

New Negative Results on Differing-Inputs Obfuscation

Mihir Bellare¹, Igors Stepanovs², and Brent Waters³

¹ Department of Computer Science and Engineering, University of California San Diego, USA. <https://cseweb.ucsd.edu/~mihir/>

² Department of Computer Science and Engineering, University of California San Diego, USA. <https://cseweb.ucsd.edu/~istepano/>

³ Department of Computer Science, University of Texas at Austin, USA. <https://www.cs.utexas.edu/~bwaters/>

Abstract. We provide the following negative results for differing-inputs obfuscation (diO): (1) If sub-exponentially secure one-way functions exist then sub-exponentially secure diO for TMs does not exist (2) If in addition sub-exponentially secure iO exists then polynomially secure diO for TMs does not exist.

1 Introduction

Differing-inputs obfuscation (diO) is a natural extension of indistinguishability obfuscation (iO). It has been conjectured that candidate constructions of iO also met diO. Based on this, diO has been exploited in applications. Garg, Gentry, Halevi and Wichs (GGHW) [28] showed that if something they called “special purpose” obfuscation exists, then diO does not. This has put diO in an ambiguous and contentious position, some people arguing that GGHW is evidence diO does not exist, others saying that perhaps it does and it is special-purpose obfuscation that does not exist. This paper uses a new approach to give powerful evidence that the first camp is right, meaning it is indeed diO that does not exist, by showing this to be true under weaker and more standard assumptions than special-purpose obfuscation. We show (1) If sub-exponentially secure one-way functions exist then sub-exponentially secure diO for TMs does not exist (2) If in addition sub-exponentially secure iO exists then polynomially secure diO for TMs does not exist.

1.1 Background

The notion of program obfuscation that is most intuitive and appealing is that an obfuscated program should be no more useful than an oracle for the program itself. Formalized as VBB obfuscation (vbbO), it was shown impossible in the sense that there is no obfuscator that will successfully VBB obfuscate all programs [36, 7]. Further negative results about vbbO were given in [33, 17]. In the face of this, Barak et al. [7] suggested other, weaker notions of obfuscation

that appeared not to succumb to their counter-examples and might therefore be achievable. The most prominent were indistinguishability obfuscation (iO) and its extension, differing-input obfuscation (diO). The first asks that obfuscations of functionally equivalent programs are indistinguishable. The second is a natural computational relaxation: even if the programs are not functionally equivalent, as long as it is hard, given the programs, to find an input on which they differ, then the obfuscations of the programs are indistinguishable. The underlying intuition is that if one can find a differing input for the programs, one can clearly distinguish their obfuscations. In iO this is excluded theoretically, by saying there does not exist such an input, while in diO it is excluded computationally, by saying such an input might exist but is hard to find. On the surface both might appear equally reasonable, since the vbbO negative results do not apply to either. But this turns out not to be true.

These intriguing notions lay dormant for many years, for two reasons. First, that one could not prove these notions unachievable did not mean they were achievable. Second, they seemed quite weak; even if they were achievable, what could one do with them? An answer to the first question came with candidate constructions of iO [27, 6, 43, 30]. An answer to the second came when Sahai and Waters showed how to use iO towards many ends [45]. Since then, applications of iO and diO have ballooned.

In these applications, a crucial role is played by *auxiliary information*. The modern definitions of iO and diO used in these applications [27, 45, 1, 20, 12] consider a program sampler S that spits out a pair P_0, P_1 of programs *together with associated auxiliary information aux* . The sampler is said to produce functionally equivalent programs if P_0 and P_1 agree on all inputs. The sampler is said to be difference-secure if an adversary given P_0, P_1, aux cannot find an input x such that $P_0(x) \neq P_1(x)$ except with small probability. The obfuscation game picks a challenge bit b and gives you (the adversary) an obfuscation \bar{P}_b of P_b under the obfuscator Obf , *together with aux* . Your task (as the adversary) is to guess b . Obf is called iO-secure if you have small advantage for all samplers producing functionally equivalent programs, and diO-secure if you have small advantage for all difference-secure samplers. Adversaries are always polynomial time, but probabilities referred to as “small” may be sub-exponentially so or negligible. Programs may be TMs or circuits. This leads to a collection of variant notions.

1.2 The GGHW result

Let Obf be an obfuscator. GGHW [28] provide a program sampler S for which they show, under certain assumptions, that diO-security of Obf fails, which means that (1) the sampler is difference secure under these assumptions, but (2) there is a way to distinguish the obfuscations under Obf of the two programs returned by the sampler given the auxiliary information. Their approach is to have the sampler first generate a signing and verification key pair (sk, vk) for a signature scheme meeting the standard notion of unforgeability [34]. The program P_1 takes a message m and candidate signature σ and accepts iff σ is a valid signature on m under vk . The program P_0 will take in the same inputs, but it

will always reject. Clearly the programs P_0 and P_1 differ exactly on the input pairs (m, σ) where σ is a valid signature of m under vk . Next, the sampler creates a third program P_2 that has hardwired the secret signing key sk and takes as input a (smaller) program \bar{P} . It hashes \bar{P} using a CRHF to get a message m , and uses sk to get a signature σ on m . It then runs the \bar{P} on (m, σ) and outputs 1 if \bar{P} accepts on these inputs. Finally, S creates auxiliary information aux consisting of an obfuscation P_2^* of P_2 . This obfuscation is not under the given obfuscator Obf , but under some other assumed “special purpose” obfuscator Obf^* whose role and properties will emerge in the following.

To serve as a counterexample it should both (1) be possible, using the auxiliary information P_2^* , to distinguish between obfuscations under Obf of P_0 and P_1 , and (2) be difficult, given P_0, P_1, P_2^* , to find an input on which P_0 and P_1 differ. The first property follows trivially from the design. An adversary given the auxiliary information P_2^* and a challenge program \bar{P} that is either an obfuscation of P_0 or P_1 can distinguish these cases by simply feeding the program \bar{P} as an input to P_2^* . If \bar{P} is an obfuscation of P_1 then, when P_2^* runs \bar{P} on the message and valid signature that P_2^* creates, \bar{P} will accept. But if \bar{P} is an obfuscation of P_0 , then P_2^* will reject.

In contrast it is much more difficult to establish the second property, namely that it is hard to find an input on which P_0, P_1 differ *even in the presence of the auxiliary information* P_2^* . The difficulty stems from the latter. In the absence of aux the property follows straightforwardly from the security of the signature scheme, as a differing input is exactly a valid message-signature pair, and would amount to a signature forgery. However, since the obfuscated differentiating program P_2^* has embedded in it the secret signing key sk it is not clear how to prove that it is hard to find signatures in the presence of P_2^* .

Recall that P_2^* was an obfuscation, under some un-specified obfuscator Obf^* , of P_2 . GGHW [28] simply conjecture that there exists some obfuscator Obf^* that will hide the secret key sk sufficiently well that it is hard to find a differing input for P_0, P_1 , meaning to find a valid message-signature pair, even given P_2^* . While they were unable to prove this conjecture under any standard obfuscation definitions such as iO or even $vbbO$, they were able to partially justify their conjecture with a heuristic analysis. Their analysis replaces the adversary’s access to the obfuscated program P_2^* with an oracle that performs the same functionality. In this world the adversary no longer has direct access to an object containing sk and GGHW are able to demonstrate differing inputs security of S by a fairly straightforward reduction to the underlying security of the signature scheme.

The GGHW result certainly creates significant questions regarding the use of diO . Arguably, the primary reason for using the diO security definition over $vbbO$ is that no impossibility results like those of [36, 7, 33, 17] are known for diO . However, if the GGHW conjecture holds, then this is no longer true and the perceived benefit of diO versus $vbbO$ is significantly reduced. (The benefit is not eliminated, since even if there exist functionalities that cannot be diO obfuscated, it is still possible that there are functionalities that can be diO obfuscated but not VBB obfuscated.) At the same time, the heuristic used to

justify the GGHW counterexample is itself much stronger than assuming diO — namely their analysis relies on modeling the differentiating obfuscated program as an oracle.

1.3 Our approach

We introduce a new approach to proving the impossibility of diO. In contrast to the prior work, we analyze our sampler under *concrete assumptions* that replace the GGHW conjecture. We now explain the intuition behind our approach as well as the obstacles we had to overcome.

Let Obf be an obfuscator that we assume, towards a contradiction, is diO-secure. At the highest level our approach is similar to GGHW. We build a program sampler \mathcal{S} that produces programs P_0, P_1 and auxiliary information P_2^* consisting of an obfuscation of a program P_2 under an obfuscator Obf^* . As in GGHW, the sampler generates a signing and verification key pair (sk, vk) for an underlying signature scheme DS , and program P_0 always rejects. Likewise, P_1 takes as input a candidate message-signature pair (m, σ) and checks its validity under the signature verification program DS.Ver with key vk . The auxiliary information continues to be the obfuscation P_2^* , under an obfuscator Obf^* , of a program \bar{P}_2 , where P_2 hardwires the secret signing key sk . \bar{P}_2 takes as input a program \bar{P} of a certain maximum length, and uses $m = \bar{P}$ as the message it signs, and runs \bar{P} on m and the signature, accepting if this accepts. The important difference now however is that Obf^* is not some new type of obfuscator as in GGHW. Rather Obf^* is *assumed to be only an iO-secure obfuscator*.

It continues to be easy, using the auxiliary information P_2^* , to distinguish between obfuscations under Obf of P_0 and P_1 . The main issue is to prove that it is difficult, given P_0, P_1, P_2^* , to find an input on which P_0 and P_1 differ. The hurdle here continues to be the same, namely that the auxiliary information program P_2^* embeds the secret signing key sk . This precludes reducing to the security of the signature scheme in an obvious way. To prove security we will show that it is computationally difficult to generate a signature on any message. We do this via a hybrid argument that steps through every possible message one by one. Since our hybrid steps through the entire message space we base our security on assumptions of sub-exponential hardness.

To execute our strategy we will replace the generic signature scheme of GGHW with a special type of puncturable signature scheme that we call a *consistent puncturable* signature scheme. Given a “master” secret key sk , it should be possible to create a punctured version sk_{m^*} of the key, for a given message m^* , that can be used to sign any message $m \neq m^*$ but even given which it is hard to produce a signature on m^* . So far this is a special type of policy-based [9], functional [21] or delegatable [5] signatures, these themselves analogues of the notions of puncturable, constrained and functional PRFs [19, 40, 21]. The additional consistency requirement is that the signatures of $m \neq m^*$ produced under the master key and the punctured key should be the same. Note that only deterministic puncturable signature schemes can be consistent, but the former is not a sufficient condition. We show in Section 3 that such signature schemes

can be built from iO and one-way functions. While making a standard signature scheme deterministic is trivial via the use of PRFs, our challenge is making the punctured and master versions of the key produce consistent signatures.

Our hybrid now proceeds as follows. We step through each program (message) P^* and show that it is computationally difficult to produce a signature on P^* . We do this by first replacing the obfuscation of P_2 with an obfuscation of a program P_{2,P^*} that works as follows. On all inputs $P \neq P^*$ the program P_{2,P^*} behaves as P_2 with the exception that it uses a punctured version of the signing key sk_{P^*} . On input P^* its output is hardwired to be whatever the output of $P_2(P^*)$ was. We observe that if indistinguishability obfuscation holds, then no poly-time attacker can distinguish between obfuscations of programs P_2 and P_{2,P^*} . This follows since the two programs share the same output on every input. On every $P \neq P^*$ the master and punctured keys will produce the same signature that they feed into P , and on input P^* program P_{2,P^*} is hardwired to behave the same as P_2 . Since it is hard to distinguish between obfuscations of these two programs, it should be no easier to output a signature on message P^* when P_2 is obfuscated to get the auxiliary information aux than it is when P_{2,P^*} is obfuscated. However, in the latter case the security of the puncturable signature scheme guarantees this is hard.

Note that since we assumed a diO-secure obfuscator Obf to start our proof by contradiction, an iO-secure obfuscator, which we use both directly and to build consistent punctured signatures, is provided for free and is not an extra assumption. This means the only assumption we need is a sub-exponentially hard one-way function. More precisely, this is the case for sub-exponential diO, while for polynomial diO the iO assumption will be an extra one.

While the text above outlines our main approach, there are several important factors that still must be taken into account. First, we notice that P_1 should be capable of verifying a signature on a message that is an obfuscation of P_1 and thus longer than P_1 itself. For this reason we need to view P_0 and P_1 as Turning Machines (TMs) that can process inputs longer than their own descriptions.

Next, our complexity leveraging argument requires that the advantage ϵ of any PT attacker on the signature scheme multiplied by the message space be negligible. To satisfy this using sub-exponential hardness assumptions we must use a verification key vk that is larger than the programs P_0, P_1 . However, this creates a circularity problem under the obvious strategy of having P_1 actually contain vk to verify the messages! We circumvent this issue by the use of a UOWHF [42], also called a target collision-resistant (TCR) hash function [10], that hashes a separate verification program as follows. We construct a program P_{ver} that takes as input a candidate message-signature pair (m, σ) and uses an embedded verification key vk to either accept or reject it. Now P_1 takes P'_{ver} as an additional input and uses it to check the candidate message-signature pairs, rather than storing vk and performing the verification itself. P_1 hardwires the hash h of P_{ver} under a TCR hash function, and rejects unless the hash h' of P'_{ver} matches its hardwired hash h . This ensures that only P_{ver} can be used to verify the signatures. We analyze security by adding a hybrid step at the beginning

using the UOWHF security. We emphasize that the argument using our UOWHF is outside of the complexity leveraging part of our hybrid.

The above is a very high-level description, and the devil is in the details that the body of the paper sorts out. The circularity issues, summarized via Fig. 6, have to be dealt with very carefully. A critical element of dealing with them is that *different primitives are run with different values of the security parameter*. Thus, while the convention is that the security parameter in a proof remains λ throughout, our constructions will feature $n(\lambda)$ as the security parameter in certain places, with n a polynomial that is carefully defined based on other parameters. Another subtlety is that the success of this program depends on the details of how sub-exponential security is defined. Specifically (cf. Section 2) we use “uniform” rather than “pointwise” definitions in the language of [8]. The latter showed them equivalent in the usual setting of negligible functions but they are not known to be equivalent in the sub-exponential setting.

1.4 Discussion and related work

Sub-exponential security. Our assumptions and conclusions both involve sub-exponential hardness and one might ask about the validity of such assumptions and the value of such conclusions. Empirical evidence, at least, says that when problems are hard, they are sub-exponentially hard. Natural problems do not appear to be polynomially but not sub-exponentially hard except in rare cases [3]. Indeed sub-exponential hardness is frequently assumed in cryptography, especially recently [35, 31, 24]. In particular it is unlikely that polynomially-secure diO exists but sub-exponentially secure diO does not, so ruling out the latter is significant in terms of evidence against diO. Similarly it is unlikely that polynomially secure OWFs exist but sub-exponentially secure ones do not, so assuming the latter is reasonable.

Bounded versus unbounded inputs. In this work we provide negative results about the existence of differing-inputs obfuscators for TMs that can take arbitrarily long inputs. Our results do not rule out the possibility of constructing diO for TMs with a-priori bounded input-lengths.

Implications. Note that [1, 20] build diO for TMs with unbounded inputs from circuit diO and SNARKs [15, 14]. This means that if SNARKs exist then our negative results for TM diO extend to circuit diO. Also [38] build diO for TMs with unbounded inputs from public-coin diO for NC^1 , fully homomorphic encryption with decryption in NC^1 and public-coin SNARKs. Our results would imply that (if the assumptions we make hold) one of these three primitives does not exist.

Constructions and applications of diO. Differing-inputs obfuscation has proven to be a powerful tool using which we have built new primitives. In some cases it has later been possible to reduce the assumption to iO or other diO variants, but sometimes at the cost of weakening the conclusion and usually at the cost of increased complexity and difficulty. All this motivates understanding whether or not diO is achievable.

diO for circuits is used in [1, 20] to achieve adaptively-secure FE (Functional Encryption) and extractable witness encryption. It was later shown in [23, 16, 41] how to build TM iO from circuit iO but the conclusion was weaker. Adaptively-secure FE from iO did emerge but the solutions were more complex than the ones from diO [47, 2].

Boyle et al. [20] show that iO implies diO for samplers outputting circuits that differ on only polynomially-many inputs. Our counter-examples and results do not apply to this type of diO. Differing input obfuscation is used as a tool in [12], via the result of [20], to give hardcore functions with polynomially-many output bits from any injective one-way function and iO, and is used as an assumption to extend this result to arbitrary one-way functions. It is used similarly as a tool in [22].

Ishai et al. [38] define public-coin diO, by relaxing the notion of diO to require that only public random coins can be used to build challenge programs and the corresponding auxiliary information. Our negative results do not apply to public-coin diO. Public-coin diO is a valuable notion but it doesn't take away from the interest in proving impossibility of diO because diO precedes public-coin diO and there are works that still rely on it, and there could be interesting new applications from diO but not from public-coin diO. Furthermore, our techniques might help understand the possibility of public-coin diO.

A variant of diO was also used as an assumption in a result in [29].

Consistent signature schemes. Some of the prior work focuses on constructing digital signature schemes with properties that are similar to the ones we require from consistent signature schemes. The known primitives include: functional signatures [21], policy-based signatures [9] and operational signatures [4], the latter subsuming the preliminary work on delegatable signatures [5]. However, none of the proposed constructions of these primitives satisfy the consistency requirement which requires that the master and punctured signing keys produce the same signatures for all messages except for the punctured message, and which is crucial for our impossibility result.

We get consistent puncturable signatures from OWFs and iO, which in our context effectively means from OWFs since our proof assumes diO towards a contradiction and thus gets iO for free. Our definition of consistent puncturable signatures is novel, but our construction follows Sahai-Waters signatures [45]. Consistent puncturable signatures are also implied by splittable signatures [41], which are built based on an injective PRG and iO. Injective PRGs are not known to be implied by OWFs so the assumption is stronger than ours. However, [18] build injective OWFs from OWFs and iO, and also say that, due to an observation of Boyle et al. [20], the injective PRG of [41] can be replaced with an injective OWF. By this route one can get consistent puncturable signatures from OWFs and iO. However our construction is direct, substantially simpler and self contained. Consistent puncturable signatures can also be constructed from constrained verifiable PRFs [26, 25]. The latter are achievable from κ -Multilinear DDH assumption. In our context, this would be an additional assumption since it is not known to be implied by diO.

Game $\text{OW}_F^{\mathcal{F}}(\lambda)$	Game $\text{TCR}_H^{\mathcal{H}}(\lambda)$	Game $\text{PPRF}_G^{\mathcal{G}}(\lambda)$
$fk \leftarrow \mathcal{F}.\text{Kg}(1^\lambda)$	$(x_0, st) \leftarrow \mathcal{H}_1(1^\lambda)$	$b \leftarrow \{0, 1\}; gk \leftarrow \mathcal{G}.\text{Kg}(1^\lambda)$
$x \leftarrow \mathcal{F}.\text{In}(\lambda)$	$hk \leftarrow \mathcal{H}.\text{Kg}(1^\lambda)$	$b' \leftarrow \mathcal{G}^{\text{CH}}(1^\lambda); \text{Return } (b = b')$
$y \leftarrow \mathcal{F}.\text{Ev}(1^\lambda, fk, x)$	$x_1 \leftarrow \mathcal{H}_2(1^\lambda, st, hk)$	$\text{CH}(x^*)$
$x' \leftarrow \mathcal{F}(1^\lambda, fk, y)$	$h_0 \leftarrow \mathcal{H}.\text{Ev}(1^\lambda, hk, x_0)$	$gk^* \leftarrow \mathcal{G}.\text{PKg}(1^\lambda, gk, x^*)$
$y' \leftarrow \mathcal{F}.\text{Ev}(1^\lambda, fk, x')$	$h_1 \leftarrow \mathcal{H}.\text{Ev}(1^\lambda, hk, x_1)$	If $b = 1$ then $r^* \leftarrow \mathcal{G}.\text{Ev}(1^\lambda, gk, x^*)$
Return $(y = y')$	$\text{win}_0 \leftarrow (x_0 \neq x_1)$	else $r^* \leftarrow \mathcal{G}.\text{Out}(\lambda)$
	$\text{win}_1 \leftarrow (h_0 = h_1)$	Return (gk^*, r^*)
	Return $(\text{win}_0 \wedge \text{win}_1)$	

Fig. 1. Games defining one-wayness of function family \mathcal{F} , target collision-resistance of function family \mathcal{H} and puncturable-PRF security of function family \mathcal{G} .

2 Preliminaries

Notation. Let $\mathbb{N} = \{0, 1, 2, \dots\}$ be the set of non-negative integers. We denote by $\lambda \in \mathbb{N}$ the security parameter and by 1^λ its unary representation. If $x \in \{0, 1\}^*$ is a string then $|x|$ denotes its length. If $x \in \{0, 1\}^*$ is a string and $\ell \in \mathbb{N}$ such that $|x| \leq \ell$ then $\langle x \rangle_\ell$ denotes the string of length ℓ that is built by padding x with leading zeros. If X is a finite set, we let $x \leftarrow X$ denote picking an element of X uniformly at random and assigning it to x . Algorithms may be randomized unless otherwise indicated. Running time is worst case. “PT” stands for “polynomial-time,” whether for randomized algorithms or deterministic ones. If A is an algorithm, we let $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs x_1, \dots and assigning the output to y . We let $y \leftarrow A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs x_1, \dots . We say that $f: \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every positive polynomial p , there exists $\lambda_p \in \mathbb{N}$ such that $f(\lambda) < 1/p(\lambda)$ for all $\lambda \geq \lambda_p$. We use the code based game playing framework of [11]. (See Fig. 1 for an example.) By $\mathcal{G}^{\mathcal{A}}(\lambda)$ we denote the event that the execution of game \mathcal{G} with adversary \mathcal{A} and security parameter λ results in the game returning **true**.

Uniform and pointwise security definitions. There are two common ways to formalize security definitions – by using different order of quantification. Let GAME be a security game, and let $\text{Adv}_{\mathcal{A}}^{\text{game}}(\lambda)$ be the advantage of a PT adversary \mathcal{A} winning in this game with security parameter λ . Consider the following two alternative definitions of sub-exponential security. A *uniform* definition requires that there is a constant $0 < \epsilon < 1$ such that for every PT adversary \mathcal{A} there exists $\lambda_{\mathcal{A}} \in \mathbb{N}$ such that $\text{Adv}_{\mathcal{A}}^{\text{game}}(\lambda) \leq 2^{-\lambda^\epsilon}$ for all $\lambda \geq \lambda_{\mathcal{A}}$. A *pointwise* definition requires that for every PT adversary \mathcal{A} there exist $0 < \epsilon < 1$ and $\lambda_{\mathcal{A}} \in \mathbb{N}$ such that $\text{Adv}_{\mathcal{A}}^{\text{game}}(\lambda) \leq 2^{-\lambda^\epsilon}$ for all $\lambda \geq \lambda_{\mathcal{A}}$. These definitions differ in the order of quantification between ϵ and \mathcal{A} . In this work, we use uniform security definitions. For the case of polynomial security, Bellare [8] proved that uniform and

pointwise definitions are equivalent. It is not known whether the equivalence also holds for the above definitions of sub-exponential security.

Circuits and Turing Machines. We say that P is a program if it is either a circuit or a Turing Machine (TM), and we denote the size of its binary representation by $|P|$. We assume that any program P takes a single input string x ; if P is defined to take multiple inputs x_1, \dots then running P on an input x is implicitly assumed to parse $(x_1, \dots) \leftarrow x$ and run $P(x_1, \dots)$.

We say that circuits C_0, C_1 are functionally equivalent, written $C_0 \equiv C_1$, if they have the same number of inputs $\ell \in \mathbb{N}$ and if $C_0(x) = C_1(x)$ holds for all $x \in \{0, 1\}^\ell$. We say that TMs M_0, M_1 are functionally equivalent, and denote it by $M_0 \equiv M_1$, if both $M_0(x)$ and $M_1(x)$ halt on all $x \in \{0, 1\}^*$ and if $M_0(x) = M_1(x)$ for all $x \in \{0, 1\}^*$.

If M is a TM and $t \in \mathbb{N}$ then $y \leftarrow \text{UTM}_M^t(x_1, \dots)$ denotes running M on inputs x_1, \dots and assigning the output to y ; if $M(x_1, \dots)$ does not halt within t steps, then $\text{UTM}_M^t(x_1, \dots)$ returns 0. If M is a TM and $x \in \{0, 1\}^*$ is a string such that M halts on input x , we use $\text{time}(M, x)$ to denote the number of steps that are required for it to halt.

Let P be any circuit or any TM that halts on all inputs. For any $s \in \mathbb{N}$ such that $|P| \leq s$ let $\text{Pad}_s(P)$ denote P padded to have size s , meaning that $\text{Pad}_s(P)$ and P are of the same type (i.e. both are circuits or TMs) and $\text{Pad}_s(P) \equiv P$. We assume that P can be padded to any size larger or equal to $|P|$.

Function families. A family of functions F specifies PT algorithms $F.\text{Kg}$ and $F.\text{Ev}$, where $F.\text{Ev}$ is deterministic. Associated to F is a collection of input sets $F.\text{In}$ and a collection of output sets $F.\text{Out}$, defining all valid inputs and outputs for each of security parameters. Key generation algorithm $F.\text{Kg}$ takes 1^λ to return a key fk . Evaluation algorithm $F.\text{Ev}$ takes $1^\lambda, fk$ and an input $x \in F.\text{In}(\lambda)$ to return $F.\text{Ev}(1^\lambda, fk, x) \in F.\text{Out}(\lambda)$. We say that F is injective if the function $F.\text{Ev}(1^\lambda, fk, \cdot): F.\text{In}(\lambda) \rightarrow F.\text{Out}(\lambda)$ is injective for all $\lambda \in \mathbb{N}$ and $fk \in [F.\text{Kg}(1^\lambda)]$.

Puncturable function families. A *puncturable* function family G specifies (beyond the usual algorithms) additional PT algorithms $G.\text{PKg}$ and $G.\text{PEv}$, where $G.\text{PEv}$ is deterministic. Punctured key generation algorithm $G.\text{PKg}$ takes 1^λ , a key $gk \in [G.\text{Kg}(1^\lambda)]$ and a target input $x^* \in G.\text{In}(\lambda)$ to return a ‘‘punctured’’ key gk^* . Punctured evaluation algorithm $G.\text{PEv}$ takes $1^\lambda, gk^*$ and an input $x \in G.\text{In}(\lambda)$ to return $G.\text{PEv}(1^\lambda, gk^*, x) \in G.\text{Out}(\lambda)$. The correctness condition requires that $G.\text{PEv}(1^\lambda, gk^*, x) = G.\text{Ev}(1^\lambda, gk, x)$ for all $\lambda \in \mathbb{N}$, $gk \in [G.\text{Kg}(1^\lambda)]$, $x^* \in G.\text{In}(\lambda)$, $gk^* \in [G.\text{PKg}(1^\lambda, gk, x^*)]$ and $x \in G.\text{In}(\lambda) \setminus \{x^*\}$.

One-way functions. Consider game OW of Fig. 1 associated to a function family F and an adversary \mathcal{F} , where $F.\text{In}(\lambda)$ is required to be finite for all $\lambda \in \mathbb{N}$. For $\lambda \in \mathbb{N}$ let $\text{Adv}_{F, \mathcal{F}}^{\text{ow}}(\lambda) = \Pr[\text{OW}_{F, \mathcal{F}}^\lambda]$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. We say that F is δ -OW-secure if for every PT adversary \mathcal{F} there exists $\lambda_{\delta, \mathcal{F}} \in \mathbb{N}$ such that $\text{Adv}_{F, \mathcal{F}}^{\text{ow}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta, \mathcal{F}}$. We say that F is sub-exponentially OW-secure if it is $2^{-(\cdot)^\epsilon}$ -OW-secure for some $0 < \epsilon < 1$.

Target collision-resistant functions. Consider game TCR of Fig. 1 associated to a function family H and an adversary \mathcal{H} . For $\lambda \in \mathbb{N}$ let $\text{Adv}_{H,\mathcal{H}}^{\text{TCR}}(\lambda) = \Pr[\text{TCR}_H^{\mathcal{H}}(\lambda)]$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. We say that H is δ -TCR-secure if for every PT adversary \mathcal{H} there exists $\lambda_{\delta,\mathcal{H}} \in \mathbb{N}$ such that $\text{Adv}_{H,\mathcal{H}}^{\text{TCR}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta,\mathcal{H}}$. We say that H is sub-exponentially TCR-secure if it is $2^{-(\cdot)^\epsilon}$ -TCR-secure for some $0 < \epsilon < 1$. Target collision-resistant hash functions were introduced by Naor and Yung [42] under the name of Universal One-Way Hash Functions (UOWHF). [10] redefined the corresponding security notion under the name of *target collision-resistance*.

TCR-secure function families can be built from one-way functions, by combining the following results. First, [44, 37] (see also [39]) proposed constructions of TCR-secure *compression function* families with fixed input and output lengths. More formally, they show how to build a function family H' such that $H'.\text{In}(\cdot) = \{0, 1\}^{p_{\text{in}}(\cdot)}$ and $H'.\text{Out}(\cdot) = \{0, 1\}^{p_{\text{out}}(\cdot)}$, where $p_{\text{in}}, p_{\text{out}}$ are some polynomials such that $p_{\text{in}}(\lambda) \geq p_{\text{out}}(\lambda)$ for all $\lambda \in \mathbb{N}$. Next, [10, 46] showed how to use any TCR-secure compression function H' with fixed input length in order to build another TCR-secure function family H for arbitrary, bounded variable-length inputs, meaning that $H.\text{In}(\lambda) = \bigcup_{i \leq p(\lambda)} \{0, 1\}^i$ and $H.\text{Out}(\lambda) = H'.\text{Out}(\lambda)$ for some function $p: \mathbb{N} \rightarrow \mathbb{N}$ all $\lambda \in \mathbb{N}$.

Puncturable PRFs. Consider game PPRF of Fig. 1 associated to a puncturable function family G and an adversary \mathcal{G} , where $G.\text{Out}(\lambda)$ is required to be finite for all $\lambda \in \mathbb{N}$ and \mathcal{G} is required to make exactly one oracle query to CH. For $\lambda \in \mathbb{N}$ let $\text{Adv}_{G,\mathcal{G}}^{\text{PPRF}}(\lambda) = 2\Pr[\text{PPRF}_G^{\mathcal{G}}(\lambda)] - 1$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. We say that G is a δ -PPRF-secure if for every PT adversary \mathcal{G} there exists $\lambda_{\delta,\mathcal{G}} \in \mathbb{N}$ such that $\text{Adv}_{G,\mathcal{G}}^{\text{PPRF}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta,\mathcal{G}}$. We say that G is sub-exponentially PPRF-secure if it is $2^{-(\cdot)^\epsilon}$ -PPRF-secure for some $0 < \epsilon < 1$. Puncturable PRFs were concurrently and independently introduced in [19, 40, 21]. They can be built by extending the standard PRF construction of Goldreich, Goldwasser and Micali [32].

Digital signature schemes. A digital signature scheme DS defines PT algorithms DS.Kg , DS.Sig , DS.Ver , where DS.Ver is deterministic. Associated to DS is a collection of input sets DS.In and a collection of output sets DS.Out , defining all valid messages and signatures for each of security parameters. Key generation algorithm DS.Kg takes 1^λ to return a signing key sk and a verification key vk . Signing algorithm DS.Sig takes $1^\lambda, sk$ and a message $m \in \text{DS.In}(\lambda)$ to return a signature $\sigma \in \text{DS.Out}(\lambda)$. Verification algorithm DS.Ver takes $1^\lambda, vk, m, \sigma$ to return a decision $d \in \{1, 0\}$ regarding whether σ is a valid signature of m under vk , where 1 is returned if σ is a valid and 0 otherwise. The correctness condition requires that $\text{DS.Ver}(1^\lambda, vk, m, \sigma) = 1$ for all $\lambda \in \mathbb{N}$, $(sk, vk) \in [\text{DS.Kg}(1^\lambda)]$, $m \in \text{DS.In}(\lambda)$ and $\sigma \in [\text{DS.Sig}(1^\lambda, sk, m)]$. We say that a digital signature scheme DS is deterministic if its signing algorithm DS.Sig is deterministic.

Obfuscators. An obfuscator is a PT algorithm Obf that on input 1^λ and a program P returns a program \bar{P} of the same type as P such that $\bar{P} \equiv P$. We say that Obf is a circuit obfuscator if it obfuscates circuits, and we say that Obf is a

Game $\text{DIFF}_S^{\mathcal{D}}(\lambda)$	Game $\text{IO}_{\text{Obf},S}^{\mathcal{O}}(\lambda)$
$(P_0, P_1, aux) \leftarrow_S \mathcal{S}(1^\lambda)$	$b \leftarrow_{\mathcal{S}} \{0, 1\}; (P_0, P_1, aux) \leftarrow_S \mathcal{S}(1^\lambda)$
$x \leftarrow_{\mathcal{D}}(1^\lambda, P_0, P_1, aux)$	$\bar{P} \leftarrow_{\mathcal{S}} \text{Obf}(1^\lambda, P_b)$
Return $(P_0(x) \neq P_1(x))$	$b' \leftarrow_{\mathcal{S}} \mathcal{O}(1^\lambda, \bar{P}, aux); \text{Return}(b = b')$

Fig. 2. Games defining difference-security of program sampler \mathcal{S} and iO-security of program obfuscator Obf relative to program sampler \mathcal{S} .

TM obfuscator if it obfuscates TMs. Note that according to our definition of functionally equivalent programs, obfuscation is not defined for TMs that do not halt on some inputs. The polynomial slowdown condition requires that for every TM obfuscator Obf there is a polynomial $p: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that for every TM M that halts on all inputs and for every input $x \in \{0, 1\}^*$, we have $\text{time}(\bar{M}, x) \leq p(\lambda, \text{time}(M, x))$ for all $\lambda \in \mathbb{N}$ and $\bar{M} \in [\text{Obf}(1^\lambda, M)]$. An analogous slowdown condition trivially holds for any PT circuit obfuscator.

In this work, we discuss indistinguishability obfuscation (iO) and differing-inputs obfuscation (diO). The study of these obfuscation notions was initiated in [7]. Later [27, 45] showed how to build and use the former, whereas [20, 1] provided results on the latter. We extend the definitional framework of [12] that uses classes of program samplers to capture different variants of security notions for iO and diO. Specifically, our definitions allow for a unified treatment of polynomial and sub-exponential security of both circuit and TM obfuscation.

Program samplers. A circuit sampler is a PT algorithm $\mathcal{S}^{\text{circ}}$ that on input 1^λ returns a triple (C_0, C_1, aux) , where C_0, C_1 are circuits of the same size, number of inputs and number of outputs, and aux is a string. A TM sampler is a PT algorithm \mathcal{S}^{tm} that on input 1^λ returns a triple (M_0, M_1, aux) , where M_0, M_1 are TMs of the same size, and aux is a string. We require that $M_0(x)$ and $M_1(x)$ halt for all $\lambda \in \mathbb{N}$, $(M_0, M_1, aux) \in [\mathcal{S}^{\text{tm}}(1^\lambda)]$ and $x \in \{0, 1\}^*$. We say that \mathcal{S} is a program sampler if it is either a circuit sampler or a TM sampler.

Classes of program samplers. We say that a program sampler \mathcal{S} produces functionally equivalent programs if $\Pr[P_0 \equiv P_1 : (P_0, P_1, aux) \leftarrow_S \mathcal{S}(1^\lambda)] = 1$ for all $\lambda \in \mathbb{N}$. Let $\mathcal{S}_{\text{eq}}^{\text{circ}}$ be the class of all circuit samplers that produce functionally equivalent circuits, and let $\mathcal{S}_{\text{eq}}^{\text{tm}}$ be the class of all TM samplers that produce functionally equivalent TMs. Consider game DIFF of Fig. 2 associated to a program sampler \mathcal{S} and an adversary \mathcal{D} . For $\lambda \in \mathbb{N}$ let $\text{Adv}_{\mathcal{S}, \mathcal{D}}^{\text{diff}}(\lambda) = \Pr[\text{DIFF}_S^{\mathcal{D}}(\lambda)]$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. We say that \mathcal{S} is δ -DIFF-secure if for every PT adversary \mathcal{D} there exists $\lambda_{\delta, \mathcal{D}} \in \mathbb{N}$ such that $\text{Adv}_{\mathcal{S}, \mathcal{D}}^{\text{diff}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta, \mathcal{D}}$. We say that \mathcal{S} is sub-exponentially DIFF-secure if it is $2^{-(\cdot)^\epsilon}$ -DIFF-secure for some $0 < \epsilon < 1$. Let $\mathcal{S}_{\delta\text{-diff}}^{\text{circ}}$ be the class of all δ -DIFF-secure circuit samplers, and let $\mathcal{S}_{\delta\text{-diff}}^{\text{tm}}$ be the class of all δ -DIFF-secure TM samplers. Informally, difference-security of a program sampler \mathcal{S} means that given its output (P_0, P_1, aux) , it is hard to find an input on which the programs P_0 and P_1 differ.

Indistinguishability obfuscation and differing-inputs obfuscation. Consider game IO of Fig. 2 associated to an obfuscator Obf , a program sampler \mathbf{S} and an adversary \mathcal{O} . For $\lambda \in \mathbb{N}$ let $\text{Adv}_{\text{Obf},\mathbf{S},\mathcal{O}}^{\text{io}}(\lambda) = 2\Pr[\text{IO}_{\text{Obf},\mathbf{S}}^{\mathcal{O}}(\lambda)] - 1$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. Let \mathbf{S} be a class of program samplers. We say that Obf is δ - \mathbf{S} -secure if for every program sampler $S \in \mathbf{S}$ and for every PT adversary \mathcal{O} there exists $\lambda_{\delta,S,\mathcal{O}} \in \mathbb{N}$ such that $\text{Adv}_{\text{Obf},\mathbf{S},\mathcal{O}}^{\text{io}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta,S,\mathcal{O}}$. We say that Obf is sub-exponentially \mathbf{S} -secure if it is $2^{-(\cdot)^\epsilon}$ - \mathbf{S} -secure for some $0 < \epsilon < 1$.

We say that Obf is a sub-exponentially secure indistinguishability obfuscator for TMs (resp. circuits) if there exists $0 < \epsilon < 1$ such that Obf is $2^{-(\cdot)^\epsilon}$ - $\mathbf{S}_{\text{eq}}^{\text{tm}}$ -secure (resp. $2^{-(\cdot)^\epsilon}$ - $\mathbf{S}_{\text{eq}}^{\text{circ}}$ -secure). We say that Obf is a differing-inputs obfuscator for TMs (resp. circuits) if for every negligible function $\gamma: \mathbb{N} \rightarrow \mathbb{R}$ there exists a negligible function $\nu: \mathbb{N} \rightarrow \mathbb{R}$ such that Obf is ν - $\mathbf{S}_{\gamma\text{-diff}}^{\text{tm}}$ -secure (resp. ν - $\mathbf{S}_{\gamma\text{-diff}}^{\text{circ}}$ -secure). Note that ν - $\mathbf{S}_{\gamma\text{-diff}}^{\text{tm}}$ -security may be unachievable if there exists an infinite number of security parameters $\lambda \in \mathbb{N}$ such that $\gamma(\lambda) > \nu(\lambda)$. We say that Obf is a sub-exponentially secure differing-inputs obfuscator for TMs (resp. circuits) if for every $0 < \epsilon_0 < 1$ and $\gamma = 2^{-(\cdot)^{\epsilon_0}}$ there exists $0 < \epsilon_1 < 1$ such that Obf is $2^{-(\cdot)^{\epsilon_1}}$ - $\mathbf{S}_{\gamma\text{-diff}}^{\text{tm}}$ -secure (resp. $2^{-(\cdot)^{\epsilon_1}}$ - $\mathbf{S}_{\gamma\text{-diff}}^{\text{circ}}$ -secure).

Note that according to our definitions, a sub-exponentially secure differing-inputs obfuscator is not necessarily a polynomially-secure differing-inputs obfuscator. Namely, the former guarantees no security with respect to δ -DIFF-secure program samplers when δ is negligible but not sub-exponentially small. This observation can be used to strengthen our definition of sub-exponentially secure diO. We chose to use the weaker definition, which is simpler to define and which makes our impossibility results stronger.

3 Consistent puncturable digital signature schemes

We start by defining *consistent puncturable digital signature schemes* that will be used for our impossibility results in Section 4. Our construction follows Sahai-Waters signatures [45], and we prove its security assuming OWF and iO.

Informally, a puncturable digital signature scheme allows to ‘puncture’ its signing key sk at an arbitrary message m^* . The resulting punctured secret key sk^* , punctured at m^* , allows to produce signatures for all messages except for m^* . The puncturability property is similar to the one of puncturable PRFs. We say that a puncturable digital signature scheme is *consistent* if its secret signing key sk and every possible punctured signing key sk^* , that can be derived from sk , deterministically produce the same signatures for all messages except for the punctured message.

We now define a security notion, informally, requiring that no PT adversary should be able to forge a valid signature for the punctured message. The natural formalization of this security notion requires *selective* unforgeability, meaning that an adversary has to choose a message m^* at which the original signing key sk should be punctured. Having received the corresponding pair of punctured signing key sk^* and verification key vk , the goal of the adversary is to produce a valid signature for m^* with respect to the verification key.

Game $\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(\lambda)$

$(m^*, st) \leftarrow \mathcal{U}_1(1^\lambda); (sk, vk) \leftarrow \text{DS.Kg}(1^\lambda)$

$sk^* \leftarrow \text{DS.PKg}(1^\lambda, sk, m^*); \sigma^* \leftarrow \mathcal{U}_2(1^\lambda, st, vk, sk^*)$

$d \leftarrow \text{DS.Ver}(1^\lambda, vk, m^*, \sigma^*); \text{Return } (d = 1)$

Fig. 3. Game defining selective unforgeability of puncturable digital signature scheme DS under chosen message attack.

Puncturable digital signature schemes. A *puncturable* digital signature scheme DS specifies (beyond the algorithms associated to digital signatures schemes) additional PT algorithms DS.PKg, DS.PSig, where DS.PSig is deterministic. Punctured key generation algorithm DS.PKg takes 1^λ , a signing key $sk \in [\text{DS.Kg}(1^\lambda)]$ and a message $m^* \in \text{DS.In}(\lambda)$ to return a “punctured” signing key sk^* . Punctured signing algorithm DS.PSig takes $1^\lambda, sk^*$ and a message $m \in \text{DS.In}(\lambda)$ to return a signature $\sigma \in \text{DS.Out}(\lambda)$. We say that puncturable digital signature scheme DS is *consistent* if $\text{DS.Sig}(1^\lambda, sk, m) = \text{DS.PSig}(1^\lambda, sk^*, m)$ for all $\lambda \in \mathbb{N}$, $(sk, vk) \in [\text{DS.Kg}(1^\lambda)]$, $m^* \in \text{DS.In}(\lambda)$, $sk^* \in [\text{DS.PKg}(1^\lambda, sk, m^*)]$ and $m \in \text{DS.In}(\lambda) \setminus \{m^*\}$. Note that DS can be consistent only if it is deterministic. More precisely, both DS.Sig and DS.PSig should be deterministic. However, determinism is a necessary but not a sufficient condition.

Punctured selective unforgeability under chosen message attack. Consider game PSUFCMA of Fig. 3 associated to a puncturable digital signature scheme DS and an adversary \mathcal{U} . For $\lambda \in \mathbb{N}$ let $\text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(\lambda) = \Pr[\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(\lambda)]$. Let $\delta: \mathbb{N} \rightarrow \mathbb{R}$ be any function. We say that DS is δ -PSUFCMA-secure if for every PT adversary \mathcal{U} there exists $\lambda_{\delta, \mathcal{U}} \in \mathbb{N}$ such that $\text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(\lambda) \leq \delta(\lambda)$ for all $\lambda \geq \lambda_{\delta, \mathcal{U}}$. We say that DS is sub-exponentially PSUFCMA-secure if it is $2^{-(\cdot)^\epsilon}$ -PSUFCMA-secure for some $0 < \epsilon < 1$.

Our construction. We build a consistent puncturable digital signature scheme DS from a PPRF G, an indistinguishability obfuscator Obf and a OWF F. Our main observation is that a PPRF key gk can be used as a secret key for DS. In order to obtain a punctured key for DS, we puncture gk accordingly. The correctness condition of puncturable PRFs guarantees that DS is consistent. We build a verification key by obfuscating a circuit that embeds the PPRF key gk and a OWF key fk . The circuit takes a message-signature pair (m, σ) and returns 1 if $\text{F.Ev}(1^\lambda, fk, \sigma) = \text{F.Ev}(1^\lambda, fk, \text{G.Ev}(1^\lambda, gk, m))$; it returns 0 otherwise.

Puncturable digital signature scheme PUNC-DS. Let $s: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial. Let G be a puncturable function family. Let F be a function family such that $\text{F.In} = \text{G.Out}$. Let Obf be a circuit obfuscator. We build a consistent puncturable digital signature scheme $\text{DS} = \text{PUNC-DS}[\text{G}, \text{F}, \text{Obf}, s]$ as follows. Let $\text{DS.In}(\lambda) = \text{G.In}(\lambda)$ and $\text{DS.Out}(\lambda) = \text{G.Out}(\lambda)$ for all $\lambda \in \mathbb{N}$, and let Fig. 4 define the puncturable digital signature scheme DS. We say that DS is *well-defined* if $s(\lambda) \geq |\mathcal{C}_{1^\lambda, gk, fk}|$ for all $\lambda \in \mathbb{N}$, $gk \in [\text{G.Kg}(1^\lambda)]$ and $fk \in [\text{F.Kg}(1^\lambda)]$.

<u>Algorithm DS.Kg(1^λ)</u> $gk \leftarrow_s \mathbf{G.Kg}(1^\lambda)$; $fk \leftarrow_s \mathbf{F.Kg}(1^\lambda)$ $\overline{C} \leftarrow_s \mathbf{Obf}(1^\lambda, \mathbf{Pad}_s(\lambda)(C_{1^\lambda, gk, fk}))$ Return (gk, \overline{C}) <u>Circuit $C_{1^\lambda, gk, fk}(m, \sigma)$</u> $\sigma' \leftarrow \mathbf{G.Ev}(1^\lambda, gk, m)$ $y' \leftarrow \mathbf{F.Ev}(1^\lambda, fk, \sigma')$ If $(y' = \mathbf{F.Ev}(1^\lambda, fk, \sigma))$ then return 1 Else return 0	<u>Algorithm DS.PKg($1^\lambda, gk, m^*$)</u> Return $\mathbf{G.PKg}(1^\lambda, gk, m^*)$ <u>Algorithm DS.Ver($1^\lambda, \overline{C}, m, \sigma$)</u> Return $\overline{C}(m, \sigma)$ <u>Algorithm DS.Sig($1^\lambda, gk, m$)</u> Return $\mathbf{G.Ev}(1^\lambda, gk, m)$ <u>Algorithm DS.PSig($1^\lambda, gk^*, m$)</u> Return $\mathbf{G.PEv}(1^\lambda, gk^*, m)$
---	---

Fig. 4. Puncturable digital signature scheme $\text{DS} = \text{PUNC-DS}[\mathbf{G}, \mathbf{F}, \text{Obf}, s]$.

The following says that a PSUFCMA-secure, consistent punctured digital signature scheme can be built assuming OWF and iO.

Theorem 1. *Let \mathbf{G} be a sub-exponentially PPRF-secure function family such that $\mathbf{G.In}(\lambda), \mathbf{G.Out}(\lambda) \subseteq \bigcup_{i \leq p_0(\lambda)} \{0, 1\}^i$ for some polynomial p_0 and all $\lambda \in \mathbb{N}$. Let \mathbf{F} be a sub-exponentially OW-secure function family such that $\mathbf{F.In} = \mathbf{G.Out}$ and $\mathbf{F.Out}(\lambda) \subseteq \bigcup_{i \leq p_1(\lambda)} \{0, 1\}^i$ for some polynomial p_1 and all $\lambda \in \mathbb{N}$. Let Obf be a sub-exponentially $\mathbf{S}_{\text{eq}}^{\text{circ}}$ -secure circuit obfuscator. Then there is a polynomial $s: \mathbb{N} \rightarrow \mathbb{N}$ such that the following is true. Let $\text{DS} = \text{PUNC-DS}[\mathbf{G}, \mathbf{F}, \text{Obf}, s]$. Then (1) DS is well-defined, and (2) DS is sub-exponentially PSUFCMA-secure.*

In order to prove that DS is PSUFCMA-secure, we show that an adversary can not find the value of $\mathbf{G.Ev}(1^\lambda, gk, m^*)$ for a challenge message m^* , even given the obfuscated verification-key circuit that contains gk . In the proof, we puncture gk at m^* to get a punctured key gk^* , and construct a functionally equivalent verification-key circuit that embeds gk^* along with $y^* = \mathbf{F.Ev}(1^\lambda, fk, \mathbf{G.Ev}(1^\lambda, gk, m^*))$. The new verification key accepts σ as a valid signature for m^* if and only if $y^* = \mathbf{F.Ev}(1^\lambda, fk, \sigma)$, whereas the verification of signatures for all other messages $m \neq m^*$ remains the same. First, we use the iO-security of Obf to switch the verification circuits. Then we use the PPRF-security of \mathbf{G} , followed by the OWF-security of \mathbf{F} to show that no adversary can find the value of $\mathbf{G.Ev}(1^\lambda, gk, m^*)$ from gk^* and y^* . The proof is given in [13].

4 Impossibility of differing-inputs obfuscation for TMs

In this section we show that differing-inputs obfuscation for Turing Machines is impossible. In order to disprove sub-exponentially secure diO for TMs, we assume only the existence of sub-exponentially secure one-way functions. Furthermore, we show that polynomially secure diO for TMs is also impossible, additionally assuming sub-exponentially secure iO.

We construct a sub-exponentially difference-secure TM sampler, meaning that given a pair of TMs produced by this sampler it is hard to find an input on

which these TMs produce different outputs. The proof of difference-security is the core part of our work. It requires to carefully specify how to choose parameters for our sampler in a way that does not introduce any circular dependencies. Besides proving difference-security, we also show that there exists an adversary that can distinguish between obfuscations of TMs that are produced by the sampler regardless of the used obfuscator. Together these claims imply the impossibility of diO for TMs.

The blueprint for impossibility results. The first attack on differing-inputs obfuscation was presented by Garg, Gentry, Halevi and Wichs (GGHW) [28]. They introduced a novel *special-purpose obfuscation* assumption and showed that it contradicts diO. Our impossibility result follows the high-level idea from their work, but we achieve it using concrete assumptions. We now explain the core ideas of our impossibility result, which closely follow those of GGHW.

We construct a TM sampler S^{tm} that returns TMs M^0, M^1 along with an auxiliary information string aux . The sampler generates a key pair (sk, vk) for a digital signature scheme DS, and its output depends on these keys. TM M^0 returns 0 on every input. TM M^1 returns 1 if and only if it gets a valid message-signature pair as input, corresponding to the verification key vk ; it returns 0 otherwise. The auxiliary information string aux is an iO-obfuscation of a TM M^{aux} . The latter embeds the signing key sk and takes a TM \bar{M} as input, which for our purpose will normally be a diO-obfuscation of M^0 or M^1 . M^{aux} returns the result of running \bar{M} on a message-signature pair that is produced using its embedded signing key sk .

In order to determine whether a TM \bar{M} is an obfuscation of M^0 or M^1 , one can run M^{aux} with \bar{M} as input. According to the construction of M^{aux} , it will return $b \in \{0, 1\}$ if and only if \bar{M} is an obfuscation of M^b . To prove difference-security of S^{tm} , we will show that it is hard to find a valid message-signature pair given (M^0, M^1, aux) . The main technical challenge of the proof is to show that aux (the obfuscation of M^{aux}) properly hides the embedded signing key sk , which does not naturally follow from the security of indistinguishability obfuscation.

Turing Machine sampler TM-SAMP. Let $s_0, \ell, n, t_0, t_1, s_1: \mathbb{N} \rightarrow \mathbb{N}$ be polynomials. Let $\text{Obf}_{\text{eq}}^{\text{tm}}, \text{Obf}_{\text{diff}}^{\text{tm}}$ be TM obfuscators. Let H be a function family such that $\text{H.In}(\lambda) = \{0, 1\}^*$ and $\text{H.Out}(\lambda) \subseteq \bigcup_{i \leq p_0(\lambda)} \{0, 1\}^i$ for some polynomial p_0 and all $\lambda \in \mathbb{N}$. Let DS be a deterministic digital signature scheme such that $\text{DS.In}(\lambda) = \{0, 1\}^{\ell(\lambda)}$ and $\text{DS.Out}(\lambda) \subseteq \bigcup_{i \leq p_1(\lambda)} \{0, 1\}^i$ for some polynomial p_1 and all $\lambda \in \mathbb{N}$. We build a TM sampler $S^{\text{tm}} = \text{TM-SAMP}[\text{Obf}_{\text{diff}}^{\text{tm}}, \text{H}, \text{DS}, \text{Obf}_{\text{eq}}^{\text{tm}}, s_0, \ell, n, t_0, t_1, s_1]$ as defined in Fig. 5. We say that S^{tm} is *well-defined* if $s_0(\lambda) \geq |M^0|$, $s_0(\lambda) \geq |M_{1^\lambda, hk, h}^1|$, $\ell(n(\lambda)) \geq \ell(\lambda)$, $t_0(\lambda) \geq \text{time}(M_{1^\lambda, vk}^{\text{ver}}, (m, \sigma))$, $t_1(\lambda) \geq \text{time}(\bar{M}, (M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \bar{M}, \sigma))$ and $s_1(\lambda) \geq |M_{1^\lambda, sk, vk}^{\text{aux}}|$ for all $\lambda \in \mathbb{N}$, $hk \in [\text{H.Kg}(1^\lambda)]$, $h \in \text{H.Out}(\lambda)$, $M \in \{M^0, M_{1^\lambda, hk, h}^1\}$, $\bar{M} \in [\text{Obf}_{\text{diff}}^{\text{tm}}(1^\lambda, \text{Pad}_{s_0(\lambda)}(M))]$, $(sk, vk) \in [\text{DS.Kg}(1^{n(\lambda)})]$, $m \in \{0, 1\}^{\ell(\lambda)}$ and $\sigma \in \text{DS.Out}(n(\lambda))$.

Core design ideas behind TM-SAMP. Note that TM $M_{1^\lambda, sk, vk}^{\text{aux}}$ takes as input an obfuscated TM \bar{M} and computes the signature σ for message $\langle \bar{M} \rangle_{\ell(n(\lambda))}$, where

<pre> TM Sampler $S^{\text{tm}}(1^\lambda)$ $(sk, vk) \leftarrow \text{DS.Kg}(1^{n(\lambda)})$ $hk \leftarrow \text{H.Kg}(1^\lambda)$ $h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$ $M_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0)$ $M_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$ $M_{\text{aux}} \leftarrow \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk, vk}^{\text{aux}})$ $aux \leftarrow \text{Obf}_{\text{eq}}^{\text{tm}}(1^{n(\lambda)}, M_{\text{aux}})$ Return (M_0, M_1, aux) TM $M_{1^\lambda, vk}^{\text{ver}}(m, \sigma)$ If $(m \neq \ell(\lambda))$ then return 0 Return $\text{DS.Ver}(1^{n(\lambda)}, vk, \langle m \rangle_{\ell(n(\lambda))}, \sigma)$ </pre>	<pre> TM $M^0(M, 1^t, m, \sigma)$ Return 0 TM $M_{1^\lambda, hk, h}^1(M, 1^t, m, \sigma)$ $h' \leftarrow \text{H.Ev}(1^\lambda, hk, M)$ If $(h' \neq h)$ then return 0 Return $\text{UTM}_M^t(m, \sigma)$ TM $M_{1^\lambda, sk, vk}^{\text{aux}}(\bar{M})$ If $(\bar{M} \neq \ell(\lambda))$ then return 0 $\sigma \leftarrow \text{DS.Sig}(1^{n(\lambda)}, sk, \langle \bar{M} \rangle_{\ell(n(\lambda))})$ $d \leftarrow \text{UTM}_{\bar{M}}^{t_1(\lambda)}(M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \bar{M}, \sigma)$ Return d </pre>
---	--

Fig. 5. TM sampler $S^{\text{tm}} = \text{TM-SAMP}[\text{Obf}_{\text{diff}}^{\text{tm}}, \text{H}, \text{DS}, \text{Obf}_{\text{eq}}^{\text{tm}}, s_0, \ell, n, t_0, t_1, s_1]$.

the latter denotes \bar{M} padded to size $\ell(n(\lambda))$. It then uses a Universal Turing Machine UTM to simulate \bar{M} on input x for the duration of $t_1(\lambda)$ steps, where $x = (M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \bar{M}, \sigma)$. The idea of computing a signature on a message that depends on \bar{M} was already proposed in GGHW [28], with the goal of avoiding a trivial attack against the difference-security of the sampler. Specifically, if a fixed message-signature pair $(m_{\text{ch}}, \sigma_{\text{ch}})$ was used for all inputs of $M_{1^\lambda, sk, vk}^{\text{aux}}$, then a difference-security adversary could construct a sequence of TMs that each reveals a single bit of $(m_{\text{ch}}, \sigma_{\text{ch}})$ when used as an input \bar{M} to $M_{1^\lambda, sk, vk}^{\text{aux}}$. This would allow adversary to recover the message-signature pair bit-by-bit.

Turing Machine $M_{1^\lambda, hk, h}^1$ takes an input $x = (M, 1^t, m, \sigma)$, where M is a TM, 1^t is the unary representation of some integer $t \in \mathbb{N}$, and (m, σ) is a message-signature pair. We use a target collision-resistant function family H in order to ensure that $M_{1^\lambda, hk, h}^1$ can return 1 only if $M = M_{1^\lambda, vk}^{\text{ver}}$. This is achieved by embedding a key hk for H and the value $h = \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$ into $M_{1^\lambda, hk, h}^1$, and by returning 0 whenever $h \neq \text{H.Ev}(1^\lambda, hk, M)$. If $M = M_{1^\lambda, vk}^{\text{ver}}$ is satisfied, then $M_{1^\lambda, hk, h}^1$ uses a Universal Turing Machine UTM to simulate M on input (m, σ) for the duration of t steps. TM $M_{1^\lambda, vk}^{\text{ver}}$ is designed to return 1 if and only if its input $x = (m, \sigma)$ is a valid message-signature pair with respect to a verification key vk for the digital signature scheme DS . Our impossibility results require the choice of DS to depend on the construction of $M_{1^\lambda, hk, h}^1$, so embedding vk directly into the latter would have introduced a circular dependency between the two. Instead we have to resort to the above approach of embedding vk into a separate TM.

According to our definitions, two TMs can be functionally equivalent only if both of them halt on all inputs. The notion of functional equivalence is further used for the definitions of program samplers and obfuscation. This means that whenever a TM needs to simulate the code of another TM, it is required to

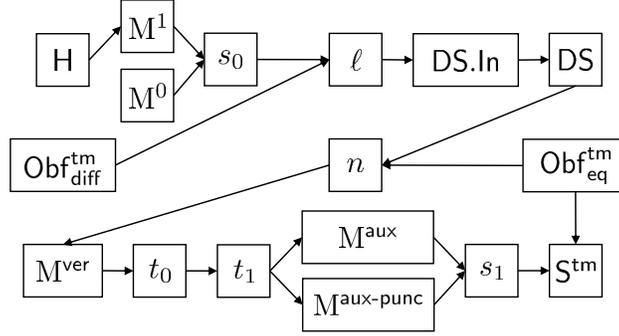


Fig. 6. Parameter dependencies in TM-SAMP for the proof of Theorem 2.

use a Universal Turing Machine UTM and specify the number of steps for the simulation. Otherwise, the simulated TMs would not be guaranteed to halt.

Parameters of TM-SAMP. Fig. 6 shows the dependencies between all schemes and parameters that will be used to instantiate the construction of TM-SAMP in Theorem 2. Let us introduce the notation that is used in this picture. For any two entities A and B, an arrow from A to B means that the construction, or the choice, of B depends on A. The relations are transitive, meaning that we do not draw a direct arrow from A to B in the case if B is already reachable from A. TM $M^{\text{aux-punc}}$ will be used only for the proof of security and is defined in Fig. 7.

The construction of TM-SAMP is parameterized by polynomials s_0, s_1, t_0, t_1, ℓ and n . Polynomials s_0, s_1 denote the size to which some of our TMs must be padded prior to obfuscating them. This stems from our definition of program samplers that are required to return programs of the same size. Polynomials t_0, t_1 are used to indicate the number of steps that must be done when simulating various TMs using a Universal Turing Machine UTM. Our definition of a well-defined instantiation of TM-SAMP specifies lower bounds for t_0, t_1 that ensure the correctness of the attack that we will design against the sub-exponential (d)iO-security of $\text{Obf}_{\text{diff}}^{\text{tm}}$ with respect to S^{tm} . Polynomial ℓ will be defined to upper-bound the size of any obfuscation \bar{M} of programs M^0 and $M_{1^\lambda, hk, h}^1$, when obfuscator $\text{Obf}_{\text{diff}}^{\text{tm}}$ is used. Note that $M_{1^\lambda, sk, vk}^{\text{aux}}$ rejects all inputs \bar{M} of size different than $\ell(\lambda)$; our attack will pad all obfuscations of M^0 and $M_{1^\lambda, hk, h}^1$ to size $\ell(\lambda)$, using the padding operator $\text{Pad}_{\ell(\lambda)}(\cdot)$ that is assumed to produce functionally equivalent TMs as per Section 2. Polynomial n is used to set security parameters for schemes DS and $\text{Obf}_{\text{eq}}^{\text{tm}}$. Specifically, if the TM sampler S^{tm} is instantiated with a security parameter $\lambda \in \mathbb{N}$, then its construction uses these two schemes, each with the security parameter $n(\lambda)$.

In order for our proof of difference-security to work, if a $2^{-(\cdot)^\epsilon}$ -security is assumed for either of DS or $\text{Obf}_{\text{eq}}^{\text{tm}}$, then the choice of polynomial n will depend on ϵ . This leads to an inconvenient dependency: DS uses $n(\lambda)$ as its security

parameter, but the choice of polynomial n depends on the choice of DS . Ideally, we would have liked to choose a digital signature scheme DS such that $\text{DS.Out}(n(\lambda)) = \{0, 1\}^{\ell(\lambda)}$, because DS is used to sign messages that are TMs of size $\ell(\lambda)$. However, since we do not know n ahead of choosing DS , we require that for all $\lambda \in \mathbb{N}$ we have $\text{DS.Out}(\lambda) = \{0, 1\}^{\ell(\lambda)}$ and $\ell(n(\lambda)) \geq \ell(\lambda)$, resulting in $\text{DS.Out}(n(\lambda)) = \{0, 1\}^{\ell(n(\lambda))}$. We then use an injective string padding to map TMs (i.e. their string representations) of length $\ell(\lambda)$ into strings of length $\ell(n(\lambda))$. The injectivity of padding is necessary for the proof of difference-security of \mathbf{S}^{tm} . In order to ensure that the requirement $\ell(n(\lambda)) \geq \ell(\lambda)$ is satisfied, we will choose polynomials ℓ, n such that $\ell(\lambda+1) \geq \ell(\lambda)$ and $n(\lambda) \geq \lambda$ for all $\lambda \in \mathbb{N}$.

Limitations and extensions. Our definition of TM samplers in Section 2 requires them to return TMs that halt on all inputs. One could argue that this definition is still insufficient for the purpose of obfuscation. Namely, a sampler can produce TMs that have significantly different running times, and it might not be reasonable to expect an obfuscator to properly hide the difference in the running times. We note that this does not hinder our results because we can artificially alter our TMs M^0 and $M_{1^\lambda, hk, h}^1$ to have the same running times, by adding void instructions to the definition of M^0 .

The construction of TM-SAMP uses a TM obfuscator $\text{Obf}_{\text{eq}}^{\text{tm}}$ that in our theorem statements will be assumed to be sub-exponentially $\mathbf{S}_{\text{eq}}^{\text{tm}}$ -secure. It is used to produce auxiliary information by obfuscating TMs $M_{1^\lambda, sk, vk}^{\text{aux}}$ and $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$. We use a TM obfuscator for readability, but we note that a sub-exponentially $\mathbf{S}_{\text{eq}}^{\text{circ}}$ -secure *circuit* obfuscator could be used instead. There are no circular dependencies preventing us from redefining these two TMs as circuits.

According to Fig. 6, the size of M^{aux} depends on the maximum size of TMs M^0 and $M_{1^\lambda, hk, h}^1$, and in particular it might be larger than these TMs. This means that our impossibility result might not hold if we restrict our attention to TM samplers whose auxiliary information strings aux are required to be shorter than the size of the corresponding TMs M^0 and M^1 . GGHW [28] circumvent this limitation in their impossibility result by using a CRHF to compute and then sign a hash of the TM that is passed inside their auxiliary-information program, rather than signing the TM itself. Our proof techniques do not seem to be compatible with such approach.

Impossibility results. We now formally state our results. Theorem 2 shows how to choose parameters for TM-SAMP such that the resulting TM sampler is simultaneously well-defined and difference-secure. Theorem 3 shows that any well-defined instantiation of TM-SAMP produces TMs that can not be securely obfuscated.

Theorem 2. *Let $\text{Obf}_{\text{diff}}^{\text{tm}}$ be a TM obfuscator. Let \mathbf{H} be a sub-exponentially TCR-secure function family such that $\mathbf{H.In}(\lambda) = \{0, 1\}^*$ and $\mathbf{H.Out}(\lambda) \subseteq \bigcup_{i \leq p_0(\lambda)} \{0, 1\}^i$ for some polynomial p_0 and all $\lambda \in \mathbb{N}$. Then there are polynomials $s_0, \ell: \mathbb{N} \rightarrow \mathbb{N}$ such that the following is true. Let DS be a sub-exponentially PSUFCMA-secure, consistent puncturable digital signature scheme such that $\text{DS.In}(\lambda) = \{0, 1\}^{\ell(\lambda)}$ and $\text{DS.Out}(\lambda) \subseteq \bigcup_{i \leq p_1(\lambda)} \{0, 1\}^i$ for some polynomial p_1*

and all $\lambda \in \mathbb{N}$. Let $\text{Obf}_{\text{eq}}^{\text{tm}}$ be a sub-exponentially $\mathbf{S}_{\text{eq}}^{\text{tm}}$ -secure TM obfuscator. Then there are polynomials $n, t_0, t_1, s_1: \mathbb{N} \rightarrow \mathbb{N}$ such that the following is true. Let $\mathbf{S}^{\text{tm}} = \text{TM-SAMP} [\text{Obf}_{\text{diff}}^{\text{tm}}, \mathbf{H}, \text{DS}, \text{Obf}_{\text{eq}}^{\text{tm}}, s_0, \ell, n, t_0, t_1, s_1]$. Then (1) \mathbf{S}^{tm} is well-defined, and (2) \mathbf{S}^{tm} is sub-exponentially DIFF-secure.

We defer the proof of Theorem 2 until after we show how to use this theorem to state and prove our main claims regarding the impossibility of differing-inputs obfuscation for TMs.

Theorem 3. Let $s_0, \ell, n, t_0, t_1, s_1: \mathbb{N} \rightarrow \mathbb{N}$ be polynomials. Let $\text{Obf}_{\text{eq}}^{\text{tm}}, \text{Obf}_{\text{diff}}^{\text{tm}}$ be TM obfuscators. Let \mathbf{H} be a function family with $\mathbf{H}.\text{In}(\lambda) = \{0, 1\}^*$ and $\mathbf{H}.\text{Out}(\lambda) \subseteq \bigcup_{i \leq p_0(\lambda)} \{0, 1\}^i$ for some polynomial p_0 and all $\lambda \in \mathbb{N}$. Let DS be a deterministic digital signature scheme such that $\text{DS}.\text{In}(\lambda) = \{0, 1\}^{\ell(\lambda)}$ and $\text{DS}.\text{Out}(\lambda) \subseteq \bigcup_{i \leq p_1(\lambda)} \{0, 1\}^i$ for some polynomial p_1 and all $\lambda \in \mathbb{N}$. Let $\mathbf{S}^{\text{tm}} = \text{TM-SAMP} [\text{Obf}_{\text{diff}}^{\text{tm}}, \mathbf{H}, \text{DS}, \text{Obf}_{\text{eq}}^{\text{tm}}, s_0, \ell, n, t_0, t_1, s_1]$. Assume that \mathbf{S}^{tm} is well-defined. Then there exists a PT adversary \mathcal{O} such that $\text{Adv}_{\text{Obf}_{\text{diff}}^{\text{tm}}, \mathbf{S}^{\text{tm}}, \mathcal{O}}^{\text{iO}}(\lambda) = 1$.

Proof (Theorem 3). We build a PT adversary \mathcal{O} against the (d)iO-security of $\text{Obf}_{\text{diff}}^{\text{tm}}$ relative to \mathbf{S}^{tm} as follows:

$$\begin{aligned} & \text{Adversary } \mathcal{O}(1^\lambda, \overline{\mathbf{M}}, \text{aux}) \\ & \overline{\mathbf{M}}_{\text{aux}} \leftarrow \text{aux}; b' \leftarrow \overline{\mathbf{M}}_{\text{aux}}(\text{Pad}_{\ell(\lambda)}(\overline{\mathbf{M}})); \text{Return } b' \end{aligned}$$

Adversary \mathcal{O} takes $1^\lambda, \overline{\mathbf{M}}, \text{aux}$ as input, where $\overline{\mathbf{M}}$ is an obfuscation of either TM M^0 or TM $M_{1^\lambda, hk, h}^1$ that was produced by the obfuscator $\text{Obf}_{\text{diff}}^{\text{tm}}$ in game $\text{IO}_{\text{Obf}_{\text{diff}}^{\text{tm}}, \mathbf{S}^{\text{tm}}}^{\mathcal{O}}(\lambda)$, and aux is an auxiliary information string. The goal of \mathcal{O} is to guess which of M^0 and $M_{1^\lambda, hk, h}^1$ was obfuscated. It should return 0 if $\overline{\mathbf{M}}$ is an obfuscation of M^0 , and it should return 1 otherwise.

Adversary \mathcal{O} parses auxiliary information string aux into a TM $\overline{\mathbf{M}}_{\text{aux}}$. The latter is an obfuscation of TM $M_{1^\lambda, sk, vk}^{\text{aux}}$, which was computed in \mathbf{S}^{tm} using obfuscator $\text{Obf}_{\text{eq}}^{\text{tm}}$. Next, \mathcal{O} pads $\overline{\mathbf{M}}$ to construct a functionally equivalent TM of size $\ell(\lambda)$ and passes it as input to $\overline{\mathbf{M}}_{\text{aux}}$. According to the construction of $M_{1^\lambda, sk, vk}^{\text{aux}}$, the latter returns 1 if and only if $\overline{\mathbf{M}}$ is an obfuscation of TM $M_{1^\lambda, hk, h}^1$. Adversary \mathcal{O} returns the same value to win the game. This concludes the proof of Theorem 3.

Next, Theorem 4 shows the impossibility of a polynomially secure diO, and Theorem 5 shows the impossibility of a sub-exponentially secure diO.

Theorem 4. Let Obf be a Turing Machine obfuscator. Assume the existence of sub-exponentially secure one-way functions and sub-exponentially secure indistinguishability obfuscation for Turing Machines. Then Obf is not a differing-inputs obfuscator.

We now prove Theorem 4. Let $\text{Obf}_{\text{eq}}^{\text{tm}}$ be a sub-exponentially $\mathbf{S}_{\text{eq}}^{\text{tm}}$ -secure TM obfuscator. Theorem 1 shows how to build a sub-exponentially PSUFCMA-secure, consistent puncturable digital signature scheme DS assuming only sub-exponentially secure OWF and sub-exponentially secure iO. For a moment, assume that we can build a TCR-secure function family \mathbf{H} with $\mathbf{H}.\text{In}(\lambda) = \{0, 1\}^*$

for all $\lambda \in \mathbb{N}$ just from sub-exponentially secure OWFs (which is not known to be true, and we address this below). Then according to Theorem 2, we can build a TM sampler \mathbf{S}^{tm} that is (1) well-defined and (2) sub-exponentially DIFF-secure. But Theorem 3 shows that there exists an efficient adversary that breaks the IO-security of Obf with respect to \mathbf{S}^{tm} . Therefore, Obf is not a differing-inputs obfuscator.

In order to build a TCR-secure function family \mathbf{H} from a sub-exponentially secure OWF, the statements of Theorem 2 and Theorem 3 can be relaxed to require $\mathbf{H}.\text{In}(\lambda) = \{0, 1\}^{2^\lambda}$ for all $\lambda \in \mathbb{N}$. This change will still ensure the correctness of \mathbf{S}^{tm} , which requires that \mathbf{H} can process inputs of length $|M_{1^\lambda, vk}^{\text{ver}}|$. The size of $M_{1^\lambda, vk}^{\text{ver}}$ in our construction is bounded polynomially in the security parameter. But the reason we have to use a hash function that can process inputs of arbitrary, super-polynomially bounded lengths is because the size of $M_{1^\lambda, vk}^{\text{ver}}$ is not known prior to fixing \mathbf{H} (as shown in Fig. 6).

As noted in Section 2, Shoup [46] shows how to build a TCR-secure function family \mathbf{H} for arbitrary, bounded variable-length inputs from any TCR-secure compression function family with fixed input size. The latter is shown to be achievable from OWFs by Rompel [44]. We note that the key size of Shoup's construction grows logarithmically with the maximum input length of the constructed function family, which is still polynomially bounded in the case of \mathbf{H} that was proposed above. Furthermore, the super-polynomial bound on the message lengths does not introduce any difficulties for the security reduction of Shoup's construction. This is because the loss of security during the reduction depends on the length of the messages that are chosen by a PT adversary, rather than by the (super-polynomial) bound on the messages supported by the scheme.

This concludes the proof of Theorem 4. Note that we ruled out the existence of polynomially-secure differing-inputs obfuscation even with respect to *sub-exponentially* secure TM samplers, which is a stronger version of difference-security than the one required by our definition of polynomially-secure differing-inputs obfuscation.

Theorem 5. *Let Obf be a Turing Machine obfuscator. Assume the existence of sub-exponentially secure one-way functions. Then Obf is not a sub-exponentially secure differing-inputs obfuscator.*

To prove Theorem 5, assume for a contradiction that Obf is a sub-exponentially secure differing-inputs obfuscator. According to our definitions, it implies the existence of sub-exponentially secure indistinguishability obfuscation. The rest of the proof is identical to the proof of Theorem 4. It results in constructing a sub-exponentially difference-secure TM sampler \mathbf{S}^{tm} that can not be securely obfuscated by Obf . Thus, we get a contradiction.

Finally, we now prove Theorem 2.

Proof (Theorem 2). We start by proving part (1) of the theorem. Specifically, we choose polynomials $s_0, \ell, n, t_0, t_1, s_1: \mathbb{N} \rightarrow \mathbb{N}$ such that \mathbf{S}^{tm} is well-defined.

We now specify polynomials $s_0, \ell: \mathbb{N} \rightarrow \mathbb{N}$. For any $\lambda \in \mathbb{N}$ let $s_0(\lambda)$ be a polynomial upper bound on $\max(|M^0|, |M_{1^\lambda, hk, h}^1|)$ where the maximum is over

all $hk \in [\text{H.Kg}(1^\lambda)]$ and $h \in \text{H.Out}(\lambda)$. For any $\lambda \in \mathbb{N}$ let $\ell(\lambda)$ be a polynomial upper bound on $\max(|\overline{M}|)$ such that $\ell(\lambda) \leq \ell(\lambda + 1)$, where the maximum is over all $hk \in [\text{H.Kg}(1^\lambda)]$, $h \in \text{H.Out}(\lambda)$, $M \in \{M^0, M_{1^\lambda, hk, h}^1\}$ and $\overline{M} \in [\text{Obf}_{\text{diff}}^{\text{tm}}(1^\lambda, \text{Pad}_{s_0(\lambda)}(M))]$. Note that the requirement that $\ell(\lambda) \leq \ell(\lambda + 1)$ for all $\lambda \in \mathbb{N}$ is trivially achievable by removing all terms with negative coefficients from the polynomial.

We now specify a constant $0 < \epsilon < 1$ for which we will prove that S^{tm} is $2^{-(\cdot)^\epsilon}$ -DIFF-secure. Let $0 < \epsilon_{\text{tcr}} < 1$ be a constant such that H is $2^{-(\cdot)^{\epsilon_{\text{tcr}}}}$ -TCR-secure. Let $0 < \epsilon_{\text{psuf}} < 1$ be a constant such that DS is $2^{-(\cdot)^{\epsilon_{\text{psuf}}}}$ -PSUFCMA-secure. Let $0 < \epsilon_{\text{io}} < 1$ be a constant such that $\text{Obf}_{\text{eq}}^{\text{tm}}$ is $2^{-(\cdot)^{\epsilon_{\text{io}}}}$ - $\text{S}_{\text{eq}}^{\text{tm}}$ -secure. Let $\epsilon = \min(\frac{1}{2}\epsilon_{\text{tcr}}, \epsilon_{\text{psuf}}, \epsilon_{\text{io}})$.

We now specify polynomial $n: \mathbb{N} \rightarrow \mathbb{N}$. For any $\lambda \in \mathbb{N}$ let $n(\lambda) = (2\lambda + \ell(\lambda) + 3)^{\lceil 1/\epsilon \rceil}$. Note that for any $\lambda \in \mathbb{N}$ we have $n(\lambda) \geq \lambda$, and earlier we required that $\ell(\lambda + 1) \geq \ell(\lambda)$ for all $\lambda \in \mathbb{N}$. It follows that $\ell(n(\lambda)) \geq \ell(\lambda)$ for all $\lambda \in \mathbb{N}$, as required for S^{tm} to be well-defined. Let Inv_n be a deterministic, PT algorithm that takes 1^λ to return the smallest $\lambda \in \mathbb{N}$ such that $n(\lambda) \geq \lambda'$. We note that n is injective, implying that $\text{Inv}_n(1^{n(\lambda)}) = \lambda$ for all $\lambda \in \mathbb{N}$.

We now specify polynomials $t_0, t_1, s_1: \mathbb{N} \rightarrow \mathbb{N}$. For any $\lambda \in \mathbb{N}$ let $t_0(\lambda)$ be a polynomial upper bound on the maximum running time of $M_{1^\lambda, vk}^{\text{ver}}(m, \sigma)$ where the maximum is over all $(sk, vk) \in [\text{DS.Kg}(1^{n(\lambda)})]$, $m \in \{0, 1\}^{\ell(\lambda)}$ and $\sigma \in \text{DS.Out}(n(\lambda))$. For any $\lambda \in \mathbb{N}$ let $t_1(\lambda)$ be a polynomial upper bound on the maximum running time of $\overline{M}(M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \overline{M}, \sigma)$ where the maximum is over all $hk \in [\text{H.Kg}(1^\lambda)]$, $h \in \text{H.Out}(\lambda)$, $M \in \{M^0, M_{1^\lambda, hk, h}^1\}$, $\overline{M} \in [\text{Obf}_{\text{diff}}^{\text{tm}}(1^\lambda, \text{Pad}_{s_0(\lambda)}(M))]$, $(sk, vk) \in [\text{DS.Kg}(1^{n(\lambda)})]$ and $\sigma \in \text{DS.Out}(n(\lambda))$. For any $\lambda \in \mathbb{N}$ let $s_1(\lambda)$ be a polynomial upper bound on $\max(|M_{1^\lambda, sk, vk}^{\text{aux}}|, |M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}|)$ where the TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ is defined in Fig. 7 and where the maximum is over all $(sk, vk) \in [\text{DS.Kg}(1^{n(\lambda)})]$, $m' \in \{0, 1\}^{\ell(\lambda)}$, $sk^* \in [\text{DS.PKg}(1^{n(\lambda)}, sk, \langle m' \rangle_{\ell(n(\lambda))})]$ and $b \in \{0, 1\}$.

We proceed to prove part (2) of Theorem 2, namely that S^{tm} is $2^{-(\cdot)^\epsilon}$ -DIFF-secure. The main challenge of the proof is to show that the signing key sk of DS can not be extracted from an obfuscation of TM $M_{1^\lambda, sk, vk}^{\text{aux}}$, meaning that the $\text{S}_{\text{eq}}^{\text{tm}}$ -secure obfuscator $\text{Obf}_{\text{eq}}^{\text{tm}}$ is sufficient to hide sk . In our proof this is implicit. The core idea of the proof is to consider the exponential number of messages from $\text{DS.In}(n(\lambda))$ and for each of them we argue that a PT adversary is unlikely to produce a signature for this message. This implies that it is hard to find an input on which TMs M^0 and $M_{1^\lambda, hk, h}^1$ return different outputs.

Let \mathcal{D} be a PT adversary. Consider the games and associated TMs of Fig. 7. Lines not annotated with comments are common to all games. Game $G_0(\lambda)$ is equivalent to $\text{DIFF}_{\text{S}^{\text{tm}}}^{\mathcal{D}}(\lambda)$, so for all $\lambda \in \mathbb{N}$ we have

$$\text{Adv}_{\text{S}^{\text{tm}}, \mathcal{D}}^{\text{diff}}(\lambda) = \Pr[G_0(\lambda)]. \quad (1)$$

Let us discuss the transitions between hybrid games that will be used in our proof. Let $\lambda \in \mathbb{N}$. In order to transition from game $G_0(\lambda)$ to game $G_{1,0}(\lambda)$ we claim that if adversary \mathcal{D} wins in game $\text{DIFF}_{\text{S}^{\text{tm}}}^{\mathcal{D}}(\lambda)$ then it must return a

Games $G_0(\lambda)$ – $G_{1,2^\ell(\lambda)}(\lambda)$	Games $G_{1,i}(\lambda)$ – $G_{1,i+1}(\lambda)$
$(sk, vk) \leftarrow \text{DS.Kg}(1^{n(\lambda)})$ $hk \leftarrow \text{H.Kg}(1^\lambda)$ $h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$ $M_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0)$ $M_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$ $M_{\text{aux}} \leftarrow \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk, vk}^{\text{aux}})$ $aux \leftarrow \text{Obf}_{\text{eq}}^{\text{tm}}(1^{n(\lambda)}, M_{\text{aux}})$ $x \leftarrow \mathcal{D}(1^\lambda, M_0, M_1, aux)$ $(M, 1^t, m, \sigma) \leftarrow x$ $d_0 \leftarrow (M_0(x) \neq M_1(x))$ $d_1 \leftarrow (M = M_{1^\lambda, vk}^{\text{ver}})$ $\text{res} \leftarrow d_0 \quad // G_0$ $\text{res} \leftarrow (d_0 \wedge d_1 \wedge m \geq 0) \quad // G_{1,0}$ \dots $\text{res} \leftarrow (d_0 \wedge d_1 \wedge m \geq 2^\ell(\lambda)) \quad // G_{1,2^\ell(\lambda)}$ Return res	$(sk, vk) \leftarrow \text{DS.Kg}(1^{n(\lambda)})$ $hk \leftarrow \text{H.Kg}(1^\lambda)$ $h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$ $M_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0)$ $M_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$ $m' \leftarrow \langle i \rangle_{\ell(\lambda)}; b \leftarrow M_{1^\lambda, sk, vk}^{\text{aux}}(m')$ $m^* \leftarrow \langle m' \rangle_{\ell(n(\lambda))}$ $sk^* \leftarrow \text{DS.PKg}(1^{n(\lambda)}, sk, m^*)$ $M_{\text{tmp}} \leftarrow M_{1^\lambda, sk, vk}^{\text{aux}}; z \leftarrow i \quad // G_{1,i}$ $M_{\text{tmp}} \leftarrow M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}; z \leftarrow i \quad // G_{1,i,A}$ $M_{\text{tmp}} \leftarrow M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}; z \leftarrow i+1 \quad // G_{1,i,B}$ $M_{\text{tmp}} \leftarrow M_{1^\lambda, sk, vk}^{\text{aux}}; z \leftarrow i+1 \quad // G_{1,i+1}$ $aux \leftarrow \text{Obf}_{\text{eq}}^{\text{tm}}(1^{n(\lambda)}, \text{Pad}_{s_1(\lambda)}(M_{\text{tmp}}))$ $x \leftarrow \mathcal{D}(1^\lambda, M_0, M_1, aux)$ $(M, 1^t, m, \sigma) \leftarrow x$ $d_0 \leftarrow (M_0(x) \neq M_1(x))$ $d_1 \leftarrow (M = M_{1^\lambda, vk}^{\text{ver}})$ Return $(d_0 \wedge d_1 \wedge m \geq z)$
<u>TM $M^0(M, 1^t, m, \sigma)$</u> Return 0 <u>TM $M_{1^\lambda, hk, h}^1(M, 1^t, m, \sigma)$</u> $h' \leftarrow \text{H.Ev}(1^\lambda, hk, M)$ If $(h' \neq h)$ then return 0 Return $\text{UTM}_{M}^t(m, \sigma)$ <u>TM $M_{1^\lambda, sk, vk}^{\text{aux}}(\bar{M})$</u> If $(\bar{M} \neq \ell(\lambda))$ then return 0 $\sigma \leftarrow \text{DS.Sig}(1^{n(\lambda)}, sk, \langle \bar{M} \rangle_{\ell(n(\lambda))})$ Return $\text{UTM}_{\bar{M}}^{t_1(\lambda)}(M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \bar{M}, \sigma)$	<u>TM $M_{1^\lambda, vk}^{\text{ver}}(m, \sigma)$</u> If $(m \neq \ell(\lambda))$ then return 0 $d \leftarrow \text{DS.Ver}(1^{n(\lambda)}, vk, \langle m \rangle_{\ell(n(\lambda))}, \sigma)$ Return d <u>TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}(\bar{M})$</u> If $(\bar{M} \neq \ell(\lambda))$ then return 0 If $(\bar{M} = m')$ then return b $\sigma \leftarrow \text{DS.PSig}(1^{n(\lambda)}, sk^*, \langle \bar{M} \rangle_{\ell(n(\lambda))})$ $d \leftarrow \text{UTM}_{\bar{M}}^{t_1(\lambda)}(M_{1^\lambda, vk}^{\text{ver}}, 1^{t_0(\lambda)}, \bar{M}, \sigma)$ Return d

Fig. 7. Games for proof of Theorem 2.

differing-input $x = (M, 1^t, m, \sigma)$ such that $M = M_{1^\lambda, vk}^{\text{ver}}$. Otherwise, one could use this adversary to break the TCR-security of H. Next, we consider an exponential number of games, going from game $G_{1,0}(\lambda)$ to game $G_{1,2^\ell(\lambda)}(\lambda)$. Each game corresponds to a unique value of message m that can be taken as input by TM $M_{1^\lambda, vk}^{\text{ver}}$. For any $i \in \{0, 1, \dots, 2^\ell(\lambda)\}$, adversary \mathcal{D} wins in game $G_{1,i}(\lambda)$ if and only if it returns $x = (M, 1^t, m, \sigma)$ such that $M = M_{1^\lambda, vk}^{\text{ver}}$, $m \geq i$ and $M_0(x) \neq M_1(x)$. According to this definition, it is impossible to win game $G_{1,2^\ell(\lambda)}(\lambda)$ because TM $M_{1^\lambda, vk}^{\text{ver}}$ rejects whenever it takes a message m as input such that $|m| \neq \ell(\lambda)$ (whereas the length of m in this game is required to be at least $\ell(\lambda) + 1$). We

now need to show that for each $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$ the success probabilities of adversary \mathcal{D} in games $G_{1,i}(\lambda)$ and $G_{1,i+1}(\lambda)$ are sub-exponentially close.

Let $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$. We split the transition from game $G_{1,i}(\lambda)$ to game $G_{1,i+1}(\lambda)$ into three steps. Specifically, we consider a sequence of games $G_{1,i}(\lambda)$, $G_{1,i,A}(\lambda)$, $G_{1,i,B}(\lambda)$ and $G_{1,i+1}(\lambda)$. Games $G_{1,i,A}(\lambda)$ and $G_{1,i,B}(\lambda)$ generate aux as an obfuscation of TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ instead of an obfuscation of TM $M_{1^\lambda, sk, vk}^{\text{aux}}$, where $m' = i$ and the used obfuscator is $\text{Obf}_{\text{eq}}^{\text{tm}}$. As opposed to TM $M_{1^\lambda, sk, vk}^{\text{aux}}$, note that TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ contains a punctured signing key sk^* for DS that is punctured at message $m^* = \langle m' \rangle_{\ell(n(\lambda))}$. Both TMs are defined to produce the same outputs on all inputs such that $\bar{M} \neq m'$, which is achieved because the punctured digital signature scheme DS is assumed to be consistent. (Recall that the latter requires that sk and sk^* return the same signatures for all messages except m^* .) Furthermore, TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ is hardwired to return $b = M_{1^\lambda, sk, vk}^{\text{aux}}(m')$ on input $\bar{M} = m'$, meaning that the TMs are functionally equivalent. We use it to claim that the success probabilities of adversary \mathcal{D} in games $G_{1,i}(\lambda)$ and $G_{1,i,A}(\lambda)$ — and in games $G_{1,i,B}(\lambda)$ and $G_{1,i+1}(\lambda)$ —are sub-exponentially close. Namely, if \mathcal{D} can distinguish between any pair of these games with a better than sub-exponentially small probability, then one can use \mathcal{D} to break the iO-security of obfuscator $\text{Obf}_{\text{eq}}^{\text{tm}}$.

It remains to discuss the transition from game $G_{1,i,A}(\lambda)$ to game $G_{1,i,B}(\lambda)$. The difference between these games is that the former requires $m \geq i$ as a part of its winning condition, whereas the later requires $m \geq i + 1$. Both of these games set aux to be an obfuscation of TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$, where sk^* is punctured at $m^* = \langle m' \rangle_{\ell(n(\lambda))}$ and $m' = i$. Note that adversary \mathcal{D} can only have a different success probability in both games if it is capable of forging a signature on message m^* given any information it might be able to extract from TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$. However, $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ does not contain any information that could help to forge the signature for message m^* (only bit b depends on the challenge signature, but \mathcal{D} can attempt to guess it). Therefore, we can use the PSUFMA-security of DS to bound the difference in adversary's success probability when transitioning between games $G_{1,i,A}(\lambda)$ and $G_{1,i,B}(\lambda)$.

Below we will prove the following claims:

Claim 1. There exists a PT adversary \mathcal{H} against the TCR-security of H such that for all $\lambda \in \mathbb{N}$ we have

$$\Pr[G_0(\lambda)] - \Pr[G_{1,0}(\lambda)] \leq \text{Adv}_{\text{H}, \mathcal{H}}^{\text{tcr}}(\lambda). \quad (2)$$

Claim 2. There exist TM samplers $S_0^{\text{tm}}, S_1^{\text{tm}}$ and a PT adversary \mathcal{O} against the iO-security of $\text{Obf}_{\text{eq}}^{\text{tm}}$ relative to S_0^{tm} and S_1^{tm} , such that for all $\lambda \in \mathbb{N}$ we have

$$\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[G_{1,i}(\lambda)] - \Pr[G_{1,i,A}(\lambda)]) \leq 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}^{\text{tm}}, S_0^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda)), \quad (3)$$

$$\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[G_{1,i,B}(\lambda)] - \Pr[G_{1,i+1}(\lambda)]) \leq 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}^{\text{tm}}, S_1^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda)). \quad (4)$$

Claim 3. There exists a PT adversary \mathcal{U} against the PSUFCMA-security of DS such that for all $\lambda \in \mathbb{N}$ we have

$$\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[G_{1,i,A}(\lambda)] - \Pr[G_{1,i,B}(\lambda)]) \leq 2^{\ell(\lambda)+1} \cdot \text{Adv}_{\text{DS},\mathcal{U}}^{\text{psufcma}}(n(\lambda)). \quad (5)$$

Finally, we claim that no adversary can win against $G_{1,2^{\ell(\lambda)}}(\lambda)$. Let $x = (M, 1^t, m, \sigma)$ be the output of adversary \mathcal{D} in game $G_{1,2^{\ell(\lambda)}}(\lambda)$. Adversary \mathcal{D} wins the game if the following three conditions are simultaneously true: $M^0(x) \neq M_{1^\lambda, hk, h}^1(x)$, $M = M_{1^\lambda, vk}^{\text{ver}}$ and $|m| > \ell(\lambda)$. The first condition requires $M_{1^\lambda, hk, h}^1(x)$ to return 1. The second condition means that $M_{1^\lambda, vk}^{\text{ver}}(m, \sigma)$ will return the output of $M_{1^\lambda, vk}^{\text{ver}}(m, \sigma)$. However, according to the third condition, the latter returns 0. Therefore, for any $\lambda \in \mathbb{N}$ we have

$$\Pr[G_{1,2^{\ell(\lambda)}}(\lambda)] = 0. \quad (6)$$

We now show that there exists $\lambda_{\mathcal{D}} \in \mathbb{N}$ such that for all $\lambda \geq \lambda_{\mathcal{D}}$ we have $\text{Adv}_{\mathcal{S}^{\text{tm}}, \mathcal{D}}^{\text{diff}}(\lambda) \leq 2^{-\lambda^\epsilon}$. By definition, this means that \mathcal{S}^{tm} is $2^{-(\cdot)^\epsilon}$ -DIFF-secure.

$$\begin{aligned} \text{Adv}_{\mathcal{S}^{\text{tm}}, \mathcal{D}}^{\text{diff}}(\lambda) &= (\Pr[G_0(\lambda)] - \Pr[G_{1,0}(\lambda)]) \\ &+ \sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[G_{1,i}(\lambda)] - \Pr[G_{1,i+1}(\lambda)]) + \Pr[G_{1,2^{\ell(\lambda)}}(\lambda)] \quad (7) \\ &\leq \text{Adv}_{\mathcal{H}, \mathcal{H}}^{\text{tcr}}(\lambda) + 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}, \mathcal{S}_0^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda)) \\ &+ 2^{\ell(\lambda)+1} \cdot \text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(n(\lambda)) + 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}, \mathcal{S}_1^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda)) \quad (8) \\ &\leq 2^{-\lambda^{\epsilon_{\text{tcr}}}} + 2^{\ell(\lambda)} \cdot \left(2^{-n(\lambda)^{\epsilon_{\text{io}}}} + 2 \cdot 2^{-n(\lambda)^{\epsilon_{\text{psuf}}}} + 2^{-n(\lambda)^{\epsilon_{\text{io}}}} \right) \quad (9) \\ &\leq 2^{-\lambda^{\epsilon_{\text{tcr}}}} + 2^{\ell(\lambda)+1} \cdot \left(2^{-n(\lambda)^{\epsilon_{\text{io}}}} + 2^{-n(\lambda)^{\epsilon_{\text{psuf}}}} + 2^{-n(\lambda)^{\epsilon_{\text{io}}}} \right) \quad (10) \\ &\leq 2^{-\lambda^{2\epsilon}} + 2^{\ell(\lambda)+1} \cdot 3 \cdot 2^{-n(\lambda)^\epsilon} \quad (11) \\ &= 2^{-\lambda^{2\epsilon}} + 2^{\ell(\lambda)+1+\log_2 3-(2\lambda+\ell(\lambda)+3)^{\lceil 1/\epsilon \rceil \cdot \epsilon}} \quad (12) \\ &\leq 2^{-\lambda^{2\epsilon}} + 2^{-(2\lambda)^\epsilon} \quad (13) \\ &\leq 2^{-(2\lambda)^\epsilon} + 2^{-(2\lambda)^\epsilon} = 2^{1-(2\lambda)^\epsilon} \quad (14) \\ &\leq 2^{-\lambda^\epsilon}. \quad (15) \end{aligned}$$

Let $\lambda_{\mathcal{H}} \in \mathbb{N}$ such that $\text{Adv}_{\mathcal{H}, \mathcal{H}}^{\text{tcr}}(\lambda) \leq 2^{-\lambda^{\epsilon_{\text{tcr}}}}$ for all $\lambda \geq \lambda_{\mathcal{H}}$. Let $\lambda_{\mathcal{U}} \in \mathbb{N}$ such that $\text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(\lambda) \leq 2^{-\lambda^{\epsilon_{\text{psuf}}}}$ for all $\lambda \geq \lambda_{\mathcal{U}}$. For $b \in \{0, 1\}$ let $\lambda_{\mathcal{S}_b^{\text{tm}}, \mathcal{O}} \in \mathbb{N}$ be such that $\text{Adv}_{\text{Obf}_{\text{eq}}, \mathcal{S}_b^{\text{tm}}, \mathcal{O}}^{\text{io}}(\lambda) \leq 2^{-\lambda^{\epsilon_{\text{io}}}}$ for all $\lambda \geq \lambda_{\mathcal{S}_b^{\text{tm}}, \mathcal{O}}$.

Equation (7) follows from Equation (1) for all $\lambda \in \mathbb{N}$. Equation (8) follows from equations (2)–(6) for all $\lambda \in \mathbb{N}$. Equation (9) holds for all $\lambda \in \mathbb{N}$ such that

$\lambda \geq \lambda_{\mathcal{H}}$ and $n(\lambda) \geq \max(\lambda_{S_0^{\text{tm}}, \mathcal{O}}, \lambda_{\mathcal{U}}, \lambda_{S_1^{\text{tm}}, \mathcal{O}})$. Equation (10) holds for all $\lambda \in \mathbb{N}$. Equation (11) is obtained by expanding ϵ according to its definition, namely by using the following relations: $2\epsilon \leq \epsilon_{\text{tcr}}$, $\epsilon \leq \epsilon_{\text{psuf}}$ and $\epsilon \leq \epsilon_{\text{io}}$. Equation (12) is obtained by expanding $n(\lambda)$ according to its definition. Equation (13) holds for all $\lambda \in \mathbb{N}$, because for any polynomial $\ell: \mathbb{N} \rightarrow \mathbb{N}$, any constant $0 < \epsilon < 1$ and all $\lambda \in \mathbb{N}$ we have

$$\begin{aligned} & \ell(\lambda) + 1 + \log_2 3 - (2\lambda + \ell(\lambda) + 3)^{\lceil 1/\epsilon \rceil \cdot \epsilon} \\ & \leq \ell(\lambda) + 1 + \log_2 3 - (2\lambda + \ell(\lambda) + 3) \\ & < -2\lambda \leq -(2\lambda)^\epsilon. \end{aligned}$$

Equation (14) holds for all $\lambda \in \mathbb{N}$ such that $\lambda^{2\epsilon} \geq (2\lambda)^\epsilon$, requiring that $\lambda \geq 2$. Equation (15) holds for all $\lambda \in \mathbb{N}$ such that $1 - 2^\epsilon \lambda^\epsilon \leq -\lambda^\epsilon$, requiring that $\lambda \geq \left(\frac{1}{2^\epsilon - 1}\right)^{1/\epsilon}$. Therefore, it suffices to set

$$\lambda_{\mathcal{D}} = \max\left(\lambda_{\mathcal{H}}, \text{Inv}_n(1^{\lambda_{S_0^{\text{tm}}, \mathcal{O}}}), \text{Inv}_n(1^{\lambda_{\mathcal{U}}}), \text{Inv}_n(1^{\lambda_{S_1^{\text{tm}}, \mathcal{O}}}), 2, \left[(2^\epsilon - 1)^{-1/\epsilon}\right]\right).$$

This completes the proof. We now prove Claims 1-3.

Proof of Claim 1. We build a PT adversary \mathcal{H} against the TCR-security of \mathbf{H} such that for all $\lambda \in \mathbb{N}$ we have $\Pr[\mathbf{G}_0(\lambda)] - \Pr[\mathbf{G}_{1,0}(\lambda)] \leq \text{Adv}_{\mathbf{H}, \mathcal{H}}^{\text{tcr}}(\lambda)$.

Adversary $\mathcal{H}_1(1^\lambda)$ $(sk, vk) \leftarrow \text{DS.Kg}(1^{n(\lambda)})$ $st \leftarrow (sk, vk)$ Return $(M_{1^\lambda, vk}^{\text{ver}}, st)$	Adversary $\mathcal{H}_2(1^\lambda, st, hk)$ $(sk, vk) \leftarrow st$; $h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$ $M_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0)$; $M_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$ $aux \leftarrow \text{Obf}_{\text{eq}}^{\text{tm}}(1^{n(\lambda)}, \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk, vk}^{\text{aux}}))$ $(M, 1^t, m, \sigma) \leftarrow \mathcal{D}(1^\lambda, M_0, M_1, aux)$; Return M
---	---

Let $x = (M, 1^t, m, \sigma)$ be an output of adversary \mathcal{D} in games $\mathbf{G}_0(\lambda)$ and $\mathbf{G}_{1,0}(\lambda)$ (note that the input distribution of \mathcal{D} is the same in both games). If these games produce different outcomes for the same x , it means that $M^0(x) \neq M_{1^\lambda, hk, h}^1(x)$ and $M \neq M_{1^\lambda, vk}^{\text{ver}}$. According to the construction of M^0 and $M_{1^\lambda, hk, h}^1$ it follows that $\text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}}) = \text{H.Ev}(1^\lambda, hk, M)$. Whenever this happens, adversary \mathcal{H} wins in game $\text{TCR}_{\mathbf{H}}^{\mathcal{H}}(\lambda)$ by returning $x_0 = M_{1^\lambda, vk}^{\text{ver}}$ and $x_1 = M$. This proves the claim.

Proof of Claim 2. We build TM samplers $S_0^{\text{tm}}, S_1^{\text{tm}}$ and a PT adversary \mathcal{O} against the iO-security of $\text{Obf}_{\text{eq}}^{\text{tm}}$ relative to S_0^{tm} and S_1^{tm} , such that for all $\lambda \in \mathbb{N}$ we have $\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[\mathbf{G}_{1,i}(\lambda)] - \Pr[\mathbf{G}_{1,i,A}(\lambda)]) \leq 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}^{\text{tm}}, S_0^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda))$ and $\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[\mathbf{G}_{1,i,B}(\lambda)] - \Pr[\mathbf{G}_{1,i+1}(\lambda)]) \leq 2^{\ell(\lambda)} \cdot \text{Adv}_{\text{Obf}_{\text{eq}}^{\text{tm}}, S_1^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda))$.

Below, on the left we (simultaneously) define the TM samplers S_0^{tm} and S_1^{tm} that differ at the commented lines and have the uncommented lines in common. On the right, we define the PT adversary \mathcal{O} .

TM Samplers $S_0^{\text{tm}}(1^{\lambda'})$, $S_1^{\text{tm}}(1^{\lambda'})$	Adversary $\mathcal{O}(1^{\lambda'}, \bar{M}, aux)$
$\lambda \leftarrow \text{Inv}_n(1^{\lambda'})$; $i \leftarrow_{\$} \{0, 1\}^{\ell(\lambda)}$	$\lambda \leftarrow \text{Inv}_n(1^{\lambda'})$
$(sk, vk) \leftarrow_{\$} \text{DS.Kg}(1^{n(\lambda)})$	$a\tilde{u}x \leftarrow \bar{M}$
$hk \leftarrow_{\$} \text{H.Kg}(1^\lambda)$; $h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$	$(\tilde{M}_0, \tilde{M}_1, vk, z) \leftarrow aux$
$\tilde{M}_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0)$; $\tilde{M}_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$	$x \leftarrow_{\$} \mathcal{D}(1^\lambda, \tilde{M}_0, \tilde{M}_1, a\tilde{u}x)$
$m' \leftarrow \langle i \rangle_{\ell(\lambda)}$; $b \leftarrow M_{1^\lambda, sk, vk}^{\text{aux}}(m')$	$(M, 1^t, m, \sigma) \leftarrow x$
$m^* \leftarrow \langle m' \rangle_{\ell(n(\lambda))}$; $sk^* \leftarrow_{\$} \text{DS.PKg}(1^{n(\lambda)}, sk, m^*)$	$d_0 \leftarrow (\tilde{M}_0(x) \neq \tilde{M}_1(x))$
$M_{\text{aux}} \leftarrow \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk, vk}^{\text{aux}})$	$d_1 \leftarrow (M = M_{1^\lambda, vk}^{\text{ver}})$
$M_{\text{aux-punc}} \leftarrow \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}})$	If $(d_0 \wedge d_1 \wedge m \geq z)$
$M_1 \leftarrow M_{\text{aux}}$; $M_0 \leftarrow M_{\text{aux-punc}}$; $z \leftarrow i$ // S_0^{tm}	Then return 1
$M_0 \leftarrow M_{\text{aux}}$; $M_1 \leftarrow M_{\text{aux-punc}}$; $z \leftarrow i + 1$ // S_1^{tm}	Else return 0
$aux \leftarrow (M_0, M_1, vk, z)$; return (M_0, M_1, aux)	

We now show that $S_0^{\text{tm}}, S_1^{\text{tm}} \in \mathcal{S}_{\text{eq}}^{\text{tm}}$, meaning that these samplers produce functionally equivalent TMs. Both samplers return TMs $M_{1^\lambda, sk, vk}^{\text{aux}}$ and $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ that are padded to size $s_1(\lambda)$. First, observe that $M_{1^\lambda, sk, vk}^{\text{aux}}$ contains a signing key sk for DS, whereas $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ contains the corresponding punctured signing key sk^* , punctured at $m^* = \langle m' \rangle_{\ell(n(\lambda))}$, and a bit b that is equal to $M_{1^\lambda, sk, vk}^{\text{aux}}(m')$. According to the definition of a consistent puncturable digital signature scheme, keys sk and sk^* produce the same signatures for all $m \in \text{DS.In}(n(\lambda)) \setminus \{m^*\}$. Note that both $M_{1^\lambda, sk, vk}^{\text{aux}}$ and $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ compute a signature for an $\ell(n(\lambda))$ -bit string $\langle \bar{M} \rangle_{\ell(n(\lambda))}$ that is built from the $\ell(\lambda)$ -bit input string \bar{M} by padding it with leading zeros, which is an injective padding. Since m^* can only be built by padding m' , these TMs are equivalent for all inputs in $\bar{M} \in \{0, 1\}^{\ell(\lambda)} \setminus \{m^*\}$. Furthermore, notice that $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$ returns $b = M_{1^\lambda, sk, vk}^{\text{aux}}(m')$ on input m' , so these TMs are equivalent for *all* inputs.

Let $\lambda \in \mathbb{N}$. For any $b \in \{0, 1\}$ consider game $\text{IO}_{\text{Obf}_{\text{eq}}, S_b^{\text{tm}}}^{\mathcal{O}}(n(\lambda))$. Let i_b denote the value of i sampled by TM sampler S_b^{tm} . For any $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$ we have $\Pr[i_b = i] = 2^{-\ell(\lambda)}$, and hence

$$\begin{aligned}
\text{Adv}_{\text{Obf}_{\text{eq}}, S_b^{\text{tm}}, \mathcal{O}}^{\text{io}}(n(\lambda)) &= 2 \cdot \Pr[\text{IO}_{\text{Obf}_{\text{eq}}, S_b^{\text{tm}}}^{\mathcal{O}}(n(\lambda))] - 1 \\
&= 2 \cdot \sum_{i=0}^{2^{\ell(\lambda)}-1} \left(\Pr[i_b = i] \cdot \Pr[\text{IO}_{\text{Obf}_{\text{eq}}, S_b^{\text{tm}}}^{\mathcal{O}}(n(\lambda)) \mid i_b = i] \right) - 1 \\
&= 2 \cdot 2^{-\ell(\lambda)} \cdot \sum_{i=0}^{2^{\ell(\lambda)}-1} \Pr[\text{IO}_{\text{Obf}_{\text{eq}}, S_b^{\text{tm}}}^{\mathcal{O}}(n(\lambda)) \mid i_b = i] - 1. \quad (16)
\end{aligned}$$

Finally, observe that for any $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$ we have the following by construction:

$$\begin{aligned}
2 \cdot \Pr[\text{IO}_{\text{Obf}_{\text{eq}}, S_0^{\text{tm}}}^{\mathcal{O}}(n(\lambda)) \mid i_0 = i] - 1 &= \Pr[G_{1,i}(\lambda)] - \Pr[G_{1,i,A}(\lambda)], \\
2 \cdot \Pr[\text{IO}_{\text{Obf}_{\text{eq}}, S_1^{\text{tm}}}^{\mathcal{O}}(n(\lambda)) \mid i_1 = i] - 1 &= \Pr[G_{1,i,B}(\lambda)] - \Pr[G_{1,i+1}(\lambda)].
\end{aligned}$$

Claim 2 follows from Equation (16) together with the two equations above.

Proof of Claim 3. We build a PT adversary \mathcal{U} against the PSUFCMA-security of DS such that for all $\lambda \in \mathbb{N}$ we have $\sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[G_{1,i,A}(\lambda)] - \Pr[G_{1,i,B}(\lambda)]) \leq 2^{\ell(\lambda)+1} \cdot \text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(n(\lambda))$.

Adversary $\mathcal{U}_1(1^{\lambda'})$	Adversary $\mathcal{U}_2(1^{\lambda'}, st, vk, sk^*)$
$\lambda \leftarrow \text{Inv}_n(1^{\lambda'})$	$\lambda \leftarrow \text{Inv}_n(1^{\lambda'}); m' \leftarrow st; b \leftarrow_{\$} \{0, 1\}$
$m' \leftarrow_{\$} \{0, 1\}^{\ell(\lambda)}$	$hk \leftarrow_{\$} \text{H.Kg}(1^\lambda); h \leftarrow \text{H.Ev}(1^\lambda, hk, M_{1^\lambda, vk}^{\text{ver}})$
$m^* \leftarrow \langle m' \rangle_{\ell(n(\lambda))}$	$M_0 \leftarrow \text{Pad}_{s_0(\lambda)}(M^0); M_1 \leftarrow \text{Pad}_{s_0(\lambda)}(M_{1^\lambda, hk, h}^1)$
$st \leftarrow m'$	$aux \leftarrow_{\$} \text{Obf}_{\text{eq}}^{\text{tm}}(1^{n(\lambda)}, \text{Pad}_{s_1(\lambda)}(M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}))$
Return (m^*, st)	$x \leftarrow_{\$} \mathcal{D}(1^\lambda, M_0, M_1, aux); (M, 1^\ell, m, \sigma) \leftarrow x$
	$d_0 \leftarrow (M_0(x) \neq M_1(x)); d_1 \leftarrow (M = M_{1^\lambda, vk}^{\text{ver}})$
	If $(d_0 \wedge d_1 \wedge m = m')$ then return σ else return \perp

Let $\lambda \in \mathbb{N}$. Consider the value of m' sampled by \mathcal{U}_1 in game $\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(n(\lambda))$. For any $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$ it holds that $\Pr[m' = i] = 2^{-\ell(\lambda)}$. Hence,

$$\begin{aligned} \text{Adv}_{\text{DS}, \mathcal{U}}^{\text{psufcma}}(n(\lambda)) &= \sum_{i=0}^{2^{\ell(\lambda)}-1} (\Pr[m' = i] \cdot \Pr[\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(n(\lambda)) \mid m' = i]) \\ &= 2^{-\ell(\lambda)} \cdot \sum_{i=0}^{2^{\ell(\lambda)}-1} \Pr[\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(n(\lambda)) \mid m' = i]. \end{aligned} \quad (17)$$

Now observe that for any $i \in \{0, 1, \dots, 2^{\ell(\lambda)} - 1\}$ we also have

$$\Pr[\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(n(\lambda)) \mid m' = i] \geq \frac{1}{2} \cdot (\Pr[G_{1,i,A}(\lambda)] - \Pr[G_{1,i,B}(\lambda)]). \quad (18)$$

Let $x = (M, 1^\ell, m, \sigma)$ be an output of adversary \mathcal{D} in games $G_{1,i,A}(\lambda)$ and $G_{1,i,B}(\lambda)$ (note that the input distribution of \mathcal{D} is the same in both games). If these games produce different outcomes for the same x , it means that $M^0(x) \neq M_{1^\lambda, hk, h}^1(x)$, $M = M_{1^\lambda, vk}^{\text{ver}}$ and $m = i$. According to the construction of M^0 and $M_{1^\lambda, hk, h}^1$ it follows that $(\langle m \rangle_{\ell(n(\lambda))}, \sigma)$ is a valid message-signature pair for the digital signature scheme DS with verification key vk .

Whenever the above happens, adversary \mathcal{U} wins in game $\text{PSUFCMA}_{\text{DS}}^{\mathcal{U}}(n(\lambda))$ by forging a valid signature σ for message m^* , given that the following two conditions are satisfied. First, it is only true if adversary \mathcal{U} sampled $m' = i$. Second, in order to build TM $M_{1^\lambda, sk^*, vk, m', b}^{\text{aux-punc}}$, adversary \mathcal{U} has to compute $b = M_{1^\lambda, sk, vk}^{\text{aux}}(m')$. Since \mathcal{U} does not know sk , instead it has to guess the value of $b \in \{0, 1\}$. Hence, \mathcal{U} can perfectly simulate the games with probability $\frac{1}{2}$.

Claim 3 follows from Equation (17) and Equation (18).

Acknowledgments

Bellare and Stepanovs were supported in part by NSF grants CNS-1526801 and CNS-1228890, ERC Project ERCC FP7/615074 and a gift from Microsoft.

Waters was supported in part by NSF grants CNS-1228599 and CNS-1414082, DARPA SafeWare, a Google Faculty Research award, the Alfred P. Sloan Fellowship, a Microsoft Faculty Fellowship and a Packard Foundation Fellowship. We thank the Eurocrypt 2016 reviewers for their comments.

References

1. P. Ananth, D. Boneh, S. Garg, A. Sahai, and M. Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/2013/689>.
2. P. Ananth, Z. Brakerski, G. Segev, and V. Vaikuntanathan. From selective to adaptive security in functional encryption. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 657–677. Springer, Heidelberg, Aug. 2015.
3. B. Applebaum, B. Barak, and A. Wigderson. Public-key cryptography from different assumptions. In L. J. Schulman, editor, *42nd ACM STOC*, pages 171–180. ACM Press, June 2010.
4. M. Backes, O. Dagdelen, M. Fischlin, S. Gajek, S. Meiser, and D. Schröder. Operational signature schemes. Cryptology ePrint Archive, Report 2014/820, 2014. <http://eprint.iacr.org/2014/820>.
5. M. Backes, S. Meiser, and D. Schröder. Delegatable functional signatures. Cryptology ePrint Archive, Report 2013/408, 2013. <http://eprint.iacr.org/2013/408>.
6. B. Barak, S. Garg, Y. T. Kalai, O. Paneth, and A. Sahai. Protecting obfuscation against algebraic attacks. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 221–238. Springer, Heidelberg, May 2014.
7. B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, Aug. 2001.
8. M. Bellare. A note on negligible functions. *Journal of Cryptology*, 15(4):271–284, 2002.
9. M. Bellare and G. Fuchsbauer. Policy-based signatures. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 520–537. Springer, Heidelberg, Mar. 2014.
10. M. Bellare and P. Rogaway. Collision-resistant hashing: Towards making UOWHFs practical. In B. S. Kaliski Jr., editor, *CRYPTO’97*, volume 1294 of *LNCS*, pages 470–484. Springer, Heidelberg, Aug. 1997.
11. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
12. M. Bellare, I. Stepanovs, and S. Tessaro. Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 102–121. Springer, Heidelberg, Dec. 2014.
13. M. Bellare, I. Stepanovs, and B. Waters. New negative results on differing-inputs obfuscation. Cryptology ePrint Archive, Report 2016/162, 2016. <http://eprint.iacr.org/2016/162>.
14. N. Bitansky, R. Canetti, A. Chiesa, S. Goldwasser, H. Lin, A. Rubinfeld, and E. Tromer. The hunting of the SNARK. Cryptology ePrint Archive, Report 2014/580, 2014. <http://eprint.iacr.org/2014/580>.

15. N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In S. Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, Jan. 2012.
16. N. Bitansky, S. Garg, H. Lin, R. Pass, and S. Telang. Succinct randomized encodings and their applications. In R. A. Servedio and R. Rubinfeld, editors, *47th ACM STOC*, pages 439–448. ACM Press, June 2015.
17. N. Bitansky and O. Paneth. On the impossibility of approximate obfuscation and applications to resetttable cryptography. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 241–250. ACM Press, June 2013.
18. N. Bitansky, O. Paneth, and D. Wichs. Perfect structure on the edge of chaos - trapdoor permutations from indistinguishability obfuscation. In E. Kushilevitz and T. Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 474–502. Springer, Heidelberg, Jan. 2016.
19. D. Boneh and B. Waters. Constrained pseudorandom functions and their applications. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, Dec. 2013.
20. E. Boyle, K.-M. Chung, and R. Pass. On extractability obfuscation. In Y. Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 52–73. Springer, Heidelberg, Feb. 2014.
21. E. Boyle, S. Goldwasser, and I. Ivan. Functional signatures and pseudorandom functions. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, Mar. 2014.
22. C. Brzuska and A. Mittelbach. Using indistinguishability obfuscation via UCEs. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 122–141. Springer, Heidelberg, Dec. 2014.
23. R. Canetti, J. Holmgren, A. Jain, and V. Vaikuntanathan. Succinct garbling and indistinguishability obfuscation for RAM programs. In R. A. Servedio and R. Rubinfeld, editors, *47th ACM STOC*, pages 429–437. ACM Press, June 2015.
24. R. Canetti, H. Lin, S. Tessaro, and V. Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 468–497. Springer, Heidelberg, Mar. 2015.
25. N. Chandran, S. Raghuraman, and D. Vinayagamurthy. Constrained pseudorandom functions: Verifiable and delegatable. Cryptology ePrint Archive, Report 2014/522, 2014. <http://eprint.iacr.org/2014/522>.
26. G. Fuchsbauer. Constrained verifiable random functions. In M. Abdalla and R. D. Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 95–114. Springer, Heidelberg, Sept. 2014.
27. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, Oct. 2013.
28. S. Garg, C. Gentry, S. Halevi, and D. Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 518–535. Springer, Heidelberg, Aug. 2014.
29. C. Gentry, S. Halevi, M. Raykova, and D. Wichs. Outsourcing private RAM computation. In *55th FOCS*, pages 404–413. IEEE Computer Society Press, Oct. 2014.
30. C. Gentry, A. B. Lewko, A. Sahai, and B. Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In V. Guruswami, editor, *56th FOCS*, pages 151–170. IEEE Computer Society Press, Oct. 2015.

31. C. Gentry, A. B. Lewko, and B. Waters. Witness encryption from instance independent assumptions. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 426–443. Springer, Heidelberg, Aug. 2014.
32. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, Oct. 1986.
33. S. Goldwasser and Y. T. Kalai. On the impossibility of obfuscation with auxiliary input. In *46th FOCS*, pages 553–562. IEEE Computer Society Press, Oct. 2005.
34. S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, Apr. 1988.
35. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013.
36. S. Hada. Zero-knowledge and code obfuscation. In T. Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 443–457. Springer, Heidelberg, Dec. 2000.
37. I. Haitner, T. Holenstein, O. Reingold, S. P. Vadhan, and H. Wee. Universal one-way hash functions via inaccessible entropy. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 616–637. Springer, Heidelberg, May 2010.
38. Y. Ishai, O. Pandey, and A. Sahai. Public-coin differing-inputs obfuscation and its applications. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 668–697. Springer, Heidelberg, Mar. 2015.
39. J. Katz and C.-Y. Koo. On constructing universal one-way hash functions from arbitrary one-way functions. Cryptology ePrint Archive, Report 2005/328, 2005. <http://eprint.iacr.org/2005/328>.
40. A. Kiayias, S. Papadopoulos, N. Triandopoulos, and T. Zacharias. Delegatable pseudorandom functions and applications. In A.-R. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, Nov. 2013.
41. V. Koppula, A. B. Lewko, and B. Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In R. A. Servedio and R. Rubinfeld, editors, *47th ACM STOC*, pages 419–428. ACM Press, June 2015.
42. M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *21st ACM STOC*, pages 33–43. ACM Press, May 1989.
43. R. Pass, K. Seth, and S. Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 500–517. Springer, Heidelberg, Aug. 2014.
44. J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394. ACM Press, May 1990.
45. A. Sahai and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In D. B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
46. V. Shoup. A composition theorem for universal one-way hash functions. In B. Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 445–452. Springer, Heidelberg, May 2000.
47. B. Waters. A punctured programming approach to adaptively secure functional encryption. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 678–697. Springer, Heidelberg, Aug. 2015.