 4.7 (Succinct Noninteractive Arguments)**.** $\Pi = (G, P, V)$ is a succinct non-interactive argument (SNARG) for an NP language L with relation R_L if following holds. (Below $|x| = n$ and by negligible we mean $\mathrm{negl}(n)$.)

- **Completeness:** For every $(x, w) \in R_L$ the probability that the verifier V rejects in the following experiment is negligible: (i) $(\mathsf{crs}, \mathsf{priv}) \leftarrow G(1^n)$, (ii) $\pi \leftarrow P(\mathsf{crs}, x, w)$, and (iii) $b \leftarrow V(\mathsf{priv}, x, \pi)$.
- **(Adaptive) Soundness:** For every efficient cheating prover P^*, the probability that the verifier V accepts in the following experiment is negligible: (i) $(\mathsf{crs}, \mathsf{priv}) \leftarrow G(1^n)$, (ii) P^* on input crs outputs both a statement x and a proof π, and (iii) $b \leftarrow V(\mathsf{priv}, x, \pi)$.
- **Succinctness:** The length of a proof $\pi \leftarrow P(\mathsf{crs}, x, w)$ generated by the prover is $n^{o(1)}$.

We mention that for the simplicity of exposition, in the above definition, we restrict the length of statements to be the same as the security parameter n, and define the succinctness property by $|\pi| \leq n^{o(1)}$. Our extension of the results of [GW11] holds for the general definition in [GW11] as well. We now recall the formal result of [GW11].

Theorem 4.8 ([GW11])**.** *Assuming the existence of sub-exponentially hard one-way functions, for any SNARG $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, the adaptive soundness of Π can not be based on any falsifiable assumption $\mathcal{C}$ through a* uniform *black-box reduction, unless $\mathcal{C}$ can already be broken (nonuniformly).*

Note that Theorem 4.8 only rules out *uniform* black-box security proof of adaptive soundness. Our goal is to extend the theorem to also rules out *nonuniform* black-box security proof. Towards this goal, we note that core of the proof of Theorem 4.8 in [GW11] is the construction of two adversarial provers, and we identify the key properties of the two adversarial provers that are needed by us.[6]

Lemma 4.9 (Lemma 4.1 in [GW11])**.** *Assuming the existence of sub-exponentially hard pseudo-random generators, for any SNARG system $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, there exist two adversarial provers A and B that satisfy the following properties.*

- *A breaks the adaptive soundness of Π but is inefficient, whereas B is efficient.*

[6] Lemma 4.1 in [GW11] is stated only w.r.t. *efficient* distinguishers (indicating them to be uniform), but since it assumes the subset-membership problem to be nonuniformly hard, the very same proof given for Lemma 4.1 handles nonuniform distinguishers as well, so we state and use this lemma in this form.

- *A and B are computationally indistinguishable in the following sense: for any polynomial-sized circuit family $\{D_n\}$ it holds that*

$$\left|\Pr[D_n^A = 1] - \Pr[D_n^B = 1]\right| = \mathrm{negl}(n).$$

- *Both A and B are randomized and* stateless *and use fresh independent randomness to generate answers to distinct queries (which are the* crs*'s).*

The third property allows us to view A (resp., B) as a deterministic algorithm that given any query $q = \mathsf{crs}$, access fresh randomness $RO(q)$ and returns $A(q, RO(q))$ (resp., $B(q, RO(q))$). We can think of $RO(\cdot)$ as a random oracle with long enough outputs length[7]. More explicitly, by $A[RO]$ we emphasize on the fact that A is using the sampled random oracle $RO \xleftarrow{\$} \mathbf{RO}$.

We are ready to extend Theorem 4.8 to rules out *nonuniform* black-box security proof.

Theorem 4.10. *Assuming the existence of sub-exponentially hard pseudorandom generators, for any SNARG system $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, the adaptive soundness of Π can not be based on any falsifiable assumption $\mathcal{C}$ through a* nonuniform *black-box reduction, unless $\mathcal{C}$ can already be broken (nonuniformly).*

Recall that $\mathbf{RO}$ denotes the distribution of random oracles, and for any partial length preserving function F suppose $\mathbf{RO}[F]$ denotes the distribution of random oracles with pre-sampled part F.

To prove Theorem 4.10 we again use Lemma 4.2 an a careful hybrid argument similar to the proof of Theorem 4.1.

of Theorem 4.10. The proof goes through a hybrid argument. We start by describing a game in which the assumption $\mathcal{C}$ is broken inefficiently (with the help of the oracle A) and then modify the game so that we obtain an efficient adversary breaking the challenge $\mathcal{C}$. In the following, w.l.o.g we assume that the security reduction S only asks distinct oracle queries.

Experiment Inef**.** In this experiment $S^A(z(A))$ interacts with Chal and breaks the assumption $\mathcal{C}$ with probability $\varepsilon \geq 1/\operatorname{poly}(n)$.[8]

To describe the next experiment we need to fix a few definitions and notations. Suppose S asks t oracle queries, and let s be such that $\varepsilon/2 = \sqrt{dt/2s}$ where d is the length of the advice given to S. Suppose A accesses the random oracle $RO \xleftarrow{\$} \mathbf{RO}$ as its random tape and uses $RO(q)$ in order to answer query q. Suppose D is an oracle algorithm with auxiliary input accessing RO as follows: $D(z(A))$ emulates the interaction of $S^A(z(A))$ with Chal and for every oracle query q, it asks $RO(q)$ to obtain the answer $A(q)$ and continues the emulation. For any sampled random oracle RO let $z(RO) = z(A(RO))$ which denotes the nonuniform advice being computed directly over the random oracle. Suppose $\mathsf{Samp}(\cdot)$ is the pre-sampling algorithm of Lemma 4.3 when applied to the advice function z (computed over the random oracle) and the oracle algorithm D and parameter ℓ specified above.

[7] By padding the queries q appropriately, we can assume w.l.o.g that $RO(\cdot)$ is length preserving.

[8] The latter holds for an infinite sequence of n, but we also focus on the same sequence.

Experiment Hyb. The experiment Hyb follows the footsteps of the second experiment of Lemma 4.3. More formally we do the following:

1. Get $RO \xleftarrow{\$} \mathbf{RO}$ and obtain $z = z(RO)$.
2. Run $\mathsf{Samp}(z)$ to get the pre-sampled part of the oracle F of size s.
3. Get $RO' \xleftarrow{\$} \mathbf{RO}[F]$.
4. Run the interaction between Chal and $S^{A[RO']}(z)$.

Experiment Eff. It is the same experiment as Hyb with the only difference that in the very last step, we run the interaction between Chal and $S^{B[F]}(z)$ where $B[F]$ is the *efficient nonuniform* oracle of Lemma 4.9 defined as follows: upon any query q answered in the partial function F it returns the answer that A would return based on F (which is known through nonuniform advice), and over any other query q it executes the randomized algorithm B and returns $B(q)$. Theorem 4.10 follows from the following two claims.

Claim 4.11. *The joint view of the parties* (Chal, S) *in experiments* Inef *and* Hyb *are* $\varepsilon/2$*-close.*

Claim 4.11 implies that $\Pr_{\mathsf{Hyb}}[\mathsf{Chal} \text{ accepts}] \geq \varepsilon/2$.

Claim 4.11 follows directly from Lemma 4.3 and the definition of the distinguishing algorithm D (as the combination of the interactive algorithms Chal and S).

Claim 4.12. *The joint view of the parties* (Chal, S) *in experiments* Hyb *and* Eff *are computationally indistinguishable.*

Concluding Theorem 4.10. Claim 4.11 and Claim 4.12 show that $\Pr_{\mathsf{Eff}}[\mathsf{Chal} \text{ accepts}] \geq \varepsilon/2 - \mathrm{negl}(n)$, and since the adversary B in Eff is an efficient circuit, we obtain Theorem 4.10.

Proving Claim 4.12. Note that the first 3 steps of Hyb and Eff (including the sampling of F) are identical in experiments Hyb and Eff. Therefore, if a distinguisher can distinguish between the experiments by advantage δ, by an averaging argument we can fix (F, z) to a value that makes the distinguisher still achieve advantage δ. By hardwiring A's answers that are determined by F into the code of S, and by combining the code of S and Chal with that of the distinguisher, we get a polynomial sized circuit that violates the second property of Lemma 4.9. □

5 Refuting Nonuniform Reductions for Fully-Black-Box Constructions

In this section we prove our results of Theorems 1.4, 1.5, and 1.6. We start by proving Theorem 1.5 and then will prove Theorems 1.4 and 1.6 based on that.

We also show how the techniques used in the proofs of these theorems can be used to obtain different proofs of nonuniform hardness for other primitives such as one-way permutations as well as extending known black-box separations in the uniform regime to the nonuniform regime.

5.1 Nonuniform Collision-Resistance of Random Hashing and Beyond

Here we prove Theorem 1.5 about the nonuniform hardness of random functions as families of collision-resistant hash functions. In fact we prove the following stronger theorem:

Theorem 5.1 (Nonuniform Collision Resistance)**.** *Let the function* $\mathbf{h}\colon [K]\times[M]\mapsto[N]$ *be chosen uniformly at random. If for* $d,t\in\mathbb{N}$ *it holds that* $dN^3<2t^5K$*, then for every* t*-query adversary* A *with* d *bits of advice* z *that may depend on* h*:* $A^h(i,z)$ *can find a collision pair for the hash function* $h_i(\cdot)=h(i,\cdot)$ *with probability at most* $3t^2/N$ *over the choice of* $h\xleftarrow{\$}\mathbf{h}, i\xleftarrow{\$}[K]$ *and the randomness of* A*.*

Instead of proving Theorem 5.1, in fact we prove a more general lemma about the nonuniform hardness of *families* of ideal primitives, if the ideal primitive is already (uniformly) secure against bounded-query adversaries.

Lemma 5.2 (Hardness of Families of Primitives)**.** *Suppose* $\mathbf{O}$ *is a randomized oracle such that any computationally unbounded* t*-query adversary* A *"wins" in the "game"* $A^{\mathbf{O}}$ *only with probability at most* ε *where the notion of winning only depends on the transcript of* $A^{\mathbf{O}}$*. Suppose* FO *is another randomized oracle which consists of* $K=2^k$ *independent samples* $O_1,\ldots,O_K$ *from* $\mathbf{O}$ *accessed through a* k*-bit prefix to the queries. We define* $A^{\mathsf{FO}}(i)$ *wins iff the transcript of the interaction of* A *with* O_i *indicates a win for* A*. Suppose* A *also receives* d *bits of advice* z *about the oracle* FO*. Then for any such advice function, and any* $s\le K$*:*

$$\Pr_{i\xleftarrow{\$}[K],\mathsf{FO}\xleftarrow{\$}\mathbf{FO}}[A^{\mathsf{FO}}(i,z(\mathsf{FO}))\ \textit{wins}]\le\varepsilon+s/K+\sqrt{td/2s}.$$

The high level idea behind the proof of Lemma 5.2 is to use Lemma 4.2 as follows. Each sampled oracle O_i could be considered as an answer returned by a random oracle over domain $[K]$. By Lemma 5.2 the oracle answers to most of the indexes in $[K]$ remain statistically close to uniform, even given a "small" advice about all of $\mathsf{FO}=[O_1,\ldots,O_K]$. So for a "typical" $i\xleftarrow{\$}[K]$, the job of the adversary to win against O_i is just as hard as the case there were no advice.

Proof. Even though the oracle O_i might be highly structured (e.g., a permutation), as a mental experiment, we can pretend that each oracle O_i of FO is sampled by first sampling a (huge) uniform random string R_i of length m, and then obtaining O_i based on R_i. Let $\mathbf{RO}$ be a random oracle from $[K]$ to $[2^m]=[M]$. Thus, any sample $RO\xleftarrow{\$}\mathbf{RO}$ determines a sample $\mathsf{F}O\xleftarrow{\$}\mathsf{FO}$. Let $z(RO)$ be the value of the advice function over the corresponding $\mathsf{F}O$.

The intuition is that by Lemma 4.3, the adversary A, even after getting the advice z will get "small information" about most of the "oracle answers" $RO(i)=R_i$; therefore, the oracle O_i remains statistically close to a fresh sample and thus the upper bound on the winning probability follows even in the presence of the advice.

More formally suppose B is an imaginary algorithm that emulates A, and whenever A makes a query to O_i, B reals *all* of the answers of the oracle O_i. By Lemma 4.3 it follows that the the view of $B^{\mathbf{RO}}$ even after in the presence of d bits of advice z about its oracle and asking up to t oracle queries (of the form $RO(i)$ to get R_i) remains $\sqrt{td/2s}$-close to another experiment in which after obtaining the advice z, we run the pre-sampling procedure of Lemma 4.3 to obtain F as part of RO, re-sample the rest of the oracle RO independently of z and F, and then execute $B(z,i)$.

Since A only observes parts of the view of B we conclude that the view of $A^{\mathsf{FO}}(i, z(\mathsf{FO}))$ is also $\sqrt{td/2s}$-close to another experiment in which $K - s$ oracles among $O_1, \ldots, O_K$ are re-sampled after A receives z. Therefore, if the index $i \xleftarrow{\$} [K]$ happens to point to a fresh re-sampled oracle O_i, by the assumption of Lemma 5.2, A can win with probability at most $\sqrt{td/2s} + \varepsilon$. Finally note that i will point to one of the pre-sampled oracles O_i only with probability s/K. □

Now we prove Theorem 5.1 based on Lemma 5.2.

of Theorem 5.1. Suppose h is a random hash function from $[M]$ to $[N]$. Let A be any t query adversary trying to find a collision in h. We claim that the probability of A winning is at most t^2/N. The reason is that the i'th query of A will collide with one of its previous queries with probability at most $(i-1)/N$, and in case none of its queries collide, if A outputs any other query blindly it will give a collision with one of its previously asked queries only with probability $1/N$, making the maximum chance of A to win to be at most $1/N + \sum_{i \in [t]}(i-1)/N \leq t^2/N$.

Suppose $\varepsilon = t^2/N$ and $s = K\varepsilon$ (i.e. $\varepsilon = s/K$). If $\sqrt{td/2s} \leq \varepsilon$, we can apply Lemma 5.2 to obtain Theorem 5.1. But $\sqrt{ts/2d} \leq \varepsilon$ is indeed equivalent to $dN^3 \leq 2Kt^5$ which is the assumption of Theorem 5.1. □

In the following, we will first describe and prove a lemma for ruling out nonuniform security reductions, and then will use this lemma to prove Theorems 1.4 and 1.6.

5.1.1 Beyond Hash Functions

Note that by Lemma 5.2, a *family* of $\mathbf{O}$ can be used to obtain a *nonuniformly* secure oracles implementing primitive $\mathsf{F}\mathcal{P}$ where $\mathsf{F}\mathcal{P}$ is a *family* version of the primitive $\mathcal{P}$. If the family version of $\mathcal{P}$ happens to be a legitimate implementation of the original primitive $\mathcal{P}$ (which as we will discuss below is the case for almost all natural cryptographic primitives!) relative to a family of $\mathbf{O}$ we obtain a nonuniformly secure construction of $\mathcal{P}$. As we will see in Section 5.2 this allows one to prove separations from $\mathcal{P}$ handling nonuniform proofs of security; we will apply this technique for the specific case of FCRHs in Section 5.3.

More formally, we can define the family version of any primitive as follows.

Definition 5.3 (Family of a Primitive)**.** For a cryptographic primitive $\mathcal{P}$ by $\mathsf{F}\mathcal{P}$ we denote another primitive called the *family* of $\mathcal{P}$ which is defined as follows.

- All the algorithms inside $\mathsf{F}\mathcal{P}$ get an extra auxiliary input z (called the family index) and output the same z as well (but z can be used during the computation of the function implementing $\mathsf{F}\mathcal{P}$.) For every $\mathsf{F}P$ which implements $\mathsf{F}\mathcal{P}$, if we fix any z as the family index (and remove the redundant represented z from the input and the output) we denote the result residual function as $\mathsf{F}P[z]$. We also use $\mathsf{F}P(m)$ for $m \in N$ to refer to $\mathsf{F}P[z]$ where the only restriction is that $|z| = m$.
- For every z we require $\mathsf{F}P[z]$ to be an implementation of $\mathcal{P}$.
- The security game of $\mathsf{F}P(m)$ is defined based on the security game of $\mathcal{P}$ as follows. In the security game of $\mathsf{F}\mathcal{P}$ the challenger also selects a $z \xleftarrow{\$} \{0,1\}^m$, the adversary Adv gets the sampled z, and Adv is supposed to break the security of $\mathsf{F}P[z]$ as an implementation of $\mathcal{P}$, where his probability of winning is taken also over the random choice of $z \xleftarrow{\$} \{0,1\}^m$.

Any primitive which is already "in the family form" (e.g., FCRHs[9]) satisfies the requirement of Definition 5.3 trivially.

Theorem 5.4. *Let $\mathcal{P}$ be a cryptographic primitive which is either of: one-way functions, one-way permutations, or public-key encryption. Then any implementation of the* family *of $\mathcal{P}$, is also an implementation of $\mathcal{P}$ itself.*

Proof.

- **One-Way Functions.** In this case any implementation of $\mathsf{F}\mathcal{P}$ will take as input $z \in \{0,1\}^m$ and $x \in \{0,1\}^n$ and outputs both z and $f(z,x) \in \{0,1\}^n$. In this case, one can interpret this function as a new function $g(\cdot)$ defined over $n+m$ bits as $g(z,x) = (z, f(z,x))$.
- **One-Way Permutations.** Interestingly, the same transformation $g(z,x) = (z, f(z,x))$ described above for the case of one-way function makes g a permutation assuming that f is a permutation for every fixed auxiliary input z.
- **Public-Key Encryption.** In this case z can be selected at random by the key-generation algorithm and be part of the public-key. All the encryption and decryption algorithms also receive z, use it, and output it. The new scheme is trivially a public-key scheme itself.

□

Nonuniformly Hard Randomized Oracles for More Primitives. The same way that Theorem 5.1 was obtained from Lemma 5.2, we can find a randomized oracle relative to which nonuniformly hard version of primitives that include their own family exist (e.g., the primitives of Theorem 5.4). In the following we present such argument for case of one-way permutations which is quite different than that of [GT00] and uses a single application of Lemma 4.3.

Theorem 5.5. *There is an oracle relative to which nonuniformly hard one-way permutations exist.*

Proof. We claim that a random *family* of permutations $\mathbf{f}\colon \{0,1\}^{5n} \mapsto \{0,1\}^n$ (where $f(z,\cdot)$ is a random permutation for every $f \xleftarrow{\$} \mathbf{f}$ and $z \in \{0,1\}^{4n}$) is nonuniformly hard to invert. This can be proved using Lemma 5.2 and a simple observation that any t query attacker can invert a random permutation over $[N]$ only with probability $O(t/N)$. However, a nice point about any family of permutations such as p is that $p(z,x) = (z, f(z,x))$ is itself a permutation! Therefore, the oracle $f \xleftarrow{\$} \mathbf{f}$ is indeed a nonuniformly-hard (to invert) one-way permutation with overwhelming probability. □

5.2 A Lemma for Ruling Out Nonuniform Security Reductions

Variants of the following lemma in the *uniform* regime has been used in some previous work [BMG07, DLMM11, MP12].

Lemma 5.6. *Let $\mathcal{P}$ and $\mathcal{Q}$ be two cryptographic primitives and $\mathcal{P}$ has security threshold zero and let $\mathbf{O}$ be a randomized oracle. Suppose the following two holds:*

1. *For* any *black-box construction $Q^{\mathbf{O}}$ of $\mathcal{Q}$ from $\mathbf{O}$ there is an adversary Adv who asks $q(n)$ oracle queries to $\mathbf{O}$ and breaks the security of $Q^{\mathbf{O}}_n$ (over security parameter n) with non-negligible probability $\varepsilon > 1/\operatorname{poly}(n)$.*

[9] Note that hash functions without the index (that can handle uniform adversaries only) do not include their family version as special case.

2. *There exists a black-box construction* $P^{\mathbf{O}}$ *of* $\mathcal{P}$ *from* $\mathbf{O}$ *such that for any* $m = n^{\Theta(1)}$ *any* $\mathrm{poly}(n)(q(n)+1)$*-query adversary* T *who receives* $\mathrm{poly}(n)$ *bits of advice about* $\mathbf{O}$ *can break* $P_m^{\mathbf{O}}$ *only with negligible probability* $\mathrm{negl}(n)$.

Then there is no black-box construction of $\mathcal{Q}$ *from* $\mathcal{P}$ *even with a nonuniform proof of security. Moreover, to rule out constructions in which* Q_n *only calls* P_n*, we only need to consider* $m = n$ *in second condition above.*

Proof. Suppose Q is a black-box construction of $\mathcal{Q}$ from $\mathcal{P}$ which has a nonuniform proof of security. Namely, for every adversary A breaking the security of Q_n^P (i.e. the security of Q^P over the security parameter n) with advantage $\varepsilon > 1/\mathrm{poly}(n)$, there is a $\mathrm{poly}(m/\varepsilon)$-sized circuit S such that $S^{P,A}$ breaks the security of P_m for some $m = n^{\Theta(1)}$.

We would like to use the security reduction to turn the assumed q-query attacker Adv against $\mathbf{Q} = Q^{\mathbf{P}}$ (where $\mathbf{P} = P^{\mathbf{O}}$ is the randomized implementation of $\mathcal{P}$ using $\mathbf{O}$) into a nonuniform attack against $\mathbf{P}$ itself and derive a contradiction, but in order to do so we first have to *fix* the oracle $\mathbf{Q}$.

Recall that the advantage ε of Adv to break $\mathbf{Q}$ is also over the random choice of $O \xleftarrow{\$} \mathbf{O}$. By an averaging argument, with probability at least $\varepsilon/2$ over the choice of $O \xleftarrow{\$} \mathbf{O}$, O has the property that Adv still breaks Q^{P^O} with advantage at least $\varepsilon/2$. We call such $O \xleftarrow{\$} \mathbf{O}$ a good oracle. For a good oracle $O \xleftarrow{\$} \mathbf{O}$ the security of the implementation Q^{P^O} of $\mathcal{Q}$ is broken by Adv^O with advantage at least $\varepsilon/2 \geq 1/\mathrm{poly}(n)$. Therefore, there is a $\mathrm{poly}(n)$-sized circuit S such that S^{P^O,Adv^O} breaks the security of P_m^O with a nonnegligible probability $\delta(m) > 1/\mathrm{poly}(m)$ for $m = n^{\Theta(1)}$.

Consider the following algorithm T: it receives the description of (S, m, n) as advice, and then runs S^{P^O,Adv^O} to break the security of P_m^O. T has the following properties.

1. Since S has size $\mathrm{poly}(m) = \mathrm{poly}(n)$, T receives at most $\mathrm{poly}(n)$ bits of advice.
2. Since P is efficient, S is $\mathrm{poly}(n)$ sized, and Adv_n asks only $q(n)$ oracle queries, T asks only $\mathrm{poly}(n) + \mathrm{poly}(n) \cdot q(n) \leq \mathrm{poly}(n) \cdot (q(n))$ queries to $\mathbf{O}$.
3. If we run T while accessing a *randomized* $\mathbf{O}$, it still breaks the security of $P_m^{\mathbf{O}}$ with advantage at least $(\varepsilon/2) \cdot \delta > 1/\mathrm{poly}(n)$, because **(a)** $\varepsilon/2 \geq 1/\mathrm{poly}(n)$ fraction of the choices of $O \xleftarrow{\$} \mathbf{O}$ are good, **(2)** for every good O, T breaks P_m^O with probability $\delta > 1/\mathrm{poly}(m)$, and **(3)** $\mathcal{P}$ has security threshold 0.

The above three properties of T show that it violates the second assumption of Lemma 5.6. □

5.3 Proving Theorems 1.4 and 1.6

Here we prove Theorems 1.4 and 1.6. Roughly speaking Theorem 1.4 can be concluded from Theorem 5.1 similar to the way we used the result of [GT00] to extend the result of [IR89] to the nonuniform regime. Formal proof is as follows. In the rest of this section we prove Theorem 1.4. In the following we prove Theorem 1.4.

of Theorem 1.4. We employ Lemma 5.6 as follows. We let the randomized oracle $\mathbf{O}$ be the random oracle. Theorem 5.1 shows that there is a construction of FCRHs relative to $\mathbf{O}$ such that any adversary T who gets $2^{o(n)}$ bits of advice about $\mathbf{O}$ and asks $2^{o(n)}$ queries to it can break it only with advantage $2^{-\Omega(n)}$. This would imply the second requirement of Lemma 5.6 with $q(n) = \mathrm{poly}(n) \leq 2^{n/10}$.

We also use the following result which is the main step toward separating one-way functions from key-agreement proved in [IR89] and provides us with the adversary Adv as needed for the first condition of Lemma 5.6.

Lemma 5.7. *Suppose* Π *is a key-agreement protocol in which Alice and Bob each ask* n *oracle queries to the random oracle* $O \xleftarrow{\$} \mathbf{RO}$*, and they agree on a key with probability* $1-\mathrm{negl}(n)$*. There is a computationally unbounded adversary* Adv *who only accesses the public messages sent between Alice and Bob, asks at most* $\mathrm{poly}(n)$ *oracle queries to* O*, and finds the key with probability* $1-1/n^2$*.*

Since both conditions of Lemma 5.6 are satisfied, Theorem 1.4 follows immediately.

□

Now we prove Theorem 1.6.

of Theorem 1.6. We will first prove Theorem 1.6 for PRGs and then will prove it for the case o digital signatures.

Part 1: Pseudorandom Generators. Suppose on the contrary that $G\colon \{0,1\}^{\ell} \mapsto \{0,1\}^{\ell+k}$ is a black-box construction of PRGs that stretches its input by k and $p = \lambda(n)\cdot(k/\log n)$ oracle queries while $1/\mathrm{poly}(n) \le \lambda(n) \le o(1)$. In this case we again use $\mathbf{O}$ to denote a random oracle, but we use a slightly different construction of FCRHs h (to be used as $\mathcal{H}$) relative to $\mathbf{O}$. Given any query (z,x) where $z \in \{0,1\}^{4n}$ and $x \in \{0,1\}^{2n}$ to choose the n-bit output we choose the last $r = \min \log^2 n, \frac{\log n}{10\lambda(n)} = \omega(\log n)$ bits of y uniformly at random, and we copy the $n - \frac{\log n}{\lambda(n)}$ first bits of x as the $n - \frac{\log n}{\lambda(n)}$ first bits of the output y.

Note that any adversary breaking collision resistance of h has to find a collision for the last r bits of h as well, so as a mental experiment we pretend that the output length of h is $r \le \log^2 n$ Since $\mathrm{poly}(n)\cdot 2^{\log^2 n} \le \mathrm{poly}(n)\cdot 2^{4n}$, Theorem 5.1 proves that the construction $\mathbf{h} = h^{\mathbf{O}}$ is nonuniformly secure in the sense that any $\mathrm{poly}(n)$ query attacker with $\mathrm{poly}(n)$ bits of advice about $\mathbf{O}$ can find a collision in $h^{\mathbf{O}}$ only with negligible probability. This shows that the second requirement of Lemma 5.6 holds for $q(n) = \mathrm{poly}(n)$.

Lemma 5.8 ([GGKT05])**.** *There is an adversary* Adv *who asks* no *queries to* $\mathbf{O}$ *but is able to distinguish the output of* $G(U_\ell)$ *from* $U_{\ell+k}$ *with advantage* $1/4$*.*

We provide a proof for sake of completeness.

Proof. The total number of random bits used in the computation of $G^h(\mathbf{U}_\ell)$ is $\ell + p\cdot\frac{\log n}{10\lambda(n)} < \ell + k$ (where $\mathbf{U}_\ell$ is a random string of length ℓ). Therefore the support set of the function $G^h(\mathbf{U}_\ell)$ has size at most $2^{\ell+k-1}$. The latter implies that a computationally unbounded adversary is able to distinguish between $G^h(\mathbf{U}_\ell)$ and $\mathbf{U}_{\ell+k}$ with advantage at least $1/4$ (by outputting 1 whenever the given y is in the support of $G^h(\mathbf{U}_\ell)$ and outputting 0 elsewhere). □

The attacker Adv of Lemma 5.8 satisfies the first requirement of Lemma 5.6. Therefore, the first part of Theorem 1.6 follows directly from Lemma 5.6.

Part 2: Signature Schemes. Now we prove the second part of Theorem 1.6. Suppose for sake of contradiction that D is a black-box construction of digital signatures for message space $\{0,1\}^n$ using $o(n)$ queries to $\mathcal{H}_n$. We use the same construction h for $\mathcal{H}_n$ as that of the proof of Theorem 1.4; namely a random FCRHs with key length $4n$, input length $2n$, and output length $2n$. By Theorem 5.1, $h^{\mathbf{O}}$ is $2^{n/10}$ secure, even if the adversary gets $2^{n/10}$ bits of advice about $\mathbf{O}$. Thus we obtain the first requirement of Lemma 5.6 for any $q(n) = 2^{o(n)}$.

Lemma 5.9 ([BMG07])**.** *For any construction of digital signatures D for message space $\{0,1\}^n$ in the random oracle model in which the key-generation, signing, and verification algorithms ask p oracle queries can be broken with advantage $1/\operatorname{poly}(n)$ by an adversary Adv who asks $\operatorname{poly}(n)2^{O(p)}$ oracle queries.*

Lemma 5.9 shows that if D asks only $o(n)$ many queries to $\mathcal{H}_n$, then Adv can break it with $q(n) = 2^{o(n)}$ queries. This implies that the first requirement of Lemma 5.6 holds as well and so Lemma 5.6 implies the second part of Theorem 1.6 directly. □

A Case of Infinitely Often Security

In this we extend theorem 1.4 to the case of infinitely often security. Namely we rule out black-box constructions of key-agreement from one-way permutations with a nonuniform black-box proof of security, even if the security is defined w.r.t. infinitely often notion of security. We first formalize such black-box security reductions.

Definition A.1. We say that a (free or black-box) construction Q of a primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ has a nonuniform (black-box) infinitely-often secure reduction, if for every implementation P of $\mathcal{P}$, and every adversary A breaking the security of $Q(P)$ (as an implementation of $\mathcal{Q}$) for all security parameters $n > n_0$ with advantage $\varepsilon \geq 1/\operatorname{poly}(n)$ for some constant n_0, there is some constant m_0 and a family of $\operatorname{poly}(n/\varepsilon)$-sized oracle circuit S (whose code might depend on (P, A)) such that $S^{P,A}$ breaks the security of P_m for any $m > m_0$. A *uniform* security reduction is defined similarly by requiring the code of $S^{P,A_n}(1^n, 1^{1/\varepsilon})$ to be uniform and independent of the choices of the oracles (P, A).

In the rest of this section we extend Theorem 1.4 to handle even security reductions according to Definition A.1.

Definition A.2. We call a randomized FCRH $\mathcal{H}$ $u(n)$-strongly hard for output length n, if for every unbounded adversary Adv who asks $u(n)$ queries to $\mathcal{H}$ and gets $u(n)$ bits of advice about (all of) $\mathcal{H}$, the probability that Adv can find a collision of output length n is at most $1/u(n)$. We call $\mathcal{H}$ $u(\cdot)$-strongly hard, if the set of output lengths that over which $\mathcal{H}$ is not $u(n)$-strongly hard is finite.

Lemma A.3. *With probability one over the choice of a length-preserving random oracle $O \leftarrow \mathbf{RO}$, relative to O there exists a family of collision-resistant hash functions $\mathcal{H} = \{h_d \colon \{0,1\}^{2n} \mapsto \{0,1\}^n \mid d \in \{0,1\}^{4n}\}_n$ that is $2^{n/4}$-strongly-hard.*

Proof. For $d \in \{0,1\}^n, x \in \{0,1\}^{2n}$ define $h_d(x) = O(d,x)|_n$ where $y|_n$ denotes the last n bits of y. Note that $\mathcal{H} = \{h_d\}$ is indeed a random hash function.

We apply Theorem 5.1 by the parameters $k = 4n, m = 2n, t = 2^{n/4} = d$. Note that **(1)** $\varepsilon = t^2/N = 2^{-n/2}$ and so $3\varepsilon < 2^{-n/4}$ (for $n > 8$) and **(2)** $d2^{3n} < 2 \cdot 2^{4n}2^{5n/4} = 2Kt^5$. In fact, by an averaging argument, with probability at least $1 - 32^{-n/4}$, the sampled h is $2^{n/4}$-secure.

For every $n_0 \in \mathbb{N}$, as a mental experiment we can sample O_n (which determines $\mathcal{H}_n$) first for every $n \neq n_0$. Since the whole sampled $O = \{O_n \mid n \neq n_0\}$ can be thought of as the description of a *fixed* adversary trying to find collision in $\mathcal{H}_{n_0}$, we conclude that with probability $1 - O(2^{-n/4})$ it holds that $\mathcal{H}_n$ is $2^{n/4}$ *strongly* hard.

Suppose E_m is the "bad" event that for some $n \geq m$ the sampled $\mathcal{H}_n$ is not $2^{n/4}$-strongly-hard. Also let E be the event that the family of hash functions $\mathcal{H}$ defined relative to O is not strongly collision-resistant in the sense that there is an infinite sequence of input lengths $\mathcal{N} = \{n_1 < n_2 \ldots\}$ such that $\mathcal{H}_n$ is not $2^{n/4}$-strongly-hard for all $n \in \mathcal{N}$. Note that $E \subseteq E_m$ holds for all $m \in \mathbb{N}$ because any $X \in E$ (determined by a full sample of O) is also $X \in E_m$ by definition. Therefore, $\Pr[E] \leq \Pr[E_m]$ for all $m \in \mathbb{N}$. By a union bound, it holds that $\Pr[E_m] \leq \sum_{n \geq m} 2^{-n/4} \leq O(2^{-m/4})$. Thus, for every $\delta < 1$ it holds that $\Pr[E_m] \leq \delta$ for sufficiently large m and so $\Pr[E] \leq \delta$.

□

of Theorem 1.4 for Infinitely Often Security. Using Lemmas 5.7 and A.3 we prove Theorem 1.4. We start by assuming that there is a black-box construction $\Pi = (I, S)$ of key-agreement protocols from FCRHs. By Lemma 5.7 there is a fixed $\text{poly}(n)$-query strategy adversary Adv that breaks any key-agreement relative to a random oracle $O \xleftarrow{\$} \mathbf{RO}$ with probability $1 - 1/n^2$. By a Markov argument, it holds that with probability $1 - 2/n^2$ over the choice of $O \leftarrow \mathbf{RO}$, the adversary breaks the key-agreement over security parameter n with probability at least $1/2$ in which case we say that the adversary succeeds for security parameter n. By the same union bound argument as that of the proof of Lemma A.3 we can conclude that: the probability of adversary not succeeding for an infinite sequence $n \in \{n_1 < n_2 \ldots\}$ is smaller than any constant $\delta > 0$. Therefore, with probability one, the sampled $O \xleftarrow{\$} \mathbf{RO}$ has the property that the constructed key-agreement relative to O can be broken over sufficiently large n. We call any such sample $O \leftarrow \mathbf{RO}$ a good one.

By definition, relative to any good O, the nonuniform security reduction S, turns the unbounded adversary Adv into another adversary Adv' who gets $\text{poly}(n)$ advice about O and breaks the collision-resistance of the hash function $\mathcal{H}$ for an infinite sequence of output lengths $n \in \{n'_1 < n'_2 \ldots\}$. Therefore, with probability one, the sampled $\mathcal{H}$ is *not* strongly secure. This contradicts the conclusion of Lemma A.3. □

References

[AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz, *On basing one-way functions on np-hardness*, Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 701–710. 5

[AS11] Sergei Artemenko and Ronen Shaltiel, *Lower bounds on the query complexity of nonuniform and adaptive reductions showing hardness amplification*, Electronic Colloquium on Computational Complexity (ECCC) **18** (2011), 16. 3

[Blu87] Manuel Blum, *How to prove a theorem so no one else can claim it*, Proceedings of the International Congress of Mathematicians, 1987, pp. 1444–1451. MR 91h:68141 8

[BMG07] Boaz Barak and Mohammad Mahmoody-Ghidary, *Lower bounds on signatures from symmetric primitives*, IEEE Symposium on Foundations of Computer Science (FOCS), 2007. 1, 4, 24, 27

[BMM09] Boaz Barak and Mohammad Mahmoody-Mahmoody, *Merkle puzzles are optimal - an $O(n^2)$-query attack on any key exchange from a random oracle*, Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings (Shai Halevi, ed.), Lecture Notes in Computer Science, vol. 5677, Springer, 2009, pp. 374–390. 1

[BMV08] Emmanuel Bresson, Jean Monnerat, and Damien Vergnaud, *Separation results on the "one-more" computational problems*, CT-RSA (Tal Malkin, ed.), Lecture Notes in Computer Science, vol. 4964, Springer, 2008, pp. 71–87. 5

[BPR+08] Dan Boneh, Periklis A. Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters, *On the Impossibility of Basing Identity Based Encryption on Trapdoor Permutations*, FOCS, 2008, pp. 283–292. 1

[Bra79] Gilles Brassard, *Relativized cryptography*, Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, 1979, pp. 383–391. 5

[BT06] Andrej Bogdanov and Luca Trevisan, *On worst-case to average-case reductions for np problems*, SIAM Journal on Computing **36** (2006), no. 4, 1119–1159. 5

[BU08] Michael Backes and Dominique Unruh, *Limits of constructive security proofs*, ASIACRYPT (Josef Pieprzyk, ed.), Lecture Notes in Computer Science, vol. 5350, Springer, 2008, pp. 290–307. 3

[BV98] Boneh and Venkatesan, *Breaking RSA may not be equivalent to factoring*, EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT, 1998. 5, 15

[DDN00a] Danny Dolev, Cynthia Dwork, and Moni Naor, *Nonmalleable cryptography*, SIAM Journal on Computing **30** (2000), no. 2, 391–437 (electronic), Preliminary version in STOC 1991. MR 1 769 364 10

[DDN00b] Danny Dolev, Cynthia Dwork, and Moni Naor, *Nonmalleable cryptography*, SIAM Journal on Computing **30** (2000), no. 2, 391–437. 12

[DLMM11] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin, *On black-box complexity of optimally-fair coin-tossing*, Theory of Cryptography Conference - TCC 2011, 2011. 1, 24

[DOP05] Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak, *On the generic insecurity of the full domain hash*, Lecture Notes in Computer Science (2005), 449–?? 6

[DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani, *Time space tradeoffs for attacks against one-way functions and PRGs*, Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings (Tal Rabin, ed.), Lecture Notes in Computer Science, vol. 6223, Springer, 2010, pp. 649–665. 4

[FF93] Joan Feigenbaum and Lance Fortnow, *Random-self-reducibility of complete sets*, SIAM Journal on Computing **22** (1993), no. 5, 994–1005. 5

[FS87] Amos Fiat and Adi Shamir, *How to prove yourself: practical solutions to identification and signature problems*, Proceedings on Advances in cryptology—CRYPTO '86 (London, UK), Springer-Verlag, 1987, pp. 186–194. 12

[FS10] Marc Fischlin and Dominique Schröder, *On the impossibility of three-move blind signature schemes*, EUROCRYPT (Henri Gilbert, ed.), Lecture Notes in Computer Science, vol. 6110, Springer, 2010, pp. 197–215. 5

[GGK03] Gennaro, Gertner, and Katz, *Lower bounds on the efficiency of encryption and digital signature schemes*, STOC: ACM Symposium on Theory of Computing (STOC), 2003. 4

[GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan, *Bounds on the efficiency of generic cryptographic constructions*, SIAM J. Comput. **35** (2005), no. 1, 217–246. 1, 4, 26

[GKM+00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan, *The Relationship between Public Key Encryption and Oblivious transfer*, FOCS, 2000, pp. 325–335. 1, 4

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff, *The knowledge complexity of interactive proof systems*, SIAM Journal on Computing **18** (1989), no. 1, 186–208. 1, 8

[GMR01] Yael Gertner, Tal Malkin, and Omer Reingold, *On the Impossibility of Basing Trapdoor Functions on Trapdoor Predicates*, FOCS, 2001, pp. 126–135. 1

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson, *Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems*, Journal of the ACM **38** (1991), no. 1, 691–729. 1, 8

[GO94] Oded Goldreich and Yair Oren, *Definitions and properties of zero-knowledge proof systems*, Journal of Cryptology **7** (1994), no. 1, 1–32. 1

[Gol93] Oded Goldreich, *A uniform-complexity treatment of encryption and zero-knowledge*, Journal of Cryptology **6** (1993), no. 1, 21–53. 1

[Gol04] _____, *Foundations of cryptography: Basic applications*, Cambridge University Press, 2004. 8

[GT00] Rosario Gennaro and Luca Trevisan, *Lower Bounds on the Efficiency of Generic Cryptographic constructions*, FOCS, 2000, pp. 305–313. 3, 4, 24, 25

[GW11] Craig Gentry and Daniel Wichs, *Separating succinct non-interactive arguments from all falsifiable assumptions*, STOC (Lance Fortnow and Salil P. Vadhan, eds.), ACM, 2011, pp. 99–108. 1, 5, 6, 19

[Hås90] Johan Håstad, *Pseudo-random generators under uniform assumptions*, Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC), 1990, pp. 387–394. 1

[HH09] Iftach Haitner and Thomas Holenstein, *On the (im)possibility of key dependent encryption*, Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings (Omer Reingold, ed.), Lecture Notes in Computer Science, vol. 5444, Springer, 2009, pp. 202–219. 4, 6

[HHRS07] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev, *Finding collisions in interactive protocols – A tight lower bound on the round complexity of statistically-hiding commitments*, Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, 2007. 1, 4

[HRS09] Iftach Haitner, Alon Rosen, and Ronen Shaltiel, *On the (im)possibility of arthur-merlin witness hiding protocols*, Theory of Cryptography, Fourth Theory of Cryptography Conference, TCC 2009, 2009. 5

[ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby, *Pseudo-random generation from one-way functions*, Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC), 1989, pp. 12–24. 1

[IR89] Russell Impagliazzo and Steven Rudich, *Limits on the provable consequences of one-way permutations*, Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC), ACM Press, 1989, pp. 44–61. 1, 3, 4, 25, 26

[KM12] Neal Koblitz and Alfred Menezes, *Another look at non-uniformity*, IACR Cryptology ePrint Archive **2012** (2012), 359, informal publication. 3

[KST99] Jeong Han Kim, Daniel R. Simon, and Prasad Tetali, *Limits on the efficiency of one-way permutation-based hash functions*, Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS), 1999, pp. 535–542. 1

[KSY11] Jonathan Katz, Dominique Schröder, and Arkady Yerukhimovich, *Impossibility of blind signatures from one-way permutations*, TCC (Yuval Ishai, ed.), Lecture Notes in Computer Science, vol. 6597, Springer, 2011, pp. 615–629. 1

[LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam, *Concurrent non-malleable commitments from any one-way function*, TCC (Ran Canetti, ed.), Lecture Notes in Computer Science, vol. 4948, Springer, 2008, pp. 571–588. 1, 3, 9, 10, 11, 12

[LTW05] Lin, Trevisan, and Wee, *On hardness amplification of one-way functions*, Theory of Cryptography Conference (TCC), LNCS, vol. 2, 2005. 1

[Lu09] Chi-Jen Lu, *On the security loss in cryptographic reductions*, EUROCRYPT (Antoine Joux, ed.), Lecture Notes in Computer Science, vol. 5479, Springer, 2009, pp. 72–87. 3

[M11] Takahiro Matsuda 0002 and Kanta Matsuura, *On black-box separations among injective one-way functions*, TCC (Yuval Ishai, ed.), Lecture Notes in Computer Science, vol. 6597, Springer, 2011, pp. 597–614. 1

[MP12] Mohammad Mahmoody and Rafael Pass, *The curious case of non-interactive commitments - on the power of black-box vs. non-black-box use of primitives*, CRYPTO (Reihaneh Safavi-Naini and Ran Canetti, eds.), Lecture Notes in Computer Science, vol. 7417, Springer, 2012, pp. 701–718. 24

[Nao03] Moni Naor, *On cryptographic assumptions and challenges*, Lecture Notes in Computer Science **2729** (2003), 96–109. 6

[Pas06] Rafael Pass, *Parallel repetition of zero-knowledge proofs and the possibility of basing cryptography on NP-hardness*, IEEE Conference on Computational Complexity, IEEE Computer Society, 2006, pp. 96–110. 5

[Pas11] _____, *Limits of provable security from standard assumptions*, Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011 (Lance Fortnow and Salil P. Vadhan, eds.), ACM, 2011, pp. 109–118. 1, 3, 5, 6, 8, 9, 11, 13, 15, 16, 17

[PTV11] Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkitasubramaniam, *Towards non-black-box lower bounds in cryptography*, Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings (Yuval Ishai, ed.), Lecture Notes in Computer Science, vol. 6597, Springer, 2011, pp. 579–596. 1, 5

[RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan, *Notions of reducibility between cryptographic primitives.*, Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Lecture Notes in Computer Science, vol. 2951, Springer, 2004, pp. 1–20. 2, 7

[RV10] Guy N. Rothblum and Salil P. Vadhan, *Are PCPs inherent in efficient arguments?*, Computational Complexity **19** (2010), no. 2, 265–304. 6

[Sch91] Claus-Peter Schnorr, *Efficient signature generation by smart cards*, J. Cryptology **4** (1991), no. 3, 161–174. 8

[Sim98] Daniel R. Simon, *Finding Collisions on a One-Way Street: Can Secure Hash Functions Be Based on General Assumptions?*, EUROCRYPT, 1998, pp. 334–345. 1

[SV10] Ronen Shaltiel and Emanuele Viola, *Hardness amplification proofs require majority*, SIAM J. Comput **39** (2010), no. 7, 3122–3154. 3

[Unr07] Dominique Unruh, *Random oracles and auxiliary input*, CRYPTO (Alfred Menezes, ed.), Lecture Notes in Computer Science, vol. 4622, Springer, 2007, pp. 205–223. 4, 16, 17

[Vah10] Yevgeniy Vahlis, *Two Is a Crowd? A Black-Box Separation of One-Wayness and Security under Correlated Inputs*, TCC, 2010, pp. 165–182. 1