


Adaptive Witness Encryption and Asymmetric Password-Based Cryptography

Mihir Bellare¹ and Viet Tung Hoang^{2,3}

¹ Department of Computer Science and Engineering,
University of California San Diego, San Diego, USA
mihir@eng.ucsd.edu

² Department of Computer Science, Georgetown University,
Washington DC, USA

³ Department of Computer Science, University of Maryland,
Washington DC, USA

Abstract. We show by counter-example that the soundness security requirement for witness encryption given by Garg, Gentry, Sahai and Waters (STOC 2013) does not suffice for the security of their own applications. We introduce adaptively-sound (AS) witness encryption to fill the gap. We then introduce asymmetric password-based encryption (A-PBE). This offers gains over classical, symmetric password-based encryption in the face of attacks that compromise servers to recover hashed passwords. We distinguish between invasive A-PBE schemes (they introduce new password-based key-derivation functions) and non-invasive ones (they can use existing, deployed password-based key-derivation functions). We give simple and efficient invasive A-PBE schemes and use AS-secure witness encryption to give non-invasive A-PBE schemes.

1 Introduction

This paper introduces (1) witness encryption with adaptive soundness security and (2) asymmetric password-based encryption (A-PBE). We show how to use (1) to achieve (2) as well as other goals.

The Problem. The security of Internet communication remains ubiquitously based on client passwords. Standards such as the widely implemented PKCS#5—equivalently, RFC 2898 [33]—specify password-based encryption (PBE). From the client password pw , one derives a hashed password $hpw = \text{PH}(sa, pw)$, where sa is a random, user-specific public salt, and PH is a deterministic password-hashing function. (In the standards, $\text{PH}(sa, pw) = H^t(sa|pw)$ where t is an iteration count and H^t denotes the t -fold iteration of cryptographic hash function H .) The server holds hpw while the client holds (sa, pw) . Now the server will encrypt under hpw using any symmetric encryption scheme, for example CBC-AES. The client can recompute hpw from (sa, pw) and decrypt using this key.

This classical form of PBE is *symmetric*: encryption and decryption are both done under the same key hpw . But this means that anyone who knows hpw can

decrypt. This is a serious vulnerability in practice because of server compromise leading to exposure of hashed passwords. The Heartbleed attack of April 2014, allowing an attacker to read large chunks of server memory that can contain sensitive client information including hashed passwords, is a recent and prominent instance. Other high-profile attacks that compromised servers to expose client information include Target (December 2013), Adobe (October 2013), LinkedIn (June 2012), RSA (March 2011), Sony (2011) and TJ Maxx (2007). According to CNBC, there were over 600 breaches in 2013 alone.

We emphasize that the problem here is not the possibility of password-recovery via a dictionary attack based on the hashed password. The problem is that S-PBE (symmetric PBE) is vulnerable *even if the password is well chosen to resist dictionary attack*. This is because possession of the hashed password is already and directly enough to decrypt any prior communications. So under S-PBE, even well-chosen passwords do not provide security in the face of server compromise.

A-PBE. We propose asymmetric password-based encryption (A-PBE). Here, encryption is done under the hashed password h_{pw} , decryption is done under the password pw , and *possession of h_{pw} does not allow decryption*. This offers significantly higher security in the face of the most important attack, namely server compromise exposing the hashed password h_{pw} .

This paper initiates a foundational treatment of A-PBE including definitions and both “invasive” and “non-invasive” schemes. At first it may appear that definitionally A-PBE is just like PKE and brings nothing new, but this is not true. Not just is security based on passwords, but in practice users pick related passwords, for example varying a base password by appending the name of the website, resulting in encryption under related keys. Our definition extends the S-PBE framework of [10]. Our security model explicitly considers encryption under *multiple* passwords, assumed to be individually unpredictable —otherwise security is not possible— but arbitrarily related to each other.

We give two proven-secure A-PBE schemes that we call APBE1 and APBE2. Their attributes are summarized in Fig. 1. APBE1 is simple, natural and as efficient as possible, but what we call invasive, is that it specifies its own password-hashing function PH. APBE2 is non-invasive, meaning able to use any, given password-hashing function. In particular it can work with in-use, standardized password hashing functions such as PKCS#5 [33] or bcrypt [34]. If one has the flexibility of changing PH and the associated password hashes then the first solution is preferable. The second solution may be easier to deploy in the face of the legacy constraint of millions of existing, PKCS#5 hashed passwords.

APBE1. We specify and analyze the following simple and natural scheme for A-PBE that we call APBE1. PH, given sa, pw , applies to them a deterministic function EX to derive a string r , uses this as coin tosses for a key-generation algorithm PKE.Kg of some standard PKE scheme to get (pk, sk) , and outputs $h_{pw} = pk$ as the hashed password. Encryption is under the encryption algorithm PKE.Enc of the PKE scheme keyed with $h_{pw} = pk$. Since PH is deterministic,

Scheme	Invasive	Assumptions
APBE1	Yes	PKE, RIP-secure hash
APBE2	No	AS-secure WE RIP-secure password hash with large stretch
		XS-secure WE ROW-secure password hash with arbitrary stretch

Fig. 1. Our A-PBE schemes. Both achieve our notion of security for related, unpredictable passwords. APBE1 has a dedicated password hash (invasive) while APBE2 can work with an arbitrary, legacy one (non-invasive). The first analysis of APBE2 assumes the password hash has large stretch, a restriction dropped in the second analysis under a stronger form of WE.

decryption under (sa, pw) can re-execute PH to get (sk, pk) and then use sk to decrypt under PKE.

A natural choice for EX is a randomness extractor [32] with seed sa . But recall that we require A-PBE to be secure even under multiple, related passwords. To achieve this, outputs of EX must be independent even if the input passwords are related, and an extractor does not guarantee this. Indeed it is not possible for this to be true information theoretically, meaning if the “independence” is required to be statistical. We instead target computational independence of the outputs of EX. We define an appropriate security goal for EX that we call related-input pseudorandomness (RIP) [29] and show that this together with security of the base PKE scheme suffices for the security of the A-PBE scheme. In practice, EX can be efficiently instantiated via HMAC [5].

Non-invasive A-PBE. APBE1 prescribes its own password-hashing algorithm under which the hashed password h_{pw} is a public key of some existing PKE scheme. In current practice, however, the hashed password is derived via the iterated hashing password-hash function of PKCS#5 [33] or alternatives such as bcrypt [34]. Right now millions of passwords are in use with these particular password-hashing functions. In the face of this legacy constraint, deployment of A-PBE would be eased by a scheme that could encrypt under an existing, given hashed password, regardless of its form. We ask whether such non-invasive A-PBE is achievable.

This turns out to be challenging, even in principle, let alone in practice. In all known PKE schemes, the secret and public keys have very specific structure and are related in very particular ways. How can we encrypt asymmetrically with the public key being just an arbitrary hash of the secret key?

The answer is witness encryption (WE), introduced by Garg, Gentry, Sahai and Waters (GGSW) [20]. We will use WE to achieve non-invasive A-PBE. For this purpose, however, we will need WE schemes satisfying an extension of the soundness security notion of GGSW [20] that we introduce and call adaptive soundness security. We define and achieve WE with adaptive soundness and apply it to achieve non-invasive A-PBE as we now discuss.

SS-Secure Witness Encryption. In a WE scheme [20] for a language $L \in \mathbf{NP}$, the encryption function WE.Enc takes a unary representation 1^λ of the security parameter $\lambda \in \mathbb{N}$, a string $x \in \{0, 1\}^*$ and a message m to return a ciphertext c . If $x \in L$ then decryption is possible given a witness w for the membership of x in L . If $x \notin L$ then the message remains private given the ciphertext. The soundness security (SS) requirement of GGSW [20] formalized the latter by asking that for any PT adversary A , any $x \notin L$ and any equal-length messages m_0, m_1 , there is a negligible function ν such that

$$\Pr[A(\text{WE.Enc}(1^\lambda, x, m_1)) = 1] - \Pr[A(\text{WE.Enc}(1^\lambda, x, m_0)) = 1] \leq \nu(\lambda)$$

for all $\lambda \in \mathbb{N}$.

AS-Secure Witness Encryption. Our (new) adaptive soundness (AS) requirement lets the adversary A , on input 1^λ , pick and return x, m_0, m_1 to the game. The latter picks a random challenge bit b and returns $\text{WE.Enc}(1^\lambda, x, m_b)$ to A , who now responds with a guess b' as to the value of b . The AS-advantage of A is defined as the probability that $(b = b')$ and $x \notin L$. We require that any PT A have negligible advantage. We note that due to the check that $x \notin L$, our game may not be polynomial time but this does not hinder our applications.

It may at first seem that adaptivity does not add strength, since soundness security already quantifies over all x, m_0, m_1 . But in fact we show that AS is strictly stronger than SS. Namely we show in Proposition 2 that AS always implies SS but SS does not necessarily imply AS. That is, any WE scheme that is AS secure is SS secure, but there exist WE schemes that are SS secure and not AS secure. Intuitively, the reason AS is strictly stronger is that SS does not allow x, m_0, m_1 to depend on λ . Our separation result modifies a SS-secure WE scheme to misbehave when $|x| \geq f(\lambda)$ for a certain poly-logarithmic function f of the security parameter. SS is preserved because for each x only finitely many values of λ trigger the anomaly. The proof that AS is violated uses the fact that $\mathbf{NP} \subseteq \mathbf{EXP}$, the constructed adversary nonetheless being polynomial time.

Having strengthened the goal, we must revisit achievability. GGHRSW [18] give an elegant and conceptually simple construction of SS-secure WE from indistinguishability obfuscation (iO). In Theorem 3 we show that the same construction achieves the stronger AS goal. Recent work has provided constructions of iO improved both along the assumptions and efficiency fronts [2, 3, 16, 24], leading to corresponding improvements for AS-secure WE. Thus AS-secure WE can be achieved without loss of efficiency or added assumptions compared to SS-secure WE.

APBE2. Our APBE2 scheme lets L be the language of pairs $(sa, \text{PH}(sa, pw))$ over the choices of sa, pw , the witness being pw . A-PBE encryption of m using the hashed password as the public key will be AS-secure witness encryption of m under $x = (sa, hpw)$. Decryption will use the witness pw .

This solution is non-invasive, as it does not prescribe or require any particular design for PH. Rather, it takes PH as given, and shows how to encrypt with public key the hashed password obtained from PH. In this way, PH can in particular

be the iterated hash design of the PKCS#5 standard [33] that already underlies millions of usages of passwords, or any other practical, legacy design. Of course, for security, we will need to make an assumption about the security of PH, but that is very different from prescribing its design. Our assumption is the same RIP security as discussed above. We note that this assumption is already, even if implicitly, made in practice for the security of in-use S-PBE, where the hashed passwords are the keys, and is shown by [10] to hold for PKCS#5 in the ROM, so it is a natural and reasonable assumption.

SS Revisited. GGSW [20,21] present constructions of PKE, IBE and ABE schemes from witness encryption, claiming that these constructions are secure assuming soundness security of the WE scheme. The need for adaptive security of our A-PBE scheme leads to the natural question of why we need a stronger condition than GGSW [20,21]. The answer is that they need it too. We point out that the theorems of GGSW [20,21] claiming security of their applications under SS are incorrect, and that SS does not in fact suffice for the security of their schemes. We do this by presenting counter-examples (cf. Section 4). Taking their PRG-based PKE construction as a representative example, we provide a WE scheme which satisfies SS yet, if used in their construction, the resulting PKE scheme will provide no security at all. We then show that the gap can be filled by using AS. Namely, we show that their PKE scheme is secure if the underlying WE scheme is AS secure and the PRG is secure. Analogous results hold for GGSW's applications to IBE and ABE. Intuitively, the weakness of SS that compromises the applications of GGSW [20,21] is that a WE scheme may satisfy SS yet behave totally insecurely, for example returning the message in the clear, when $|x| = \lambda$. But in applications, x will have length related to λ , so SS is not enough. AS does not have this weakness because x can depend on λ .

Better Security for APBE2. Define the stretch of a password-hashing function as the difference between its output length and input length, and denote it by s . Our result of Theorem 5 proving the security of APBE2 requires that 2^{-s} is negligible, meaning the output length is somewhat more than the input length. This captures situations in which passwords are, say 12-character ASCII strings (input length is 78-bit) and the password hashing function is iterated SHA1 (output length is 160-bit). However, when passwords are longer, say 24-character, then passwords should offer more security. To fill this gap we offer a second analysis of the security of APBE2 that removes the restriction on the stretch, allowing it now to be arbitrary. For this purpose we strengthen the assumption on the WE scheme from AS to a notion of adaptive extractability we call XS. As a side benefit, the prior assumption on the password hashing function (RIP security, asking that password hashes are pseudorandom) is reduced to ROW security, asking merely that the password hashing function is one way.

XS is an adaptive variant of the notion of extractability from GKPVZ [26]. XS asks that, given an adversary violating the security of the encryption under $x \in \{0, 1\}^*$, one can extract a witness w for the membership of $x \in L$, even when x depends on the security parameter. We show that XS implies AS and

also that XS-secure WE can be achieved based on extractable (aka. differing-input) obfuscations [1, 4, 14].

Some works [15, 19] cast doubts on the achievability of extractable witness encryption or extractable iO with *arbitrary auxiliary inputs*. Our result however requires a very particular auxiliary input and the attacks in these works do not apply.

A-PBE as PKE. The standard model for public-key encryption (PKE) is that the user (receiver) publishes a public encryption key and stores the corresponding secret key securely. In practice, however, the secret key is often not stored in computer memory but instead derived from a password stored in human memory. Reasons this is advantageous include *security* and *mobility*. Computer-stored keys are vulnerable to exfiltration by malware. Meanwhile, users tend to have numerous devices including cellphones and tablets on which they want to decrypt. They may also use web-based services such as gmail on untrusted client machines. Passwords are more flexible and secure than stored keys in such settings.

A-PBE captures this more real-world PKE model. Our definitions allow us to evaluate security in the setting of actual use, namely when secret keys are possibly correlated passwords. Our schemes provide solutions with provable guarantees. We note that A-PBE is the model of the recently proposed gmail end-to-end encryption system, evidencing practical relevance of the goal.

Password-Based Signatures. Beyond A-PBE, we view this paper as initiating a study of asymmetric password-based cryptography. In this light we also introduce and treat password-based signatures with both invasive and non-invasive solutions to mirror the case of A-PBE.

Password-based authentication is currently done using a MAC keyed by the hashed password. It is thus subject to the same weakness as S-PBE, namely that compromise of the server through Heartbleed or other attacks leads to compromise of hashed passwords, resulting in compromise of the authentication. In the password-based signatures we suggest, one signs under the password pw and verifies under the hashed password $hpw = PH(sa, pw)$. Possession of the hashed password does not compromise security.

We can give a simple solution analogous to the one for A-PBE, namely apply a RIP function EX to the password and salt to get coins; run a key-generation of a standard digital signature scheme on these to get a signing key and verification key; set the password hash to the verification key; to sign given the password, re-generate the signing and verifying keys and sign under the former. This, however is invasive, prescribing its own password-hashing function. It is a good choice if one has the flexibility of implementing a new password hashing function, but as discussed above, deployment in the face of legacy PKCS#5 password hashes motivates asking whether a non-invasive solution, meaning one that can utilize any given password hashing function, is possible. As with A-PBE, this is a much more challenging question. We can show how to obtain a non-invasive password-based signature scheme by using key-versatile signatures [9].

The latter are effectively witness signatures meeting strong simulatability and extractability conditions [9, 17] and allow us to obtain password-based signatures analogous to how we obtained A-PBE from WE. The only assumption needed on the password hashing function PH is that it is one-way.

Discussion and GGSW Updates. A good definition for WE should have two properties: (1) *Usability*, meaning it suffices to prove security of applications, and (2) *Achievability*, meaning proposed and natural constructions, which in this case mainly means the iO-based one of GGHRWSW [18], can be shown to meet the definition. Our AS definition has both properties, making it viable. We have shown that SS lacked the usability property.

Here we have referred to the original GGSW STOC paper [20] and the corresponding original full ePrint version [21]. Subsequent to seeing prior versions of our paper, the GGSW authors updated their paper on ePrint [22, 23]. They acknowledge the gap we found. They also propose their own, modified definitions in an attempt to fill this gap.

Beyond (and despite) the fact that these updated definitions are subsequent to ours, they remain problematic. We showed that their first proposed definition, which we call SS2 [22], is unachievable. (Because the negligible function is not allowed to depend on the adversary. See Appendix A.) We communicated this to the authors. They then updated SS2 to SS3 [23]. But we explain in Appendix A that SS3 has limitations with regard to achievability. While one might of course propose still further modifications to their definition it is not clear why this is a productive route for the community in the face of the fact that, with AS, we have—and had prior to the GGSW updates—a definition that provides both usability and achievability.

Recently KNY [31] gave a definition, that we call SS5, in the quantifier style of SS1, SS2 and SS3. We discuss it also in Appendix A where we show that it is unachievable. (Because, like SS2, the negligible function doesn't depend on the adversary.)

These developments are an indication that neither the gap we find, nor the AS definition we propose to fill it, are trivial, that quantifier-based definitions are error-prone, and that our counter-examples for SS remain important to understand and guide definitional choices. Demonstrating the last, beyond [22, 23], further work subsequent to ours, and definitionally influenced by ours, includes [25].

We believe the idea of witness encryption is important and useful and we view our work as advancing its cause. Precision in definitions, proofs and details is particularly important in our field because we claim proven security. Reaching such precision can require iteration and definitional adjustments and increments, and our work, in this vein, helps towards greater impact and clarity for the area of witness encryption.

2 Preliminaries

By $\lambda \in \mathbb{N}$ we denote the security parameter and by 1^λ its unary representation. We denote the number of coordinates of a vector \mathbf{x} by $|\mathbf{x}|$, and the length of

a string $x \in \{0, 1\}^*$ by $|x|$. Algorithms are 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 $y \leftarrow^s A(x_1, \dots)$ be the resulting of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We say that $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every positive polynomial p , there exists $n_p \in \mathbb{N}$ such that $f(n) < 1/p(n)$ for all $n > n_p$. An adversary is an algorithm or a tuple of algorithms.

We use the code based game playing framework of [11]. For an example of a game see Fig. 2. By $G^A(\lambda)$ we denote the event that the execution of game G with adversary A and security parameter λ results in output true, the game output being what is returned by GAME.

Unpredictability. Let $A = (A_1, \dots)$ be a tuple of algorithms where A_1 , on input the unary representation 1^λ of the security parameter $\lambda \in \mathbb{N}$, returns a vector \mathbf{pw} . Let $\text{Guess}_A(\lambda)$ denote the maximum, over all i, pw , of $\Pr[\mathbf{pw}[i] = pw]$, the probability over $\mathbf{pw} \leftarrow^s A_1(\lambda)$. We say that A is *unpredictable* if the function $\text{Guess}_A(\cdot)$ is negligible.

3 Adaptive Witness Encryption

We begin by recalling the notion of witness encryption of GGSW [20] and their soundness security requirement. We then give a different security notion called adaptive soundness. We show that it is strictly stronger than the original, which means we must address achieving it. We show that it is achievable via indistinguishability obfuscation.

NP Relations. For $R : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\text{true}, \text{false}\}$, we let $R(x) = \{ w : R(x, w) \}$ be the *witness set* of $x \in \{0, 1\}^*$. We say R is an **NP**-relation if it is computable in PT and there is a polynomial $R.wl : \mathbb{N} \rightarrow \mathbb{N}$, called the witness length of R , such that $R(x) \subseteq \{0, 1\}^{R.wl(|x|)}$ for all $x \in \{0, 1\}^*$. We let $\mathcal{L}(R) = \{ x : R(x) \neq \emptyset \} \in \mathbf{NP}$ be the language defined by R .

WE Syntax and Correctness. A witness encryption (WE) scheme WE for $\overline{L} = \mathcal{L}(R)$ defines a pair of PT algorithms $WE.\text{Enc}, WE.\text{Dec}$. Algorithm $WE.\text{Enc}$ takes as input the unary representation 1^λ of a security parameter $\lambda \in \mathbb{N}$, a string $x \in \{0, 1\}^*$, and a message $m \in \{0, 1\}^*$, and outputs a ciphertext c . Algorithm $WE.\text{Dec}$ takes as input a string w and a ciphertext c , and outputs $m \in \{0, 1\}^* \cup \{\perp\}$. Correctness requires that $WE.\text{Dec}(w, WE.\text{Enc}(1^\lambda, x, m)) = m$ for all $\lambda \in \mathbb{N}$, all $x \in L$, all $w \in R(x)$ and all $m \in \{0, 1\}^*$.

Soundness Security. The soundness security (SS) condition of GGSW [20] says that for any PT adversary A , any $x \in \{0, 1\}^* \setminus L$ and any equal-length $m_0, m_1 \in \{0, 1\}^*$, there is a negligible function ν such that for all $\lambda \in \mathbb{N}$ we have

$$\Pr[A(WE.\text{Enc}(1^\lambda, x, m_1)) = 1] - \Pr[A(WE.\text{Enc}(1^\lambda, x, m_0)) = 1] < \nu(\lambda) . \quad (1)$$

GAME $AS_{WE,L}^A(\lambda)$

$(x, m_0, m_1, St) \leftarrow A(1^\lambda)$; $b \leftarrow \{0, 1\}$; $c \leftarrow WE(1^\lambda, x, m_b)$; $b' \leftarrow A(St, c)$

Return $((b = b') \wedge (x \notin L))$

Fig. 2. Game AS defining adaptive soundness of witness encryption scheme WE

In the following, it is useful to let $\text{Adv}_{WE,L,x,m_0,m_1,A}^{\text{SS}}(\lambda)$ denote the probability difference in (1). Then the soundness condition can be succinctly and equivalently stated as follows: WE is $\text{SS}[L]$ -secure if for any PT adversary A , any $x \in \{0, 1\}^* \setminus L$ and any equal-length $m_0, m_1 \in \{0, 1\}^*$, the function $\text{Adv}_{WE,L,x,m_0,m_1,A}^{\text{SS}}(\cdot)$ is negligible. It is convenient, in order to succinctly and precisely express relations between notions, to let $\text{SS}[L]$ denote the set of all correct witness encryption schemes that are $\text{SS}[L]$ -secure.

Adaptive Soundness. Our security definition associates to witness encryption scheme WE, language $L \in \text{NP}$, adversary A and $\lambda \in \mathbb{N}$ the game $AS_{WE,L}^A(\lambda)$ of Fig. 2. Here the adversary, on input 1^λ , produces instance x , messages m_0, m_1 , and state information St . It is required that $|m_0| = |m_1|$. The game picks a random challenge bit b and computes a ciphertext c via $WE.\text{Enc}(1^\lambda, x, m_b)$. The adversary is now given c , along with its state information St , and outputs a prediction b' for b . The game returns true if the prediction is correct, meaning $b = b'$, and also if $x \notin L$. We let $\text{Adv}_{WE,L,A}^{\text{AS}}(\lambda) = 2\Pr[AS_{WE,L}^A(\lambda)] - 1$. We say that WE has adaptive soundness security for L , or is $\text{AS}[L]$ -secure, if for every PT A the function $\text{Adv}_{WE,L,A}^{\text{AS}}(\cdot)$ is negligible. We let $\text{AS}[L]$ denote the set of all correct witness encryption schemes that are $\text{AS}[L]$ -secure.

Due to the check that $x \notin L$, our game does not necessarily run in PT. This, however, will not preclude applicability. The difference between AS and SS is that in the former, x, m_0, m_1 can depend on the security parameter and on each other. Given that SS quantifies over all x, m_0, m_1 , this may not at first appear to make any difference. But we will see that it does and that AS is strictly stronger than SS.

AS is a game-based definition while SS is phrased in a more “quantifier-based” style that mimics the soundness condition in interactive proofs [28]. The game-based AS notion is better suited for applications because the latter are also underlain by game-based definitions. Indeed we’ll see that SS does not suffice for applications.

A Useful Transform. In several proofs, we’ll employ the following transform. Given a WE scheme $WE \in \text{SS}[L]$ and a PT function $f: \mathbb{N} \rightarrow \mathbb{N}$, our transform returns another WE scheme WE_f . The constructed scheme, formally specified in Fig. 3, misbehaves, returning the message in the clear, when $|x| \geq f(\lambda)$, and otherwise behaves like WE. The following says that if f is chosen to satisfy certain conditions then $\text{SS}[L]$ -security is preserved, meaning $WE_f \in \text{SS}[L]$. In our uses of the transform we will exploit the fact that WE_f will fail to have other security properties or lead to failure of applications that use it.

$\text{WE}_f.\text{Enc}(1^\lambda, x, m)$	$\text{WE}_f.\text{Dec}(w, c)$
If $ x \geq f(\lambda)$ then return $(0, m)$	$(b, t) \leftarrow c$
Else return $(1, \text{WE}.\text{Enc}(1^\lambda, x, m))$	If $b = 0$ then return t else return $\text{WE}.\text{Dec}(w, t)$

Fig. 3. Witness encryption scheme WE_f for $L \in \mathbf{NP}$, derived from $\text{WE} \in \text{SS}[L]$ and a PT-computable function $f : \mathbb{N} \rightarrow \mathbb{N}$

Lemma 1. Let $L \in \mathbf{NP}$ and $\text{WE} \in \text{SS}[L]$. Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a non-decreasing, PT-computable function such that $\lim_{\lambda \rightarrow \infty} f(\lambda) = \infty$. Consider witness encryption scheme WE_f derived from WE and f as shown in Fig. 3. Then $\text{WE}_f \in \text{SS}[L]$.

Proof. Let A be a PT adversary. Let $x \in \{0, 1\}^* \setminus L$ and let $m_0, m_1 \in \{0, 1\}^*$ have equal length. Let PT adversary B , on input ciphertext c , return $b' \leftarrow A((1, c))$. Let $S(x) = \{ \lambda \in \mathbb{N} : f(\lambda) \leq |x| \}$. Then for all $\lambda \in \mathbb{N} \setminus S(x)$ we have $\text{Adv}_{\text{WE}, L, x, m_0, m_1, B}^{\text{ss}}(\lambda) = \text{Adv}_{\text{WE}_f, L, x, m_0, m_1, A}^{\text{ss}}(\lambda)$. The assumption that $\text{WE} \in \text{SS}[L]$ means that $\text{Adv}_{\text{WE}, L, x, m_0, m_1, B}^{\text{ss}}(\cdot)$ is negligible. But the assumptions on f mean that the set $S(x)$ is finite. Consequently, the function $\text{Adv}_{\text{WE}_f, L, x, m_0, m_1, A}^{\text{ss}}(\cdot)$ is negligible as well. \square

Relations. We show that adaptive soundness implies soundness but not vice versa, meaning adaptive soundness is a strictly stronger requirement.

Proposition 2. Let $L \in \mathbf{NP}$. Then: (1) $\text{AS}[L] \subseteq \text{SS}[L]$, and (2) If $\{0, 1\}^* \setminus L$ is infinite and $\text{SS}[L] \neq \emptyset$ then $\text{SS}[L] \not\subseteq \text{AS}[L]$.

Claim (1) above says that any witness encryption scheme WE that is $\text{AS}[L]$ -secure is also $\text{SS}[L]$ -secure. Claim (2) says that the converse is not true. Namely, there is a witness encryption scheme WE such that WE is $\text{SS}[L]$ -secure but not $\text{AS}[L]$ -secure. This separation assumes some $\text{SS}[L]$ -secure witness encryption scheme exists, for otherwise the claim is moot. It also assumes that the complement of L is not trivial, meaning is infinite, which is true if L is \mathbf{NP} -complete and $\mathbf{P} \neq \mathbf{NP}$, hence is not a strong assumption.

Proof (Proposition 2). For part (1), assume we are given WE that is $\text{AS}[L]$ -secure. We want to show that WE is $\text{SS}[L]$ -secure. Referring to the definition of soundness security, let A be a PT adversary, let $x \in \{0, 1\}^* \setminus L$ and let $m_0, m_1 \in \{0, 1\}^*$ have equal length. We want to show that the function $\text{Adv}_{\text{WE}, L, x, m_0, m_1, A}^{\text{ss}}(\cdot)$ is negligible. We define the adversary B_{x, m_0, m_1} as follows: Let $B_{x, m_0, m_1}(1^\lambda)$ return $(x, m_0, m_1, \varepsilon)$ and let $B_{x, m_0, m_1}(t, c)$ return $b' \leftarrow A(c)$. Here, B_{x, m_0, m_1} has x, m_0, m_1 hardwired in its code, and, in its first stage, it returns them, along with $\text{St} = \varepsilon$ as state information. In its second stage, it simply runs A . Note that even though B_{x, m_0, m_1} has hardwired information, this information is finite and not dependent on the security parameter, so the hardwiring does not require non-uniformity. Now it is easy to see that for all $\lambda \in \mathbb{N}$ we have $\text{Adv}_{\text{WE}, L, B_{x, m_0, m_1}}^{\text{as}}(\lambda) = \text{Adv}_{\text{WE}, L, x, m_0, m_1, A}^{\text{ss}}(\lambda)$. The assumption

that WE is AS[L]-secure means that $\text{Adv}_{\text{WE},L,B_x,m_0,m_1}^{\text{as}}(\cdot)$ is negligible, hence so is $\text{Adv}_{\text{WE},L,x,m_0,m_1,A}^{\text{ss}}(\cdot)$, as desired.

For part (2), the assumption $\text{SS}[L] \neq \emptyset$ means there is some $\text{WE} \in \text{SS}[L]$. By way of Lemma 1, we can modify it to $\text{WE}_f \in \text{SS}[L]$ as specified in Fig. 3, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is some non-decreasing, PT-computable function such that $\lim_{\lambda \rightarrow \infty} f(\lambda) = \infty$. Now we want to present an attacker A violating AS[L]-security of WE_f . The difficulty is that A needs to find $x \notin L$ of length $f(\lambda)$, but $L \in \mathbf{NP}$ and A must be PT. We will exploit the fact that $\mathbf{NP} \subseteq \mathbf{EXP}$ and pick f to be a poly-logarithmic function related to the exponential time to decide L , so that if there exists an $x \notin L$ of length $f(\lambda)$ then A can find it by exhaustive search in PT. Our assumption that the complement of L is infinite means that A succeeds on infinitely many values of λ .

Proceeding to the details, since $L \in \mathbf{NP} \subseteq \mathbf{EXP}$, there is a constant $d \geq 1$ and a deterministic algorithm M such that for every $x \in \{0, 1\}^*$, we have $M(x) = 1$ if and only if $x \in L$, and M 's running time is $\mathcal{O}(2^{|x|^d})$. Define f by $f(\lambda) = \lfloor \lg^{1/d}(\lambda) \rfloor$ for all $\lambda \in \mathbb{N}$. Let $\text{WE} \in \text{SS}[L]$ and let WE_f be the witness encryption scheme derived from WE and f as specified in Fig. 3. By Lemma 1, $\text{WE}_f \in \text{SS}[L]$. Now we show that $\text{WE}_f \notin \text{AS}[L]$. Let $m_0, m_1 \in \{0, 1\}^*$ be arbitrary, distinct, equal-length messages. Consider the following adversary A :

$A(1^\lambda)$ $k \leftarrow f(\lambda); x \leftarrow 0^k$ For all $s \in \{0, 1\}^k$ do If $(M(s) \neq 1)$ then $x \leftarrow s$ Return $(x, m_0, m_1, \varepsilon)$	$A(t, c)$ $(b, m) \leftarrow c$ If $((b = 0) \wedge (m = m_1))$ then return 1 Return 0
---	---

Each execution of M takes time $\mathcal{O}(2^{k^d}) = \mathcal{O}(\lambda)$. The For loop goes through all $s \in \{0, 1\}^k$ in lexicographic order and thus M is executed at most $2^k \leq \lambda$ times. So A is PT. For any $\lambda \in \mathbb{N}$, if $\{0, 1\}^{f(\lambda)} \setminus L \neq \emptyset$ then $\text{Adv}_{\text{WE}_f,L,A}^{\text{as}}(\lambda) = 1$. Since $\{0, 1\}^* \setminus L$ is infinite, f is non-decreasing, and $\lim_{t \rightarrow \infty} f(t) = \infty$, there are infinitely many values λ such that $\text{Adv}_{\text{WE}_f,L,A}^{\text{as}}(\lambda) = 1$, and thus $\text{WE}_f \notin \text{AS}[L]$, as claimed. □

Indistinguishability Obfuscation. We say that two circuits C_0 and C_1 are *functionally equivalent*, denoted $C_0 \equiv C_1$, if they have the same size, the same number n of inputs, and $C_0(x) = C_1(x)$ for every input $x \in \{0, 1\}^n$. An obfuscator \mathbf{P} defines PT algorithms $\mathbf{P}.\text{Ob}, \mathbf{P}.\text{Ev}$. Algorithm $\mathbf{P}.\text{Ob}$ takes as input the unary representation 1^λ of a security parameter λ and a circuit C , and outputs a string c . Algorithm $\mathbf{P}.\text{Ev}$ takes as input strings c, x and returns $y \in \{0, 1\}^* \cup \{\perp\}$. We require that for any circuit C , any input x , and any $\lambda \in \mathbb{N}$, it holds that $\mathbf{P}.\text{Ev}(x, \mathbf{P}.\text{Ob}(1^\lambda, C)) = C(x)$. We say that \mathbf{P} is iO-secure if $\text{Adv}_{\mathbf{P},A}^{\text{iO}}(\lambda) = 2 \Pr[\text{IO}_{\mathbf{P}}^A(\lambda)] - 1$ is negligible for every PT adversary A , where game IO is defined at Fig. 4. This definition is slightly different from the notion in [4, 18]—the adversary is non-uniform and must produce functionally equivalent circuits C_0 and C_1 —but the former definition is implied by the latter.

$\begin{array}{l} \text{GAME IO}_P^A(\lambda) \\ (C_0, C_1, \text{St}) \leftarrow_s A(1^\lambda); b \leftarrow_s \{0, 1\}; c \leftarrow_s \text{P.Ob}(1^\lambda, C_b) \\ b' \leftarrow_s A(\text{St}, c); \text{Return } (b = b') \wedge (C_0 \equiv C_1) \end{array}$
--

Fig. 4. Game IO defining security of an indistinguishability obfuscator P

Achieving AS-Security. Our AS security notion is strictly stronger than the SS one of GGSW [20], but we'll show that the iO-based WE scheme of [18] is AS-secure. Proceeding to the details, let R be an NP-relation. For each $x, m \in \{0, 1\}^*$, let $R_{x,m}$ be a circuit that, on input $w \in \{0, 1\}^{\text{R.wl}(|x|)}$, returns m if $R(x, w)$ and returns $0^{|m|}$ otherwise. Let P be an indistinguishability obfuscator, defining a PT obfuscation algorithm P.Ob and a PT evaluation algorithm P.Ev. We define WE scheme $\text{WE}_R[\text{P}]$ as follows: algorithm $\text{WE}_R[\text{P}].\text{Enc}(1^\lambda, x, m)$ returns $c \leftarrow_s \text{P.Ob}(1^\lambda, R_{x,m})$; and algorithm $\text{WE}_R[\text{P}].\text{Dec}(w, c)$ returns $m \leftarrow_s \text{P.Ev}(w, c)$.

Theorem 3. Let R be an NP-relation and let $L = \mathcal{L}(\text{R})$. Let P be an indistinguishability obfuscator. Construct $\text{WE}_R[\text{P}]$ as above. If P is iO-secure then $\text{WE}_R[\text{P}] \in \text{AS}[L]$.

Proof. Let A be a PT adversary attacking the AS[L]-security of $\text{WE}_R[\text{P}]$. Wlog, assume that A produces distinct m_0 and m_1 . Note that $R_{x,m_0} \equiv R_{x,m_1}$ if and only if $x \notin L$. Consider the following PT adversary B attacking iO-security of P:

$$\begin{array}{l|l} \frac{B(1^\lambda)}{(x, m_0, m_1, \text{St}) \leftarrow_s A(1^\lambda)} & \frac{B(\text{St}, c)}{b' \leftarrow_s A(\text{St}, c); \text{Return } b'} \\ \text{Return } (R_{x,m_0}, R_{x,m_1}, \text{St}) & \end{array}$$

Then $\Pr[\text{AS}_{\text{WE}_R[\text{P}],L}^A(\cdot)] = \Pr[\text{IO}_P^B(\cdot)]$ and thus $\text{Adv}_{\text{WE}_R[\text{P}],L,A}^{\text{as}}(\cdot) = \text{Adv}_{P,B}^{\text{iO}}(\cdot)$. \square

4 Insufficiency of Soundness Security

GGSW [20] present constructions of several primitives from witness encryption, including PKE, IBE and ABE for all circuits. They claim security of these constructions assuming soundness security of the underlying witness-encryption scheme. We observe here that these claims are wrong. Taking their PRG-based PKE scheme as a representative example, we present a counter-example, namely a witness-encryption scheme satisfying soundness security such that the PKE scheme built from it is insecure. Similar counter-examples can be built for the other applications in GGSW [20]. Briefly, the problem is that a witness encryption scheme could fail to provide any security when $|x|$ is equal to, or related in some specific way to, the security parameter, yet satisfy SS security because the latter requirement holds x fixed and lets λ go to ∞ . We show that the gap can be filled, and all the applications of GGSW recovered, by using adaptive soundness in place of soundness security. We'll begin by recalling the well-known notions of PRG and PKE.

<p style="text-align: center; margin: 0;"><u>GAME PRG_G^A(λ)</u></p> <p style="margin: 0;">$s \leftarrow \{0, 1\}^\lambda ; x_1 \leftarrow G(s)$</p> <p style="margin: 0;">$x_0 \leftarrow \{0, 1\}^{\ell(\lambda)} ; b \leftarrow \{0, 1\}$</p> <p style="margin: 0;">$b' \leftarrow A(1^\lambda, x_b) ; \text{Return } (b = b')$</p>	<p style="text-align: center; margin: 0;"><u>GAME INDCPA_{PKE}^A(λ)</u></p> <p style="margin: 0;">$(pk, sk) \leftarrow \text{PKE.Kg}(1^\lambda) ; b \leftarrow \{0, 1\}$</p> <p style="margin: 0;">$b' \leftarrow A^{\text{LR}}(1^\lambda, pk) ; \text{Return } (b = b')$</p> <hr style="border: 0; border-top: 1px solid black; margin: 5px 0;"/> <p style="text-align: center; margin: 0;"><u>LR(m_0, m_1)</u></p> <p style="margin: 0;">$c \leftarrow \text{PKE.Enc}(pk, m_b) ; \text{Return } c$</p>
---	--

Fig. 5. **Left:** Game PRG defining security of a pseudorandom generator G . Here $\ell : \mathbb{N} \rightarrow \mathbb{N}$ is the expansion factor of G . **Right:** Game INDCPA defining INDCPA security of a PKE scheme PKE. For each oracle query, the messages $m_0, m_1 \in \{0, 1\}^*$ must have the same length.

<p style="text-align: center; margin: 0;"><u>PKE.Kg(1^λ)</u></p> <p style="margin: 0;">$sk \leftarrow \{0, 1\}^\lambda ; x \leftarrow G(sk)$</p> <p style="margin: 0;">$pk \leftarrow (\lambda, x) ; \text{Return } (pk, sk)$</p>	<p style="text-align: center; margin: 0;"><u>PKE.Enc(pk, m)</u></p> <p style="margin: 0;">$(\lambda, x) \leftarrow pk$</p> <p style="margin: 0;">$\text{Return } \overline{\text{WE}}.\text{Enc}(1^\lambda, x, m)$</p>	<p style="text-align: center; margin: 0;"><u>PKE.Dec(sk, c)</u></p> <p style="margin: 0;">$\text{Return } \overline{\text{WE}}.\text{Dec}(c, sk)$</p>
--	---	---

Fig. 6. GGSW’s PKE scheme $\text{PKE}[G, \overline{\text{WE}}]$, where G is a length-doubling PRG and $\overline{\text{WE}}$ is a witness encryption scheme for $L_G = \{ G(s) : s \in \{0, 1\}^* \}$

Primitives. A pseudorandom generator (PRG) [12, 37] is a PT deterministic algorithm G that takes any string $s \in \{0, 1\}^*$ as input and return a string $G(s)$ of length $\ell(|s|)$, where the function $\ell : \mathbb{N} \rightarrow \mathbb{N}$ is call the *expansion factor* of G . We say that G is secure if $\text{Adv}_{A,G}^{\text{PRG}}(\lambda) = 2 \Pr[\text{PRG}_A^G(\lambda)] - 1$ is negligible, for every PT adversary A , where game PRG is defined in Fig. 5.

A public-key encryption (PKE) scheme PKE defines PT algorithms PKE.Kg, PKE.Enc, PKE.Dec, the last deterministic. Algorithm PKE.Kg takes as input 1^λ and outputs a public encryption key pk and a secret decryption key sk . Algorithm PKE.Enc takes as input pk and a message $m \in \{0, 1\}^*$, and outputs a ciphertext c . Algorithm PKE.Dec(sk, c) outputs $m \in \{0, 1\}^* \cup \{\perp\}$. Scheme PKE is INDCPA-secure [6, 27] if $\text{Adv}_{\text{PKE}, A}^{\text{ind-cpa}}(\cdot) = 2 \Pr[\text{INDCPA}_{\text{PKE}}^A(\cdot)] - 1$ is negligible for every PT adversary A , where game INDCPA is defined in Fig. 5.

SS does not Suffice for GGSW’s PKE Scheme. Let G be a PRG that is length doubling, meaning $|G(s)| = 2|s|$ for every $s \in \{0, 1\}^*$. Let $L_G = \{ G(s) : s \in \{0, 1\}^* \}$. This language is in NP. Let $\overline{\text{WE}} \in \text{SS}[L_G]$ be a $\text{SS}[L_G]$ -secure WE scheme. The PKE scheme $\text{PKE}[G, \overline{\text{WE}}]$ of GGSW is shown in Fig. 6. We claim that $\text{SS}[L_G]$ -security of $\overline{\text{WE}}$ is insufficient for PKE to be INDCPA-secure. We show this by counter-example, meaning we give an example of a particular WE scheme $\overline{\text{WE}} \in \text{SS}[L_G]$ such that $\text{PKE}[G, \overline{\text{WE}}]$ is not INDCPA. We assume there exists some $\overline{\text{WE}} \in \text{SS}[L_G]$, else the question is moot. Let $f(\lambda) = 2\lambda$ for every $\lambda \in \mathbb{N}$. Now let $\overline{\text{WE}} = \text{WE}_f$ be the WE scheme of Fig. 3 obtained from WE and f . Lemma 1 tells us that $\text{WE}_f \in \text{SS}[L_G]$. Now we claim that $\text{PKE}[G, \text{WE}_f]$ is not INDCPA. The reason is that when $\text{PKE.Enc}(pk, m)$ runs $\text{WE}_f.\text{Enc}(1^\lambda, x, m)$, we have $|x| = 2\lambda = f(\lambda)$. By definition of $\text{WE}_f.\text{Enc}$, the latter returns $(0, m)$ as the ciphertext, effectively sending the message in the clear.

AS Security Suffices for GGSW’s PKE. We now show that the gap can be filled using AS. That is, we prove that if G is a secure PRG and \overline{WE} is $AS[L_G]$ -secure, then $PKE[G, \overline{WE}]$ is IND CPA-secure:

Theorem 4. Let $G : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a length-doubling PRG. Let $L_G = \{G(s) : s \in \{0, 1\}^*\}$. If G is a secure PRG and $\overline{WE} \in AS[L_G]$ then $PKE[G, \overline{WE}]$ is IND CPA-secure.

The proof is in [7]. It follows the template of the proof of GGSW [20]. First one uses the PRG security of G to move to a game where x is random. Since G is length doubling, such an x is not in L_G with high probability. At this point GGSW [20] (incorrectly) claim that the result follows from the $SS[L_G]$ -security of \overline{WE} . We instead use the $AS[L_G]$ -security of \overline{WE} , providing a reduction with an explicit construction of an AS adversary.

To obtain similar counter-examples showing the inadequacy of SS for the other applications of GGSW (namely IBE and ABE for all circuits), one can follow the template of our PKE attack, by choosing a lower bound $f(\lambda)$ for the length of the string $x = X(\lambda)$ given to the witness encryption. Since $X(\lambda)$ is generated from some cryptographic primitive π (for example, in IBE, π is a unique signature scheme), the security of π requires that $X(\lambda)$ have super-logarithmic length. Hence there is a constant $C > 0$ such that $|X(\lambda)| \geq C \lg(\lambda)$ for all $\lambda \in \mathbb{N}$, and therefore we can let $f(\lambda) = \lfloor C \lg(\lambda) \rfloor$.

5 Asymmetric Password-Based Encryption

In this Section we introduce the new primitive of asymmetric password-based encryption (A-PBE). We then provide a non-invasive, WE-based A-PBE scheme we call APBE2, with two security analyses. First we prove security of APBE2 under AS-security of the WE scheme. Then under XS-security of the WE scheme we provide another proof that shows the scheme to admit better “stretch,” leading to better security for some real password distributions. In [7] we provide a simple and fast, but invasive, A-PBE scheme, called APBE1. Our model and definitions are of interest beyond our schemes because they capture PKE in the real-world setting where secret keys are based on passwords and may thus be related.

A-PBE Syntax and Security. An *asymmetric password-based encryption* (A-PBE) scheme F specifies PT algorithms $F.Ph, F.Enc, F.Dec$, the first and the last deterministic. It also specifies a password-length function $F.pl : \mathbb{N} \rightarrow \mathbb{N}$, a salt-length function $F.sl : \mathbb{N} \rightarrow \mathbb{N}$, and a hash-length function $F.hl : \mathbb{N} \rightarrow \mathbb{N}$. Algorithm $F.Ph$ takes as input the unary representation 1^λ of security parameter λ , a salt $sa \in \{0, 1\}^{F.sl(\lambda)}$, and a password $pw \in \{0, 1\}^{F.pl(\lambda)}$, and returns a hashed password $hpw = F.Ph(1^\lambda, sa, pw) \in \{0, 1\}^{F.hl(\lambda)}$. Algorithm $F.Enc$ takes as input $1^\lambda, hpw, sa$ and a message $m \in \{0, 1\}^*$, and outputs a ciphertext c . Finally, given (pw, c) , algorithm $F.Dec$ returns $m \in \{0, 1\}^* \cup \{\perp\}$. We require that

$$F.Dec(pw, F.Enc(1^\lambda, F.Ph(1^\lambda, sa, pw), sa, m)) = m$$

for every $m \in \{0, 1\}^*, \lambda \in \mathbb{N}, sa \in \{0, 1\}^{F.sl(\lambda)}$, and $pw \in \{0, 1\}^{F.pl(\lambda)}$.

<u>GAME APBE_F^A(λ)</u>	<u>GAME RIP_H^A(λ)</u>
$\mathbf{pw} \leftarrow_{\$} A_1(1^\lambda); b \leftarrow_{\$} \{0, 1\}$	$\mathbf{pw} \leftarrow_{\$} A_1(1^\lambda); b \leftarrow_{\$} \{0, 1\}$
For $i = 1$ to $ \mathbf{pw} $ do	For $i = 1$ to $ \mathbf{pw} $ do
$\mathbf{sa}[i] \leftarrow_{\$} \{0, 1\}^{\text{F.sl}(\lambda)}$	$\mathbf{sa}[i] \leftarrow_{\$} \{0, 1\}^{\text{H.kl}(\lambda)}$
$\mathbf{hpw}[i] \leftarrow \text{F.Ph}(1^\lambda, \mathbf{sa}[i], \mathbf{pw}[i])$	$\mathbf{hpw}[i] \leftarrow \text{H}(1^\lambda, \mathbf{sa}[i], \mathbf{pw}[i])$
$b' \leftarrow_{\$} A_2^{\text{LR}}(1^\lambda, \mathbf{sa}, \mathbf{hpw})$	If $b = 0$ then $\mathbf{hpw}[i] \leftarrow_{\$} \{0, 1\}^{\text{H.ol}(\lambda)}$
Return ($b = b'$)	$b' \leftarrow_{\$} A_2(1^\lambda, \mathbf{sa}, \mathbf{hpw});$ Return ($b = b'$)
<u>LR(m_0, m_1, i)</u>	
$c \leftarrow_{\$} \text{F.Enc}(1^\lambda, \mathbf{hpw}[i], \mathbf{sa}[i], m_b)$	
Return c	

Fig. 7. Left: Game APBE defining security of an A-PBE scheme F . **Right:** Game RIP defining RIP security for a hash family H .

An adversary A is a pair of PT algorithms (A_1, A_2) . Adversary $A_1(1^\lambda)$ generates a vector of passwords \mathbf{pw} , each entry a $\text{F.pl}(\lambda)$ -bit string. It is required that A is unpredictable as defined in Section 2. Note that passwords —entries of the vector \mathbf{pw} — may be correlated, even though each individually is unpredictable, to capture the fact that individual users often pick related passwords for their different accounts. We say that A-PBE scheme F is secure if $\text{Adv}_{F,A}^{\text{apbe}}(\cdot) = 2 \Pr[\text{APBE}_F^A(\cdot)] - 1$ is negligible for every PT unpredictable adversary A , where game $\text{APBE}_F^A(\lambda)$ is defined in Fig. 7. In this game, $A_1(1^\lambda)$ first generates its vector \mathbf{pw} of passwords. The game picks a challenge bit $b \leftarrow_{\$} \{0, 1\}$ and a vector of random salts \mathbf{sa} . Adversary A_2 is given \mathbf{sa} and the vector \mathbf{hpw} of hashed passwords. It can then query its oracle LR with equal-length, distinct messages m_0, m_1 , and an index i , to get $\text{F.Enc}(1^\lambda, \mathbf{hpw}[i], \mathbf{sa}[i], m_b)$. Finally A_2 outputs a prediction b' for b . The game returns true if the prediction is correct, meaning $b = b'$, and false otherwise.

Achieving A-PBE. If we have the luxury of prescribing our own password hashing function PH then we can provide a fast and simple A-PBE scheme, that we call APBE1, based on any PKE scheme. See [7]. However, this solution is invasive, asking for the deployment of a new PH, which may not be possible due to existing legacy passwords and password-hashing functions. We thus ask if it is possible to design a secure A-PBE scheme that is non-invasive. This means we take F.Ph as given and aim to achieve security by making reasonable assumptions about its security without prescribing its design, assumptions that in particular are met by the F.Ph function of PKCS#5 or other standards. This turns out to be more challenging. We now provide the APBE2 scheme that accomplishes this using WE.

Non-invasive A-PBE. We view ourselves as given a function family H with key, input and output length functions $\text{H.kl}, \text{H.il}, \text{H.ol}$. Our goal is to design an A-PBE scheme F such that F.Ph is H . In particular, we could let H be the password

$F[H, WE].Ph(1^\lambda, sa, pw)$	$F[H, WE].Enc(1^\lambda, hpw, sa, m)$	$F[H, WE].Dec(pw, c)$
$hpw \leftarrow H(1^\lambda, sa, pw)$	$x \leftarrow (1^\lambda, sa, hpw)$	$m \leftarrow WE.Dec(pw, c)$
Return hpw	$c \leftarrow_s WE(1^\lambda, x, m)$	Return m
	Return c	

Fig. 8. A-PBE scheme $F = APBE2[H, WE]$ associated to hash family H and witness encryption scheme WE for L_H

hashing function family from PKCS#5 [33] or bcrypt [34], thereby obtaining A-PBE without change in the existing hashed passwords. We begin by reviewing the security assumption on H .

Related-Input Pseudorandomness. Let H be a function family. This means that H is a deterministic, PT function taking 1^λ , a key $k \in \{0, 1\}^{H.kl(\lambda)}$ and an input $x \in \{0, 1\}^{H.il(\lambda)}$ to return $H(1^\lambda, k, x) \in \{0, 1\}^{H.ol(\lambda)}$. Here $H.kl, H.il, H.ol: \mathbb{N} \rightarrow \mathbb{N}$ are the key, input and output lengths associated to H , respectively. We say that H is related-input pseudorandom (RIP) if $Adv_{H,A}^{rip}(\cdot) = 2Pr[RIP_H^A(\cdot)] - 1$ is negligible for every PT unpredictable adversary $A = (A_1, A_2)$, where game RIP_H^A is shown in Fig. 7. Informally, this means that the hashed passwords should be indistinguishable from random strings, even in the presence of the salts. We note that this is exactly the property needed for classical S-PBE (symmetric PBE) to be secure, for it uses the hashed password as the symmetric key. Thus, the assumption can be viewed as already made and existing, even if implicitly, in current usage of passwords for S-PBE. We note that RIP security of H is implied by UCE security of H relative to statistically unpredictable sources [8].

The APBE2 Scheme. Let

$$L_H = \{ (1^\lambda, sa, H(1^\lambda, sa, pw)) : \lambda \in \mathbb{N}, sa \in \{0, 1\}^{H.kl(\lambda)}, pw \in \{0, 1\}^{H.il(\lambda)} \} .$$

This language is in NP. Let WE be a witness encryption scheme for L_H . We associate to H and WE the A-PBE scheme $F = APBE2[H, WE]$ specified in Fig. 8. We let $F.pl = H.il$, $F.sl = H.kl$ and $F.hl = H.ol$. The construction lets the salt play the role of the key for H , the password being the input and the hashed password the output.

Security of APBE2 under AS. Theorem 5 below says that if H is RIP and WE is $AS[L_H]$ -secure then $APBE2[H, WE]$ is a secure A-PBE scheme. The proof is in [7].

Theorem 5. Let H be a function family such that $2^{H.il(\cdot) - H.ol(\cdot)}$ is a negligible function. If H is RIP and $WE \in AS[L_H]$ then $F = APBE2[H, WE]$ is a secure A-PBE scheme.

The key feature of this result is that it is non-invasive, meaning it puts conditions on the hash family H that suffice for security rather than mandating any

GAME $\text{XS}_{\text{WE},R}^{A,E}(\lambda)$	GAME $\text{ROW}_H^A(\lambda)$
$(x, m_0, m_1, \text{St}) \leftarrow_{\$} A(1^\lambda); b \leftarrow_{\$} \{0, 1\}$	$\mathbf{pw} \leftarrow_{\$} A_1(1^\lambda)$
$c \leftarrow_{\$} \text{WE.Enc}(1^\lambda, x, m_b)$	For $i = 1$ to $ \mathbf{pw} $ do
$b' \leftarrow_{\$} A(\text{St}, c)$	$\mathbf{sa}[i] \leftarrow_{\$} \{0, 1\}^{\text{H.kl}(\lambda)}$
$w \leftarrow_{\$} E(1^\lambda, x, m_0, m_1, \text{St}, c)$	$\mathbf{hpw}[i] \leftarrow H(1^\lambda, \mathbf{sa}[i], \mathbf{pw}[i])$
Return $((b = b') \wedge \neg R(x, w))$	$(w, i) \leftarrow_{\$} A_2(1^\lambda, \mathbf{sa}, \mathbf{hpw})$
	Return $(\mathbf{hpw}[i] = H(1^\lambda, \mathbf{sa}[i], w))$

Fig. 9. Left: Game XS defining extractable security of witness encryption scheme WE. **Right:** Game ROW defining ROW security of H.

particular design of H. Practical and standardized key-derivation functions may be assumed to satisfy concrete versions of these asymptotic conditions.

Arbitrary Stretch. Define the stretch $\text{H.s}(\cdot) = \text{H.ol}(\cdot) - \text{H.il}(\cdot)$ of password hashing function H as the difference between its output length and its input length. Theorem 5 requires that $2^{-\text{H.s}(\cdot)}$ is negligible, meaning the output length of the hash must be somewhat longer than the input length. This captures situations in which passwords are, say 12-character ASCII strings (input length is 78-bit) and H is iterated SHA1 (output length is 160-bit). However, when passwords are longer, say 24-character, then Theorem 5 doesn't apply. This is unsatisfying, because intuitively, longer passwords should offer better security. In this section, we formalize a stronger security requirement for witness encryption called XS that allows us to remove the assumption on the stretch of H.

XS-Secure Witness Encryption. The security requirements for SS and AS are for $x \notin L$, no security requirement being made if $x \in L$. Extractable witness encryption [26] is a requirement for all $x \in \{0, 1\}^*$, asking that if the adversary violates privacy of encryption under x then one can extract a witness for the membership of $x \in L$. Intuitively, the only way to violate privacy is to know a witness. We provide a formalization of extraction security that we call XS. It strengthens the formalization of GKPVZ [26] in being adaptive, in the vein of AS, but weakens it by not involving auxiliary inputs. The formalizations also differ in other details.

Let R be an NP-relation and let $L = \mathcal{L}(R)$. Let WE be a witness encryption scheme for L. We say that WE is XS[L]-secure if for any PT adversary A there is a corresponding PT algorithm E such that $\text{Adv}_{\text{WE},R,A,E}^{\text{XS}}(\lambda) = 2 \Pr[\text{XS}_{\text{WE},R}^{A,E}(\lambda)] - 1$ is negligible, where game $\text{XS}_{\text{WE},R}^{A,E}$ is defined at the leftpanel of Fig. 9. Let XS[L] denote the set of correct, XS[L]-secure witness encryption schemes for L.

Intuitively, XS[L] security implies AS[L] security for any $L \in \text{NP}$, because in the former notion, if the adversary produces $x \notin L$ then no witness exists, so no extractor E (even a computationally unbounded one) can find one. Proposition 6 below formally confirms this. The proof is in [7].

Proposition 6. For any NP-relation R, it holds that $\text{XS}[\mathcal{L}(R)] \subseteq \text{AS}[\mathcal{L}(R)]$.

Extractable obfuscation (xO), also known as differing-input obfuscation, was defined in [1, 4, 14]. BCP [14] show that it implies extractable witness encryption meeting the definition of GKPVZ [26]. In [7], we give an alternative definition of xO and show that it implies $\text{XS}[\mathcal{L}(\mathbf{R})]$ -secure witness encryption, for any NP relation R. The construction is the same $\text{WE}_R[\mathbf{P}]$ in Section 3, where the obfuscator P is assumed to be xO-secure, instead of just being iO-secure.

Related-Input One-Wayness. We now formalize another hardness assumption, related-input one-wayness, on hash function family H. Informally we demand that if the adversary is given the hashed passwords and the salts, it can't compute a preimage of any hashed password. This is exactly the intuitive requirement for password-hashing functions: if passwords are well-chosen to resist dictionary attacks, then no adversary should be able to recover some password from the hashed ones. It's a variant of the notion of one-wayness under correlated products of [35]. Formally, we say that H is related-input one-way (ROW) if $\text{Adv}_{\mathbf{H},A}^{\text{row}}(\lambda) = \Pr[\text{ROW}_{\mathbf{H}}^A(\lambda)]$ is negligible for all PT unpredictable adversary $A = (A_1, A_2)$, where game $\text{ROW}_{\mathbf{H}}^A$ is shown at the right panel of Fig. 9.

Security of APBE2 under XS. The following establishes the security of $\mathbf{F} = \text{APBE2}[\mathbf{H}, \text{WE}]$ without any restrictions or assumptions on the stretch of H. See [7] for the proof.

Theorem 7. If H is ROW and $\text{WE} \in \text{XS}[L_{\mathbf{H}}]$ then $\mathbf{F} = \text{APBE2}[\mathbf{H}, \text{WE}]$ is a secure A-PBE scheme.

A Further Versions of SS

A good definition for WE security should have two properties: (1) *Usability*, meaning it should suffice to prove security of applications, and (2) *Achievability*, meaning it should be provably achieved by the natural constructs, which in this case means the iO-based one of GGHRSW [18]. Our AS definition has both properties. We have shown that SS [20, 21] lacks (1).

After seeing a prior version of our work, GGSW updated the ePrint version of their paper [22]. Here they acknowledge the gaps we find. They then propose their own modification of SS, that we call SS2, in an attempt to fill the gaps. This was unnecessary because AS had already been put forth and shown to fill the gap, but GGSW appeared to want a definition in their quantifier-based style rather than our game-based style. They viewed the problem in SS as arising from the “order of quantification” and attempted to address it by changing this order. SS2 quantified the negligible function first, making it universal. We explain below that SS2 is unachievable, meaning no WE scheme can be SS2 secure. (More precisely our result is that SS2-secure WE is unachievable for any NP-complete language unless the polynomial-time hierarchy collapses. Our proof uses the fact that statistically-secure WE is not achievable [20].) We pointed this out to GGSW in a personal communication. They acknowledged this and further updated their definition to one we call SS3 [22], which used another

order of quantification. Below we show that SS3 remains limited in terms of achievability. This is because it does not seem possible to show that the iO-based WE construction of GGHRSW [18] meets it under the definition of iO-security that is commonly used in other applications of iO [13,30,36] and that we have shown suffices for AS-secure WE.

The updated GGSW papers [22,23] characterize the gap we find as having to do with the “order of quantifiers” in the SS definition, and their fixes attempt to change quantifier order. However, the issue is not quantifier order but, more subtly, the relation between x and λ . More broadly, game-based definitions are a better fit in this domain than quantifier-based ones. This is because applications one wants to achieve with WE, as well as primitives one wants to use to achieve WE, are both themselves underlain by game-based definitions. Reductions are thus facilitated, and less error-prone, with a game-based WE definition. A quantifier-based one leads to mismatches. In particular, under certain quantifier orders, one gets definitions like SS that do not provide usability, and when one changes the order, one gets definitions like SS3 that are too strong and challenge achievability. Intuitively, the latter is because the quantification ends up demanding security even on inputs that no adversary could ever find. This does not mean a viable quantifier-based definition is impossible. Indeed, below, we suggest SS4, a quantifier-based definition of WE that recovers achievability under weak iO in the non-uniform case. But the game-based AS is simpler and more user friendly, and does not require non-uniformity to be achieved under weak iO.

Below we also consider a recent definition of soundness security of KNY [31]. We call it SS5. It is similar to SS2 and consequently also unachievable.

The above indicates that the problems we find with SS, and the fix we deliver with AS, are not trivial. Certainly it is easy, once the problem has been pointed out, to propose alternatives, but our work remains important in having pointed out the need for alternatives and in guiding the choice of, and verifying, these alternatives.

We believe that WE is an important and useful notion and that our work helps advance its cause via precise definitions that satisfy the usability and achievability conditions above. We believe it is important for our field that work like this is published, and that such work is not damaging for the GGSW authors but rather advances the primitive they proposed.

SS2. WE scheme WE is SS2[L]-secure according to [22] if there exists a negligible function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that for any PT adversary A , any $x \in \{0, 1\}^* \setminus L$, any equal-length $m_0, m_1 \in \{0, 1\}^*$, and any $\lambda \in \mathbb{N}$ we have

$$\begin{aligned} & \text{Adv}_{\text{WE}, L, x, m_0, m_1, A}^{\text{ss2}}(\lambda) \\ &= \Pr[A(\text{WE.Enc}(1^\lambda, x, m_1)) = 1] - \Pr[A(\text{WE.Enc}(1^\lambda, x, m_0)) = 1] < \nu(\lambda) . \end{aligned}$$

We claim this notion is unachievable, meaning, *no* WE scheme is SS2[L]-secure. The reason is that ν is universal and in particular not allowed to depend on the adversary. More formally let L be an NP-complete language. Let WE by any WE

scheme and let ν be any negligible function. We show that if the polynomial-time hierarchy does not collapse then there is a PT adversary A as well as $x \in \{0, 1\}^* \setminus L$, equal-length m_0, m_1 and $\lambda \in \mathbb{N}$ such that $\text{Adv}_{\text{WE}, L, x, m_0, m_1, A}^{\text{ss2}}(\lambda) \geq \nu(\lambda)$. This shows that WE is not SS2[L]-secure.

For probability distribution functions $\mu, \mu' : D \rightarrow [0, 1]$, let

$$\|\mu - \mu'\| = \frac{1}{2} \sum_{x \in D} |\mu(x) - \mu'(x)|$$

be the statistical distance between μ and μ' . For any λ, x, m let $\mu_{\lambda, x, m}$ be the distribution of $\text{WE.Enc}(1^\lambda, x, m)$. Results from GGSW [20] imply that, unless the polynomial hierarchy collapses, there exists a string $x' \in \{0, 1\}^* \setminus L$, equal-length messages m'_0, m'_1 and a constant $\lambda_0 \in \mathbb{N}$ such that $\|\mu_{\lambda_0, x', m'_1} - \mu_{\lambda_0, x', m'_0}\| \geq \nu(\lambda_0)$. Consider the following adversary A . On input a ciphertext c , if c is not in the domain of $\mu_{\lambda_0, x', m'_1}$ then A outputs a random guess. Otherwise, A outputs 1 if $\mu_{\lambda_0, x', m'_1}(c) > \mu_{\lambda_0, x', m'_0}(c)$, and outputs 0 otherwise. Note that the test as to whether c is in the domain of $\mu_{\lambda_0, x', m'_1}$ only takes polynomial time because λ_0, x', m'_1 are fixed, and all computations related to them are constant time, and similarly for computations of $\mu_{\lambda_0, x', m'_0}(\cdot)$. Thus A runs in polynomial time. But $\text{Adv}_{\text{WE}, L, x', m'_0, m'_1, A}^{\text{ss2}}(\lambda_0) = \|\mu_{\lambda_0, x', m'_1} - \mu_{\lambda_0, x', m'_0}\| \geq \nu(\lambda_0)$.

SS3. A WE scheme WE is SS3[L]-secure [23] if for any PT adversary A , there exists a negligible function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that for any $x \in \{0, 1\}^* \setminus L$ and any $\lambda \in \mathbb{N}$,

$$\begin{aligned} & \text{Adv}_{\text{WE}, L, x, A}^{\text{ss3}}(\lambda) \\ &= \Pr[A(\text{WE.Enc}(1^\lambda, x, 1)) = 1] - \Pr[A(\text{WE.Enc}(1^\lambda, x, 0)) = 1] < \nu(\lambda) . \end{aligned}$$

A first nit is that this considers only encryption of a 1-bit message but for applications one has to encrypt many bits, and it is not stated how security is defined in this case. More importantly, however, SS3 has limitations with regard to achievability. Specifically, it seems unlikely one can show that iO implies SS3[L]-secure WE via the natural GGHRWS construction that worked for both SS and AS and under the definition of iO that is used for other applications [13, 30, 36] and we have shown suffices for AS-secure WE. We now explain, referring to our formulation of the definition in Section 3. Let L be an NP language. In Section 3 we recalled the GGHRWS construction $\text{WE} = \text{WE}_R[\text{P}]$ of a WE scheme from an indistinguishability obfuscator P . Now assume we are given an arbitrary PT adversary A attacking the SS3[L]-security of WE. To prove security, we need to build an adversary B attacking the iO-security of P . Adversary B , given 1^λ , needs to efficiently find and output circuits of the form $R_{x, m}$ that we defined in Section 3, where x intuitively is an input where the WE security “breaks.” But how is B to find such an x efficiently? There seems to be no way. Even if we allow B to be non-uniform, its advice string has length polynomial in λ , and thus it can’t tell what is the “best” x because the set $\{0, 1\}^* \setminus L$ is infinite. In the case of AS, this was not a problem because A handed back an x on

which it succeeded. Also for the original SS, it is not a problem because the entire claim pertains to only one, fixed x that can be assumed known to B . An approach we might consider for SS3 is the following. Given any string $x \notin L$, we can build an adversary B_x such that $\text{Adv}_{\text{WE},L,x,A}^{\text{ss3}}(\lambda) \leq \text{Adv}_{\text{P},B_x}^{\text{iO}}(\lambda)$ for all $\lambda \in \mathbb{N}$. Now the assumed iO security gives us a negligible function ν_x such that $\text{Adv}_{\text{P},B_x}^{\text{iO}}(\cdot) < \nu_x(\cdot)$. But the SS3 notion wants a single negligible function ν that is independent of x . It's unclear how to get ν from the set $\{\nu_x : x \notin L\}$ since the latter set is infinite. One natural idea is to set $\nu(\lambda) = \sup_{x \notin L} \{\nu_x(\lambda)\}$. But this doesn't work. For example, consider $\nu_x(\lambda) = 1$ if $\lambda < |x|$, and $\nu_x(\lambda) = 0$ otherwise. For any fixed x , the function ν_x is negligible, but $\nu(\lambda) = 1$ for every $\lambda \in \mathbb{N}$, meaning ν is not negligible. So one appears to need to construct an iO adversary B independent of x , but it is unclear how to do that.

We note the proof does seem possible under some stronger notions of iO from [18]. However, iO is a strong assumption no matter what and it is desirable that applications use as weak a form of it as possible. Also in further work subsequent to ours, GLW [25] claim to achieve SS3 via a direct construction. However, this requires sub-exponential hardness assumptions. (They call it complexity leveraging.) Beyond this, trying to achieve SS3 is an unnecessary route to follow, since AS already provides the properties we want, namely it suffices for applications and is achieved even under weak iO.

SS4. As we have seen, the game-based AS definition fulfills the usability and achievability conditions for a good definition. However, GGSW appear to want a quantifier-based definition in the style of SS. But their SS2, SS3 attempts have been inadequate. Here we accordingly suggest a quantifier-based definition that we call SS4 which does satisfy, usability and is less limited than SS3 with regard to achievability, namely weak iO does suffice for it, as long as this is assumed for non-uniform adversaries. In particular it is implied by the non-uniform generalization of AS and thus can be built from weak, non-uniform iO by our results. We explain why it suffices for the PKE application of GGSW.

We say that a WE scheme WE is SS4[L]-secure if for any PT adversary A and any polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$, there exists a negligible function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that, for any string $x \in \{0, 1\}^* \setminus L$, and any $\lambda \in \mathbb{N}$, if $|x| \leq \ell(\lambda)$ then

$$\begin{aligned} & \text{Adv}_{\text{WE},L,\ell,x,A}^{\text{ss4}}(\lambda) \\ &= \Pr[A(1^\lambda, x, \text{WE.Enc}(1^\lambda, x, 1)) = 1] - \Pr[A(1^\lambda, x, \text{WE.Enc}(1^\lambda, x, 0)) = 1] \\ &< \nu(\lambda) . \end{aligned}$$

We now show this notion is implied by non-uniform AS. Given any SS4 adversary A and any polynomial ℓ , one can build another non-uniform AS adversary B as follows. For each $\lambda \in \mathbb{N}$, let $x_\lambda \in \{0, 1\}^* \setminus L$ be a string such that $|x_\lambda| \leq \ell(\lambda)$ and $\text{Adv}_{\text{WE},L,\ell,x,A}^{\text{ss4}}(\lambda) \leq \text{Adv}_{\text{WE},L,\ell,x_\lambda,A}^{\text{ss4}}(\lambda)$ for all $x \in \{0, 1\}^* \setminus L$ with $|x| \leq \ell(\lambda)$. Adversary $B(1^\lambda)$ outputs $(x_\lambda, 0, 1, \varepsilon)$, and $B(\text{St}, c)$ runs $A(1^\lambda, x_\lambda, c)$. Then for any string $x \in \{0, 1\}^* \setminus L$, $\lambda \in \mathbb{N}$, if $|x| \leq \ell(\lambda)$ then

$$\text{Adv}_{\text{WE},B}^{\text{as}}(\lambda) = \text{Adv}_{\text{WE},L,\ell,x_\lambda,A}^{\text{ss4}}(\lambda) \geq \text{Adv}_{\text{WE},L,\ell,x,A}^{\text{ss4}}(\lambda) .$$

We now briefly explain why SS4 is enough for GGSW’s PKE scheme, but under message space $\{0, 1\}$, because SS4 only allows encrypting a single bit. The IND-CPA adversary is assumed to make only a single query $(0, 1)$. The proof will follow the template in [7] but with a change in constructing WE adversary D from an INDCPA adversary A . Let $\ell(\lambda) = 2\lambda$ for every $\lambda \in \mathbb{N}$. Adversary $D(1^\lambda, x, c)$ runs $A(1^\lambda, pk)$ with $pk = (\lambda, x)$. When the latter makes its query, the former returns c . Finally, D outputs the same guess as A .

For both SS3 and SS4, in the PKE application, to encrypt an n -bit message, one has to make n calls to WE to encrypt n bits individually, exacerbating the inefficiency of the scheme. If one modifies SS3 and SS4 for encrypting equal-length m_0, m_1 of arbitrary length instead of $m_0 = 0$ and $m_1 = 1$, then the PKE still can only encrypt bit-by-bit. The reason is that, an INDCPA adversary A is allowed to choose any equal-length m_0, m_1 but in SS3 and SS4, the WE adversary D has no control of the messages m_0, m_1 , and thus one can’t construct D from A . This again shows that AS is superior to SS3 and SS4 in terms of usability.

SS5. In a recent paper, KNY [31] define the following variant of SS, which we call SS5. A scheme is SS5[L]-secure if for any security parameter λ , any equal-length messages $m_0, m_1 \in \{0, 1\}^{\text{poly}(\lambda)}$, any PT adversary A , and any $x \notin L$, we have

$$\begin{aligned} & \text{Adv}_{\text{WE}, L, x, m_0, m_1, A}^{\text{ss5}}(\lambda) \\ &= \Pr[A(\text{WE.Enc}(1^\lambda, x, m_1)) = 1] - \Pr[A(\text{WE.Enc}(1^\lambda, x, m_0)) = 1] < \text{negl}(\lambda) . \end{aligned}$$

This definition doesn’t specify where to place the (existential) quantifier for the negligible function negl , but the only meaningful position in the context of what is written is to place it prior to the (universal) quantification of the security parameter. (We certainly don’t want a different negligible function for every value of λ .) But if so, the function negl is independent of the adversary A . The same argument against SS2 can be used to show that SS5 is unachievable.

SS5 again demonstrates that quantifier-based notions for WE are error-prone. KNY’s definition [31] is problematic, although it is subsequent to our work and all of GGSW’s revisions.

Acknowledgments. We thank Krzysztof Pietrzak and Georg Fuchsbauer for discussions about witness encryption. Bellare is supported in part by NSF grants CNS-0904380, CCF-0915675, CNS-1116800 and CNS-1228890. Part of the work was done when Hoang was working at UCSD, and supported in part by NSF grants CNS-0904380, CCF-0915675, CNS-1116800 and CNS-1228890.

References

1. Ananth, P., Boneh, D., Garg, S., Sahai, A., Zhandry, M.: Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689 (2013). <http://eprint.iacr.org/2013/689>
2. Ananth, P.V., Gupta, D., Ishai, Y., Sahai, A.: Optimizing obfuscation: avoiding barrington’s theorem. In: Ahn, G.-J., Yung, M., Li, N. (eds.) ACM CCS 2014, pp. 646–658. ACM Press, November 2014

3. Barak, B., Garg, S., Kalai, Y.T., Paneth, O., Sahai, A.: Protecting obfuscation against algebraic attacks. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 221–238. Springer, Heidelberg (2014)
4. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (2001)
5. Bellare, M., Canetti, R., Krawczyk, H.: Keying hash functions for message authentication. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 1–15. Springer, Heidelberg (1996)
6. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public-key encryption schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998)
7. Bellare, M., Hoang, V.T.: Adaptive witness encryption and asymmetric password-based cryptography. Cryptology ePrint Archive, Report 2013/704 (2013). <http://eprint.iacr.org/2013/704>
8. Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating random oracles via UCEs. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 398–415. Springer, Heidelberg (2013)
9. Bellare, M., Meiklejohn, S., Thomson, S.: Key-versatile signatures and applications: RKA, KDM and joint Enc/Sig. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 496–513. Springer, Heidelberg (2014)
10. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance security and its application to password-based cryptography. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 312–329. Springer, Heidelberg (2012)
11. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)
12. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing* **13**(4), 850–864 (1984)
13. Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 480–499. Springer, Heidelberg (2014)
14. Boyle, E., Chung, K.-M., Pass, R.: On extractability obfuscation. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 52–73. Springer, Heidelberg (2014)
15. Boyle, E., Pass, R.: Limits of extractability assumptions with distributional auxiliary input. Cryptology ePrint Archive, Report 2013/703 (2013). <http://eprint.iacr.org/2013/703>
16. Brakerski, Z., Rothblum, G.N.: Virtual black-box obfuscation for all circuits via generic graded encoding. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 1–25. Springer, Heidelberg (2014)
17. Chase, M., Lysyanskaya, A.: On signatures of knowledge. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 78–96. Springer, Heidelberg (2006)
18. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS, pp. 40–49. IEEE Computer Society Press, October 2013
19. Garg, S., Gentry, C., Halevi, S., Wichs, D.: On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 518–535. Springer, Heidelberg (2014)

20. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC, pp. 467–476. ACM Press, June 2013
21. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. Cryptology ePrint Archive, Report 2013/258, version 20130508:202916, May 8, 2013
22. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. Cryptology ePrint Archive, Report 2013/258, version 20140211:224937, February 11, 2014
23. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. Cryptology ePrint Archive, Report 2013/258, version 20140418:025904, April 18, 2014
24. Gentry, C., Lewko, A., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption. Cryptology ePrint Archive, Report 2014/309 (2014). <http://eprint.iacr.org/2014/309>
25. Gentry, C., Lewko, A., Waters, B.: Witness encryption from instance independent assumptions. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 426–443. Springer, Heidelberg (2014)
26. Goldwasser, S., Kalai, Y.T., Popa, R.A., Vaikuntanathan, V., Zeldovich, N.: How to run turing machines on encrypted data. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 536–553. Springer, Heidelberg (2013)
27. Goldwasser, S., Micali, S.: Probabilistic encryption. *Journal of Computer and System Sciences* **28**(2), 270–299 (1984)
28. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. *SIAM Journal on Computing* **18**(1), 186–208 (1989)
29. Goyal, V., O’Neill, A., Rao, V.: Correlated-input secure hash functions. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 182–200. Springer, Heidelberg (2011)
30. Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: full domain hash from indistinguishability obfuscation. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 201–220. Springer, Heidelberg (2014)
31. Komargodski, I., Naor, M., Yorgev, E.: Secret-sharing for NP. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 254–273. Springer, Heidelberg (2014)
32. Nisan, N., Zuckerman, D.: Randomness is linear in space. *Journal of Computer and System Sciences* **52**(1), 43–52 (1996)
33. PKCS #5: Password-based cryptography standard (RFC 2898). RSA Data Security Inc, Version 2.0, September 2000
34. Provos, N., Mazières, D.: A future-adaptable password scheme. In: USENIX Annual Technical Conference, FREENIX Track, pp. 81–91 (1999)
35. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009)
36. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th ACM STOC, pp. 475–484. ACM Press, May/June 2014
37. Yao, A.C.-C.: Theory and applications of trapdoor functions (extended abstract). In: 23rd FOCS, pp. 80–91. IEEE Computer Society Press, November 1982