Black-Box Impossibilities of Obtaining 2-Round Weak ZK and Strong WI from Polynomial Hardness

Susumu Kiyoshima

NTT Research, Sunnyvale, CA, USA.
susumu.kiyoshima@ntt-research.com

Abstract. We study the problem of obtaining 2-round interactive arguments for NP with *weak zero-knowledge (weak ZK)* [Dwork et al., 2003] or with *strong witness indistinguishability (strong WI)* [Goldreich, 2001] under polynomially hard falsifiable assumptions. We consider both the *delayed-input* setting [Jain et al., 2017] and the standard non-delayed-input setting, where in the delayed-input setting, (i) prover privacy is only required to hold against *delayed-input verifiers* (which learn statements in the last round of the protocol) and (ii) soundness is required to hold even against *adaptive provers* (which choose statements in the last round of the protocol).

Concretely, we show the following black-box (BB) impossibility results by relying on standard cryptographic primitives.

1. It is impossible to obtain 2-round delayed-input weak ZK arguments under polynomially hard falsifiable assumptions if BB reductions are used to prove soundness. This result holds even when non-black-box techniques are used to prove weak ZK.
2. It is impossible to obtain 2-round non-delayed-input strong WI arguments and 2-round publicly verifiable delayed-input strong WI arguments under polynomially hard falsifiable assumptions if a natural type of BB reductions, called "oblivious" BB reductions, are used to prove strong WI.
3. It is impossible to obtain 2-round delayed-input strong WI arguments under polynomially hard falsifiable assumptions if BB reductions are used to prove both soundness and strong WI (the BB reductions for strong WI are required to be oblivious as above). Compared with the above result, this result no longer requires public verifiability in the delayed-input setting.

1 Introduction

Zero-knowledge (ZK) proofs and arguments have been extensively used in cryptography due to their powerful security. Informally, their security guarantees that an honest prover can convince a verifier of the validity of a statement without revealing anything beyond it. More formally, the *zero-knowledgeness* (ZK) guarantees that for any verifier there exists a (efficient) *simulator* such that for any distinguisher, the output of the simulator (which is given a statement and is executed alone) is indistinguishable from the output of the verifier (which interacts with an honest prover that proves the validity of the statement).

The powerful security of ZK protocols however comes with a cost: it is known that ZK protocols require at least three rounds for any language outside of BPP [18]. This lower bound limits the applicability of ZK protocols since many applications require that the number of interactions is at most two rounds.

Fortunately, it has been shown that by weakening the definition of ZK, we can obtain useful security notions that can be achieved in less than three rounds.[1] Such security notions include *witness indistinguishability (WI)* [14, 11], *witness hiding (WH)* [14, 5], *strong WI* [17, 25], *weak ZK* [12, 5], *super-polynomial-time simulation (SPS) ZK* [31], and *ZK against bounded-size verifiers* [4].

Still, the state-of-the-art is not satisfactory since many of the existing 2-round constructions for them are based on super-polynomially hard assumptions (i.e., assumptions against adversaries that run in fixed super-polynomial time) [25, 5, 31, 27, 2, 4, 26, 1, 20, 28]. Indeed, for some of the above security notions (such as strong WI and weak ZK as explained below), no 2-round construction is currently known under polynomially hard standard assumptions. This situation is frustrating since for WI, it has long been known that 2-round (or even non-interactive) constructions can be obtained from polynomially hard standard assumptions [11, 22].

In this work, we study whether the use of super-polynomially hard assumptions is unavoidable in these existing 2-round protocols, focusing on the cases of weak ZK and strong WI.

Weak ZK. Weak ZK is defined identically with ZK except that the order of the quantifier is reversed, i.e., it is now required that for any verifier V^* and any distinguisher D, there exists a simulator S (which may depend on both V^* and D) such that the distinguisher D cannot distinguish the output of the simulator S from the output of the verifier V^*. Weak ZK is weaker than ZK but still implies WI and WH.

Currently, two positive results are known about 2-round weak ZK, where one is shown in the *delayed-input setting* [25]—i.e., in the setting where (i) an honest verifier can create its first-round message without knowing the statement to be proven, (ii) soundness is required to hold even against any *adaptive prover*, which can choose the statement to prove in the last round of the protocol (i.e., after seeing the verifier's first-round message), and (iii) weak ZK is only required to hold against any *delayed-input verifier*, which creates its first-round message without knowing the statement to be proven. Note that the delayed-input setting and the standard (non-delayed-input) setting are incomparable since the former considers soundness against stronger provers whereas the latter considers weak ZK against stronger verifiers.

In the delayed-input setting, Jain et al. [25] constructed a 2-round argument that satisfies *distributional ϵ-weak ZK* for any inverse polynomial ϵ, where

[1] Throughout this paper, we focus on interactive proofs/arguments for all NP.

distributional ϵ-weak ZK is weaker than the standard weak ZK in that (i) the simulator is only required to work for random statements that are sampled from a distribution $\mathcal{D}$ and (ii) the distinguishing gap between the verifier's output and the simulator's output is only bounded by the inverse polynomial ϵ (the simulator is allowed to depend on both $\mathcal{D}$ and ϵ). The security of their protocol is proven under a quasi-polynomially hard assumption.

In the standard setting, Bitansky et al. [5] constructed a 2-round argument that is ϵ-weak ZK for any inverse polynomial ϵ under super-polynomially hard assumptions.[2]

Strong WI. Strong WI guarantees that for any two indistinguishable distributions $\mathcal{D}^0, \mathcal{D}^1$ over statements, no verifier can distinguish a proof for a random statement $x \leftarrow \mathcal{D}^0$ from a proof for a random statement $x \leftarrow \mathcal{D}^1$. A typical application of strong WI is proof of honest behaviors: for example, when a strong WI protocol is used to prove that a commitment is correctly generated, it directly guarantees that the hiding property of the commitment is preserved. (In contrast, the standard WI does not guarantee anything when the commitment is perfectly binding.)

In the delayed-input setting, Jain et al. [25] constructed a 2-round strong WI argument under a quasi-polynomially hard assumption. In the standard setting, the above-mentioned result about 2-round weak ZK [5] also holds for 2-round strong WI since ϵ-weak ZK implies strong WI.

1.1 Our Results

At a high level, we show impossibility results about obtaining 2-round weak ZK and strong WI protocols under "standard assumptions" by using "standard techniques." Following previous works (e.g., [16]), we formalize "standard assumptions" and "standard techniques" by using *falsifiable assumptions* and *black-box (BB) reductions*, respectively. Roughly speaking, (polynomially hard) falsifiable assumptions are the assumptions that are modeled as interactive games between a polynomial-time adversary and a polynomial-time challenger, where a falsifiable assumption (C, c) is considered true if no polynomial-time adversary can make the challenger C output 1 with probability non-negligibly higher than the threshold $c \in [0, 1]$. Essentially all standard cryptographic assumptions are falsifiable, including both general assumptions (e.g., the existence of one-way functions) and concrete ones (e.g., the RSA, DDH, and LWE assumptions). Regarding BB reductions, we consider two types of BB reductions, one is for soundness and the other is for strong WI. These two types are explained below with our results.

BB impossibility of 2-round weak ZK. Our first impossibility result is about obtaining 2-round weak ZK protocols while using BB reductions in the

[2] Weak ZK is defined slightly differently in Bitansky et al. [5], where essentially ϵ-weak ZK (for any inverse polynomial ϵ) is referred as "weak ZK." We follow other prior works [12, 8, 25] and require the distinguishing gap to be negligible.

proof of soundness. Here, BB reductions are defined for soundness as follows: for a 2-round weak ZK argument (P, V), we say that *the soundness of* (P, V) *is proven by a BB reduction based on a falsifiable assumption* (C, c) if there exists a polynomial-time oracle machine (or *BB reduction*) R such that for any verifier V^* that breaks the soundness of (P, V), the machine R^{V^*} breaks the assumption (C, c).

Theorem (informal). *Assume the existence of one-way functions. Then, there exists an NP language* $\boldsymbol{L}$ *such that if (i) there exists a 2-round delayed-input distributional* ϵ*-weak ZK argument for* $\boldsymbol{L}$ *and (ii) its adaptive soundness is proven by a BB reduction based on a falsifiable assumption* (C, c)*, then the assumption* (C, c) *is false.*

(The formal statement is given as Theorem 1 in Section 7.) We note that using BB reductions in the proof of soundness is quite common, and in particular, BB reductions are used in the proof of soundness in the above-mentioned two positive results of 2-round weak ZK [25, 5].[3] (In fact, to the best of our knowledge, currently there do not exist any non-BB technique that can be used to prove the soundness of 2-round interactive arguments.) This result therefore matches with the positive result of [25] (note that this result holds even for the distributional ϵ-weak ZK version of weak ZK) and thus explains why the use of super-polynomial-time hardness is required in [25]. Finally, we note that this result holds even when non-BB techniques are used in the proof of weak ZK.

Let us explain informally what this result says about the difficulty of obtaining 2-round weak ZK protocols under polynomially hard assumptions. First, in the delayed-input setting, this result directly explains the difficulty: to overcome this result, we need to prove the soundness of 2-round arguments by using non-BB techniques,[4] but given the state-of-the-art, this approach seems to require novel techniques. Second, even in the standard setting, this result partially explains the difficulty: to overcome this result, we need to consider protocols that are inherently not adaptively sound, and thus, we need to be careful when using the popular *FLS paradigm* [13]. Indeed, if we naively use the FLS paradigm (where the verifier sets up a "trapdoor statement" in the first round and the prover gives a WI proof in the second round to prove that either the actual statement is true or the trapdoor statement is true), it is often the case that the first-round message is independent of the statement and as a result adaptive soundness holds whenever soundness holds.

BB impossibility of 2-round strong WI (non-delayed-input or publicly verifiable). Our second impossibility result is about obtaining 2-round strong WI protocols while using a certain type of BB reductions in the proof of strong WI. Specifically, we consider BB reductions that we call *oblivious BB reductions*,

[3] In [5], weak ZK is proven by a non-black-box technique, but soundness is proven by a BB reduction.

[4] It is easy to verify that for interactive proofs (rather than arguments) in the delayed-input setting, the classical impossibility result of 2-round ZK [18] can be extended to 2-round weak ZK.

which are defined roughly as follows: for a 2-round strong WI protocol (P, V), we say that *the strong WI of* (P, V) *is proven by an oblivious BB reduction based on a falsifiable assumption* (C, c) if there exists a polynomial-time oracle machine (or oblivious BB reduction) R such that for any verifier V^* that breaks the strong WI of (P, V) w.r.t. some distributions $\mathcal{D}^0, \mathcal{D}^1$, the machine $R^{V^*, \mathcal{D}^0, \mathcal{D}^1}$ either breaks the assumption (C, c) or distinguishes the distributions $\mathcal{D}^0$ and $\mathcal{D}^1$. We note that R is oblivious to the distributions $\mathcal{D}^0, \mathcal{D}^1$ in the sense that R is defined before the distributions $\mathcal{D}^0, \mathcal{D}^1$ are specified.[5] (We emphasize that R is given oracle access to $\mathcal{D}^0, \mathcal{D}^1$.)

Theorem (informal). *Assume the existence of CCA-secure PKE. Then, there exists an NP language* $\boldsymbol{L}$ *such that the following hold.*

1. *If there exists a 2-round (non-delayed-input) strong WI protocol for* $\boldsymbol{L}$ *and its strong WI is proven by an oblivious BB reduction based on a falsifiable assumption* (C, c)*, then the assumption* (C, c) *is false.*
2. *If there exists a 2-round publicly verifiable delayed-input strong WI protocol*[6] *for* $\boldsymbol{L}$ *and its strong WI is proven by an oblivious BB reduction based on a falsifiable assumption* (C, c)*, then the assumption* (C, c) *is false.*

(The formal statement is given as Theorem 3 and Theorem 4 in Section 7.) We note that obliviousness is a natural property for BB reductions, and for example oblivious reductions are used in the above-mentioned positive result of 2-round strong WI [25] and in the trivial proof showing that ZK implies strong WI [17, Proposition 4.6.3]. (Indeed, we are not aware of any non-oblivious reduction that can be used to prove strong WI for NP w.r.t. all distributions.) We also note that the second part of this result in particular holds for strong WI versions of *ZAPs* [11] and *ZAP arguments* [1, 20, 28].

Let us explain informally what this result says about the difficulty of obtaining 2-round strong WI protocols under polynomially hard assumptions. In particular, since the only way to overcome this result is to use non-BB or non-oblivious techniques in the proof of strong WI (as long as we consider non-delayed-input or publicly verifiable protocols), let us explain informally the difficulty of using these two types of techniques.

- Let us first see the difficulty of using non-BB techniques. We first note that for witness hiding, there exists a non-BB technique [5] such that (i) it can be used to prove the prover privacy of 2-round arguments under polynomially hard assumptions and (ii) we can use it while proving soundness under polynomially hard assumptions (such as the existence of *witness encryption schemes* [15]). Unfortunately, the usage of this technique in the witness hiding setting strongly relies on a certain property of witness hiding (concretely,

[5] This type of obliviousness is considered previously for witness hiding [23].

[6] that is, a 2-round delayed-input strong WI protocol such that anyone can decide whether a proof is accepting or not given the protocol transcript (without knowing the verifier randomness).

the property that a successful cheating verifier against witness hiding outputs a witness for the statement). As a result, it is currently unclear whether we can use this (or any other) non-BB technique in the setting of strong WI while proving soundness under polynomially hard assumptions.

- Let us next see the difficulty of using non-oblivious techniques. The main difficulty is that when we consider strong WI that holds for all NP w.r.t. all distributions over statements, we currently do not have any technique that makes non-oblivious use of distributions. As a result, it is currently unclear whether any non-oblivious technique is useful to obtain 2-round strong WI under polynomially hard assumptions.

BB impossibility of 2-round strong WI (delayed-input). Our third impossibility result is about obtaining 2-round strong WI arguments while using BB reductions in the proofs of soundness and strong WI. The motivation behind this result is to give an impossibility result about 2-round privately verifiable delayed-input strong WI protocols (for which the above result does not hold).

Theorem (informal). *Assume the existence of trapdoor permutations. Then, there exists an NP language* $\boldsymbol{L}$ *such that if (i) there exists a 2-round delayed-input strong WI argument for* $\boldsymbol{L}$*, (ii) its soundness is proven by a BB reduction based on a falsifiable assumption* (C, c)*, and (iii) its strong WI is proven by an oblivious BB reduction based on a falsifiable assumption* (C', c')*, then either the assumption* (C, c) *or the assumption* (C', c') *is false.*

(The formal statement is given as Theorem 2 in Section 7.) We note that this result matches with the positive result of [25] since BB reductions are used for both soundness and strong WI in the result of [25] (the one for strong WI is oblivious). Thus, this result explains why the use of super-polynomial-time hardness is required in [25].

Let us explain informally what this result says about the difficulty of obtaining 2-round strong WI protocols under polynomially hard assumptions. Compared with the above result, this result holds even for 2-round privately verifiable delayed-input strong WI protocols, but it holds only when BB reductions are used in the proof of soundness. Still, it seems reasonable to think that this result explains the difficulty of obtaining 2-round strong WI protocols almost as strongly as the above one since, as in the case of 2-round weak ZK, novel techniques are likely to be required to obtain 2-round strong WI protocols without using BB reductions in the proof of soundness.

Summary. In Table 1, we summarize the settings that we consider in our impossibility results (standard v.s. delayed-input) for each combination of the types of reductions (BB and non-BB reductions for soundness and weak ZK, and oblivious BB, non-oblivious BB, and non-BB reductions for strong WI). For example, "delayed-input" in the cell that corresponds to BB for soundness and BB for weak ZK indicates that one of our results (concretely, the first result) shows the impossibility of 2-round delayed-input weak ZK arguments when BB techniques are used for both soundness and weak ZK.

Table 1. Summary of the settings that we consider in our impossibility results.

		weak ZK		strong WI	
		BB	non-BB	obl. BB	non-obl. BB & non-BB
Sound	BB	delayed-input	delayed-input	standard delayed-input	
	non-BB			standard pub-verifiable delayed-input	

2 Our Techniques

2.1 BB Impossibility of 2-Round Delayed-Input Weak ZK

We first explain how we obtain our BB impossibility result about 2-round delayed-input weak ZK. This result is technically less involved and is used in a non-modular way in one of our BB impossibility results about strong WI.

At a very high level, we obtain our result about weak ZK by obtaining a BB impossibility result about (t, ϵ)-*zero-knowledge* [8], which is defined identically with the standard zero-knowledge except that (i) the definition is parameterized by a polynomial t and an inverse polynomial ϵ, (ii) the running time of the distinguisher is bounded by t, and (iii) the distinguishing gap is bounded by ϵ (the simulator is allowed to depend on both t and ϵ). Note that (t, ϵ)-ZK is defined with the same order of quantifier as the standard ZK (i.e., in the form "$\forall V^* \exists S \forall D \ldots$") and thus seems much stronger than weak ZK. Nonetheless, it is known that weak ZK implies (t, ϵ)-ZK for every polynomial t and inverse polynomial ϵ (with no modification to the protocol) [8]. Thus, to obtain a BB impossibility result on weak ZK, it suffices to obtain it on (t, ϵ)-ZK.

Before explaining how we obtain a BB impossibility result about (t, ϵ)-ZK, let us explain a subtle difference between (t, ϵ)-ZK and the standard ZK. Specifically, we note that in (t, ϵ)-ZK (in particular, the one that is defined in [8] and shown to be implied by weak ZK), the indistinguishability between a real proof and simulation is only guaranteed to hold against uniform distinguishers, i.e., distinguishers that take no auxiliary input other than the one that is given to the verifier and the simulator.

Somewhat surprisingly, this subtle difference causes difficulties when we try to obtain impossibility results about (t, ϵ)-ZK by using known techniques. Indeed, the classical impossibility result of 2-round ZK [18] does not hold for (t, ϵ)-ZK exactly due to this difference. Also, known techniques in BB impossibility literature, such as those that have been used for the BB impossibility of other 2-round interactive protocols [16, 7, 9], also require non-uniform indistinguishability and thus cannot be used for (t, ϵ)-ZK directly.

Roughly speaking, we overcome the difficulties as follows. First, we observe that weak ZK implies (t, ϵ)-ZK with non-uniform indistinguishability if we allow the simulator of (t, ϵ)-ZK to run in a "pre-processing" manner, i.e., in a manner that the simulator is computationally unbounded before receiving the statement. (More specifically, the simulator is separated into two parts, a *pre-processing simulator* and a *main simulator*, where the pre-processing simulator is computationally unbounded and creates short trapdoor information without knowing the statement, and the main simulator takes the statement along with the trapdoor information and simulates the verifier's output in polynomial time.) Second, we observe that the *meta-reduction* techniques, which have been used for the BB impossibility of other 2-round interactive protocols [16, 7, 9], can be used naturally to obtain a BB impossibility result about 2-round delayed-input pre-processing (t, ϵ)-ZK. More details are explained below.

Step 1. Showing that weak ZK implies pre-processing (t, ϵ)-ZK. We first note that, as already observed in [8], weak ZK implies (t, ϵ)-ZK with non-uniform indistinguishability if we allow the simulator of (t, ϵ)-ZK to be non-uniform, i.e., if we only require that for each auxiliary input z_V to the verifier there exists an auxiliary input z_S to the simulator such that on input z_S, the simulator works for any (non-uniform) distinguisher. Now, the problem is that it is in general not possible to compute a "good" z_S from z_V efficiently. Thus, we give the simulator unbounded computing power so that it can compute a good z_S from z_V by brute force. To make sure that the simulator can compute a good z_S before receiving the statement, we further weaken the definition of (t, ϵ)-ZK and consider the distributional version of it, where the simulator is only required to work for random statements that are sampled from a certain distribution. Since it is now sufficient for the simulator to find a good z_S for random statements, the simulator can find it before obtaining the actual statement.

Step 2. Showing BB impossibility of pre-processing (t, ϵ)-ZK. We obtain a BB impossibility result about 2-round delayed-input pre-processing (t, ϵ)-ZK by appropriately modifying a proof that is given in [7, 9] for the BB impossibility of 2-round *super-polynomial-simulation (SPS) ZK*, where the simulator is allowed to run in fixed super-polynomial time T.[7] To see how we modify the proof of [7, 9], consider for example a step in the proof where it is shown that the simulator creates an accepting proof for a false statement. In [7, 9], this property is shown by (i) first observing that the simulator creates an accepting proof for a true statement due to the indistinguishability of simulation (note that an honest prover does so with probability 1 by completeness) and then (2) observing that the simulator creates an accepting proof even for a false statement due to the indistinguishability between true and false statements (since the simulator runs in super-polynomial time T, it is assumed that true and false statements are indistinguishable in $\mathsf{poly}(T)$ time). Clearly, when the simulator is computationally unbounded, the second step of this argument fails since the simulator can distinguish true and false statements by brute force. Nevertheless, in the

[7] In SPS ZK, the simulator is usually computationally bounded by a fixed moderate super-polynomial (e.g., a quasi-polynomial) but it can use its super-polynomial-time computing power arbitrarily. In pre-processing (t, ϵ)-ZK, the simulator is computationally unbounded but it can use its super-polynomial-time computing power only before receiving the statement.

pre-processing model, we can still show the same property by relying on the non-uniform polynomial-time indistinguishability of true and false statements. To see this, observe that the non-uniform indistinguishability guarantees that no polynomial-time algorithm can distinguish true and false statements even when it is given any auxiliary input that is computed independently of the statement. This guarantee is clearly sufficient to show that when the main simulator in the pre-processing model creates an accepting proof for a true statement, it creates an accepting proof even for a false statement.

2.2 BB Impossibility of 2-Round Strong WI

We next explain how we obtain our BB impossibility results about 2-round strong WI.

Non-interactive strong WI. First, as a warm-up, we explain how we can obtain a BB impossibility result about non-interactive strong WI. In particular, we show that the strong WI of non-interactive arguments cannot be proven by oblivious BB reductions based on falsifiable assumptions.

At a high level, the proof proceeds as follows. Recall that an oblivious BB reduction R_{SWI} for strong WI has the following property: for any verifier V^* that breaks strong WI w.r.t. some distributions $\mathcal{D}^0$, $\mathcal{D}^1$ over statements (meaning that V^* can distinguish a proof π for statement $x \leftarrow \mathcal{D}^0$ and a proof π for statement $x \leftarrow \mathcal{D}^1$), the reduction $R_{\mathrm{SWI}}^{V^*}$ either breaks the underlying assumption (C, c) or distinguishes $\mathcal{D}^0$ and $\mathcal{D}^1$.[8] First, we observe that $R_{\mathrm{SWI}}^{V^*}$ breaks the assumption (C, c) rather than distinguishes $\mathcal{D}^0$ and $\mathcal{D}^1$. Assume for contradiction that $R_{\mathrm{SWI}}^{V^*}$ distinguishes $\mathcal{D}^0$ and $\mathcal{D}^1$, and assume without loss of generality that V^* aborts when it receives a proof that is not accepting. Now, intuitively, the assumption that $R_{\mathrm{SWI}}^{V^*}$ can distinguish $x \leftarrow \mathcal{D}^0$ and $x \leftarrow \mathcal{D}^1$ seems to imply that R_{SWI} sends x to V^* along with an accepting proof (since otherwise V^* seems useless); if so, we can use R_{SWI} to break soundness by arguing that even when x is a false statement, R_{SWI} still sends x to V^* along with an accepting proof. A problem is that R_{SWI} might distinguish $x \leftarrow \mathcal{D}^0$ and $x \leftarrow \mathcal{D}^1$ by sending a related statement x' to V^* without directly sending x. We solve this problem by designing a "non-malleable" language $\mathbf{L}$, which guarantees that R_{SWI} cannot distinguish $x \leftarrow \mathcal{D}^0$ and $x \leftarrow \mathcal{D}^1$ even when it sends a related statement x' to V^*. After showing $R_{\mathrm{SWI}}^{V^*}$ breaks the assumption (C, c), we conclude that the assumption (C, c) must be false by observing that we can design as V^* a specific cheating verifier that breaks strong WI w.r.t. $\mathcal{D}^0$, $\mathcal{D}^1$ efficiently.

More specifically, the proof proceeds as follows. Consider an NP language $\mathbf{L}$ that contains all the encryptions of 0 and 1 of a CCA-secure public-key encryption scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, i.e., $\mathbf{L} := \{(\mathsf{pk}, \mathsf{ct}) \mid \exists r \text{ s.t. } \mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, 0; r) \text{ or } \mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, 1; r)\}$. Also, for each public key pk of PKE and each $b \in \{0, 1\}$, consider the distribution $\mathcal{D}^b_{\mathsf{pk}}$ that outputs a random encryption of b

[8] Formally, R_{SWI} also has oracle access to $\mathcal{D}^0$ and $\mathcal{D}^1$, but we ignore it for simplicity in this overview.

under the public-key pk, i.e., $\mathcal{D}^b_{\mathsf{pk}} := \{(\mathsf{pk}, \mathsf{ct}) \mid \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, b)\}$. (We emphasize that $\mathcal{D}^b_{\mathsf{pk}}$ always outputs ct that is encrypted with the hardwired public key pk.) Assume that there exist a non-interactive argument (P, V) for $\mathbf{L}$ and an oblivious BB reduction $R_{\rm SWI}$ for showing the strong WI of (P, V) based on a falsifiable assumption (C, c). Note that this assumption implies that for any public key pk and any verifier V^* that breaks the strong WI of (P, V) w.r.t. $\mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}$, the reduction $R^{V^*}_{\rm SWI}$ either breaks the assumption (C, c) or distinguishes $\mathcal{D}^0_{\mathsf{pk}}$ and $\mathcal{D}^1_{\mathsf{pk}}$. Now, our goal is to show that the assumption (C, c) is false. Toward this goal, for each public-key–secret-key pair $(\mathsf{pk}, \mathsf{sk})$, we consider the following verifier $V^*_{\rm SWI} = V^*_{\rm SWI}[\mathsf{pk}, \mathsf{sk}]$ against the strong WI of (P, V).

- **Verifier $V^*_{\mathbf{swi}}$:** Given a statement $(\mathsf{pk}', \mathsf{ct})$ and a proof π from the prover, return the decryption result $b \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$ to the prover if $\mathsf{pk} = \mathsf{pk}'$ holds and π is an accepting proof for $(\mathsf{pk}', \mathsf{ct})$, and return $\bot$ otherwise.

Note that for any $(\mathsf{pk}, \mathsf{sk})$, the verifier $V^*_{\rm SWI}$ breaks the strong WI w.r.t. $\mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}$ due to the correctness of PKE. Thus, for any $(\mathsf{pk}, \mathsf{sk})$, the reduction $R^{V^*_{\rm SWI}}_{\rm SWI}$ either breaks the assumption (C, c) or distinguishes $\mathcal{D}^0_{\mathsf{pk}}$ and $\mathcal{D}^1_{\mathsf{pk}}$. Now, we observe that the assumption (C, c) is false unless we can use the reduction $R^{V^*_{\rm SWI}}_{\rm SWI}$ to break either the CCA security of PKE or the soundness of (P, V). Consider the following three cases for random pk.

- **Case 1. $R^{V^*_{\mathbf{swi}}}_{\mathbf{swi}}$ breaks the assumption (C, c).** In this case, it follows immediately that the assumption (C, c) is false since we can emulate $V^*_{\rm SWI}$ for $R_{\rm SWI}$ efficiently by using sk for random $(\mathsf{pk}, \mathsf{sk})$.
- **Case 2. $R^{V^*_{\mathbf{swi}}}_{\mathbf{swi}}(\mathsf{pk}, \mathsf{ct})$ distinguishes whether $(\mathsf{pk}, \mathsf{ct}) \leftarrow \mathcal{D}^0_{\mathsf{pk}}$ or $(\mathsf{pk}, \mathsf{ct}) \leftarrow \mathcal{D}^1_{\mathsf{pk}}$, and $R_{\mathbf{swi}}$ does not send $(\mathsf{pk}, \mathsf{ct})$ to $V^*_{\mathbf{swi}}$ along with an accepting proof π.** In this case, we can use $R^{V^*_{\rm SWI}}_{\rm SWI}$ to break the CCA security of PKE since we can efficiently emulate $V^*_{\rm SWI}$ for $R_{\rm SWI}$ in the CCA-security game (i.e., by using the decryption oracle).
- **Case 3. $R^{V^*_{\mathbf{swi}}}_{\mathbf{swi}}(\mathsf{pk}, \mathsf{ct})$ distinguishes whether $(\mathsf{pk}, \mathsf{ct}) \leftarrow \mathcal{D}^0_{\mathsf{pk}}$ or $(\mathsf{pk}, \mathsf{ct}) \leftarrow \mathcal{D}^1_{\mathsf{pk}}$, and $R_{\mathbf{swi}}$ sends $(\mathsf{pk}, \mathsf{ct})$ to $V^*_{\mathbf{swi}}$ along with an accepting proof π.** In this case, we can use $R_{\rm SWI}$ to break the soundness of (P, V). Indeed, the CCA security of PKE guarantees that even when ct is a false statement (e.g., an encryption of 2), $R_{\rm SWI}$ still sends ct to $V^*_{\rm SWI}$ along with an accepting proof. Thus, we can straightforwardly design an attacker against the soundness of (P, V) by efficiently emulating $V^*_{\rm SWI}$ for $R_{\rm SWI}$ by using sk for random $(\mathsf{pk}, \mathsf{sk})$.

Note that in the above, it is important that the reduction $R_{\rm SWI}$ is oblivious, i.e., is black-box about the distributions. This is because when we rely on the CCA security of PKE, we require that a single reduction works for every $(\mathsf{pk}, \mathsf{sk})$.

2-round strong WI: non-delayed-input or publicly verifiable. Next, we explain the main difficulty that arises when we consider 2-round protocols. In general, when we consider a BB reduction for the strong WI of 2-round interactive arguments, we need to think that the reduction can "rewind" the given verifier V^*, i.e., it can control the randomness of V^* so that it can force V^* to reuse the same verifier message in multiple queries. In this case, the above argument for non-interactive strong WI fails when we try to use the reduction R_{SWI} to break the soundness of (P, V). To see this, note that the soundness attacker first receives a verifier message from the external verifier and needs to forward it to the internally emulated R_{SWI} as an oracle response from V^*_{SWI}. Now, if the reduction R_{SWI} can force V^*_{SWI} to reuse this verifier message in multiple queries (possibly for different statements when we consider the delayed-input setting), we can no longer efficiently emulate V^*_{SWI} for R_{SWI} since we cannot decide whether the reduction R_{SWI} creates an accepting proof or not.

We can easily avoid this difficulty if we consider the standard (non-delayed-input) strong WI and (possibly delayed-input) publicly verifiable strong WI. First, in the case of publicly verifiable strong WI, it is easy to see that the above argument for non-interactive strong WI still works with no modification since we can still emulate V^*_{SWI} for R_{SWI} efficiently even when the same first message is reused. Second, in the case of the standard strong WI, we can effectively prevent the reuse of verifier messages since we can consider a verifier that obtains all the randomness by applying PRF on the statement at the beginning.

Thus, it remains to consider privately verifiable delayed-input strong WI.

2-round strong WI: (possibly privately verifiable) delayed-input. In this case, we cannot obtain a BB impossibility result that is as strong as the one for non-interactive strong WI since there exists a positive result [25] whose strong WI is proven by a BB reduction based on a falsifiable assumption.[9] We thus consider a weaker form of BB impossibility result by assuming that soundness is also proven by a BB reduction based on a falsifiable assumption.

Our high-level strategy is to show that strong WI implies (a weak form of) weak ZK and then reuse our BB impossibility result about weak ZK. Toward showing that strong WI implies weak ZK, let us fix any verifier V^*_{WZK} and distinguisher D_{WZK} against the weak ZK of (P, V), and consider the following strong WI verifier $V^*_{\text{SWI}} = V^*_{\text{SWI}}[\mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$ (which can be seen as a generalization of $V^*_{\text{SWI}}[\mathsf{pk}, \mathsf{sk}]$, which we consider in the non-interactive case above).

Verifier $V^*_{\textbf{swi}}$:

1. Invoke V^*_{WZK} and let it interact with the external prover. Let $(\mathsf{pk}', \mathsf{ct})$ denote the statement given from the prover and out_V denote the output of V^*_{WZK}.
2. If $\mathsf{pk} = \mathsf{pk}'$ holds and D_{WZK} is convinced by the external prover (i.e., D_{WZK} outputs 1 on $((\mathsf{pk}', \mathsf{ct}), \mathsf{out}_V)$), return the decryption result $b \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$ to the prover. Otherwise, return $\bot$.

Note that V^*_{SWI} returns a meaning response only when it receives a proof that convinces D_{WZK}. Now, at a high level, by arguing similarly to the case of non-interactive strong WI (with this new version of V^*_{SWI}), we show that the assumption (C, c) is false unless we can use the reduction $R_{\text{SWI}}^{V^*_{\text{SWI}}}$ either to break the CCA security of PKE or to obtain a weak ZK simulator that convinces D_{WZK}.

[9] The soundness is proven based on quasi-polynomially hard assumptions.

Unfortunately, although our strategy is intuitively simple, we need to overcome various problems because of subtle differences from the case of non-interactive strong WI (where we use R_{SWI} to break the soundness of (P, V) rather than to obtain a weak ZK simulator).

1. Unlike the case that we use the reduction R_{SWI} to break the soundness of (P, V) (where it suffices to construct a prover that obtains sk as auxiliary input to emulate V^*_{SWI} for R_{SWI} efficiently), we need to construct a weak ZK simulator that is not given sk and still is able to emulate V^*_{SWI} for R_{SWI}—this is because for our proof of weak ZK BB impossibility to go through, we need to make sure that the simulator cannot distinguish true statements (encryptions of 0 or 1) and false statements (encryptions of 2) so that we can show that the simulator creates an accepting proof for a false statement as mentioned at the end of Section 2.1. To overcome this problem, we assume that the CCA-secure encryption PKE in the definition of the language $\mathbf{L}$ is *puncturable* in the following sense: the CCA security holds even when the adversary is given a *punctured secret key* that can be used to emulate the decryption oracle unless the target ciphertext is queried. (It is easy to see that the classical CCA-secure encryption by Dolev et al. [10] satisfies such a property.) Then, we consider a simulator that takes as auxiliary input a punctured secrete key $\mathsf{sk}_{\{\mathsf{ct}\}}$ that corresponds to the statement $(\mathsf{pk}, \mathsf{ct})$ (i.e., $\mathsf{sk}_{\{\mathsf{ct}\}}$ is a key that can be used to emulate the decryption oracle unless ct is queried). The simulator can now emulate V^*_{SWI} for R_{SWI} efficiently by using $\mathsf{sk}_{\{\mathsf{ct}\}}$ and yet it cannot distinguish true and false statements as required.
2. Unlike the case that we use the reduction R_{SWI} to break the soundness of (P, V) (where it suffices to show that we can use R_{SWI} to create a convincing proof for a single (false) statement), we need to show that we can use R_{SWI} to create a convincing proof (w.r.t. V^*_{WZK} and D_{WZK}) for any (true) statement. This is in general hard to show since R_{SWI} might work only for a non-negligible fraction of the statements (this is because the reduction R_{SWI} is only guaranteed to have non-negligible advantage even when it is combined with a verifier V^* that breaks strong WI with very high advantage). To overcome this problem, we consider a weaker definition of distributional weak ZK where (i) the simulator is given polynomially many statements that are sampled from a distribution over $\mathbf{L}$ and (ii) the simulator is only required to give a simulated proof for one of these statements. Now, by properly defining the distribution, we can show that if the simulator is given sufficiently many statements, with high probability the simulator can find a statement for which the reduction R_{SWI} works, so it can create a convincing proof for one of the statements. Furthermore, our BB impossibility of weak ZK can be easily extended to this distributional weak ZK setting.
3. Unlike the case that we use the reduction R_{SWI} to break the soundness of (P, V) (where it suffices to show that R_{SWI} creates a proof that is convincing with non-negligible probability), we need to show that R_{SWI} creates a proof that is convincing with probability as high as an honest proof. To overcome this problem, we modify V^*_{SWI} in such a way that (i) V^*_{SWI} approximates (by

sampling) the probability that an honest prover convinces D_{WZK} for a random statement, and also approximates the probability that the external prover convinces D_{WZK}, and (ii) V^*_{SWI} returns the decryption result $b \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$ only when the latter is sufficiently high compared with the former. Now, we can show that R_{SWI} creates a proof that convinces D_{WZK} with probability as high as an honest proof since otherwise R_{SWI} cannot obtain meaningful responses from V^*_{SWI}.

3 Preliminaries

We denote the security parameter by n. For any random variable X, we use $\mathrm{Supp}(X)$ to denote the support of X. For any NP language $\mathbf{L}$, we use $\mathbf{R_L}$ to denote its witness relation. For any pair of (possibly probabilistic) interactive Turing machines (P, V), we use $\langle P(w), V(z)\rangle(x)$ for any $x, w, z \in \{0,1\}^*$ to denote the random variable representing the output of V in an interaction between $P(x, w)$ and $V(x, z)$. Specifically, since we only consider such P and V that participate in a 2-round interaction where V starts the interaction, $\langle P(w), V(z)\rangle(x)$ represents the value out_V that is generated in the following process: $m_1 \leftarrow V(x, z)$; $m_2 \leftarrow P(x, w, m_1)$; $\mathsf{out}_V \leftarrow V(m_2)$.[10]

Unless explicitly stated, we assume that cryptographic primitives are secure against non-uniform adversaries. Following the standard convention, we think that a Turing machine runs in polynomial time if its running time is polynomially bounded in the length of its first input (which is often implicitly the security parameter). For any two sequences of random variables (or distributions) $\mathcal{X} = \{X_i\}_{i\in\mathbb{N}}, \mathcal{Y} = \{Y_i\}_{i\in\mathbb{N}}$, we use $\mathcal{X} \approx \mathcal{Y}$ to denote that $\mathcal{X}$ and $\mathcal{Y}$ are computationally indistinguishable.

3.1 (δ, γ)-Approximation

Definition 1. *For any $p, \delta, \gamma \in [0,1]$, a probabilistic algorithm* Algo *is said to give* a (δ, γ)-approximation *of p if the output $\tilde{p}$ of* Algo *satisfies* $\Pr[|\tilde{p} - p| \leq \delta] \geq 1 - \gamma$.

It is easy to see (using a Chernoff Bound) that for any $\delta, \gamma \in [0,1]$ and any distribution $\mathcal{D}$ over $\{0,1\}$, a (δ, γ)-approximation of $p := \Pr[b = 1 \mid b \leftarrow \mathcal{D}]$ can be obtained by taking $k := \Theta(\delta^{-2} \log \gamma^{-1})$ samples from $\mathcal{D}$ and computing the relative frequency in which 1 is sampled.

3.2 2-Round Interactive Argument

Basic definitions. Let us recall the definitions of interactive arguments [19, 6] and their delayed-input version [25], focusing on 2-round ones.

[10] It should be understood that the secret state that is generated in the first invocation of V is implicitly inherited by the second invocation of V.

Definition 2 (Interactive argument). *For any NP language $\boldsymbol{L}$, a pair of interactive Turing machines (P, V) is called* a 2-round interactive argument for $\mathbf{L}$ *if it satisfies the following.*

- ***Completeness.*** *There exists a negligible function* negl *such that for every* $(x, w) \in \boldsymbol{R_L}$, $\Pr\left[\langle P(w), V\rangle(x) = 1\right] \geq 1 - \mathsf{negl}(|x|)$.
- ***Soundness.*** *For every* PPT *interactive Turing machine P^*, there exists a negligible function* negl *such that for every $x \in \{0,1\}^* \setminus \boldsymbol{L}$ and $z \in \{0,1\}^*$*, $\Pr\left[\langle P^*(z), V\rangle(x) = 1\right] \leq \mathsf{negl}(|x|)$.

Definition 3 (Delayed-input interactive argument). *A 2-round interactive argument (P, V) for an NP language $\boldsymbol{L}$ is called* delayed-input *if it satisfies the following.*

- ***Completeness.*** *There exists a negligible function* negl *such that for every* $(x, w) \in \boldsymbol{R_L}$,
$$\Pr\left[\mathsf{out} = 1 \;\middle|\; \begin{array}{l} m_1 \leftarrow V(1^{|x|});\ m_2 \leftarrow P(x, w, m_1) \\ \mathsf{out} \leftarrow V(x, m_2) \end{array}\right] \geq 1 - \mathsf{negl}(|x|) \ .$$
- ***Adaptive soundness.*** *For every* PPT *interactive Turing machine P^*, there exists a negligible function* negl *such that for every $n \in \mathbb{N}$ and $z \in \{0,1\}^*$,*
$$\Pr\left[\begin{array}{l} \mathsf{out} = 1 \\ \wedge\ x \in \{0,1\}^n \setminus \boldsymbol{L} \end{array} \;\middle|\; \begin{array}{l} m_1 \leftarrow V(1^n);\ (x, m_2) \leftarrow P^*(1^n, z, m_1) \\ \mathsf{out} \leftarrow V(x, m_2) \end{array}\right] \leq \mathsf{negl}(n) \ .$$

Notation. For a 2-round delayed-input interactive argument (P, V) for an NP language $\mathbf{L}$, an interactive Turing machine V^* is called a *delayed-input verifier* if for any $(x, w) \in \mathbf{R_L}$, it interacts with $P(x, w)$ in behalf of V in the manner defined in the definition of the correctness above (i.e., in the manner that V^* receives x in the last round of the interaction). For a delayed-input verifier V^*, the notation $\langle P(w), V^*(z)\rangle(x)$ is overloaded naturally, i.e., it denotes the value out_V that is generated in the following process: $m_1 \leftarrow V^*(1^{|x|}, z)$; $m_2 \leftarrow P(x, w, m_1)$; $\mathsf{out}_V \leftarrow V^*(x, m_2)$.

Strong witness indistinguishability. Next, let us recall the definition of strong witness indistinguishability (strong WI) [17]. Since we focus on negative results, we give a definition that is slightly weaker than the one given in [17, Definition 4.6.2].

Definition 4 ((delayed-input) strong WI). *An interactive argument (resp., a delayed-input interactive argument) (P, V) for an NP language $\boldsymbol{L}$ is called* strongly witness indistinguishable *(resp.,* delayed-input strongly witness indistinguishable*) if the following holds: for every $\{(\mathcal{X}_n^0, \mathcal{W}_n^0)\}_{n\in\mathbb{N}}, \{(\mathcal{X}_n^1, \mathcal{W}_n^1)\}_{n\in\mathbb{N}}$ and $\{z_n\}_{n\in\mathbb{N}}$ where each $(\mathcal{X}_n^b, \mathcal{W}_n^b)$ is a joint distribution that ranges over $\boldsymbol{R_L} \cap (\{0,1\}^n \times \{0,1\}^*)$ and each z_n is a string in $\{0,1\}^*$, if it holds $\{\mathcal{X}_n^0\}_{n\in\mathbb{N}} \approx \{\mathcal{X}_n^1\}_{n\in\mathbb{N}}$, then for every* PPT *verifier (resp.* PPT *delayed-input verifier) V^* there exists a negligible function* negl *such that for every $n \in \mathbb{N}$,*
$$\left|\begin{array}{l} \Pr\left[\langle P(w), V^*(z_n)\rangle(x) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n^0, \mathcal{W}_n^0)\right] \\ -\Pr\left[\langle P(w), V^*(z_n)\rangle(x) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n^1, \mathcal{W}_n^1)\right] \end{array}\right| \leq \mathsf{negl}(n) \ .$$

Delayed-input weak zero-knowledge. Next, let us recall the definition of weak zero-knowledge (weak ZK) [12, 8], focusing on the delayed-input version of it while considering non-uniform indistinguishability. Since we focus on negative results, we give a weaker, distributional (t, ϵ) version of the definition [8, 25].

Definition 5 (delayed-input distributional weak (t, ϵ)-zero-knowledge). *Let $\boldsymbol{L}$ be an NP language, t be a polynomial, and ϵ be an inverse polynomial. Then, a delayed-input interactive argument (P, V) for $\boldsymbol{L}$ is said to be* delayed-input distributional weak (t, ϵ)-zero-knowledge *if for every sequence of joint distributions $\mathcal{D}_{xw} = \{(\mathcal{X}_n, \mathcal{W}_n)\}_{n\in\mathbb{N}}$ such that each $(\mathcal{X}_n, \mathcal{W}_n)$ ranges over $\boldsymbol{R_L} \cap (\{0,1\}^n \times \{0,1\}^*)$, every* PPT *delayed-input verifier V^*, and every probabilistic t-time distinguisher D, there exists a* PPT *simulator S and an $n_0 \in \mathbb{N}$ such that for every $n > n_0$, $z_V \in \{0,1\}^*$, and $z_D \in \{0,1\}^*$, it holds*

$$\left| \begin{array}{l} \Pr\left[D(x, z_D, \langle P(w), V^*(z_V)\rangle(x)) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n, \mathcal{W}_n)\right] \\ - \Pr\left[D(x, z_D, S(x, z_V, z_D)) = 1 \mid x \leftarrow \mathcal{X}_n\right] \end{array} \right| \leq \epsilon(n) \enspace .$$

Special-purpose (weak) zero-knowledge. Next, let us introduce two new prover privacy notions for interactive arguments, where one is a weaker version of ZK and the other is a weaker version of weak ZK. We note that these nations should be viewed just as useful tools for our negative results; they are not intended to give any intuitively meaningful security.

First, we introduce special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional pre-processing (t, ϵ)-zero-knowledge. For editorial simplicity, we focus on deterministic verifiers below.

Definition 6 (special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional pre-processing (t, ϵ)-zero-knowledge). *Let $\boldsymbol{L}$ be an NP language, N, t be polynomials, ϵ be an inverse polynomial, and $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ be a sequence of joint distributions such that each $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ ranges over $(\boldsymbol{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$. Then, a 2-round delayed-input interactive argument (P, V) for $\boldsymbol{L}$ is said to be* special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional pre-processing (t, ϵ)-zero-knowledge *if for every deterministic polynomial-time delayed-input verifier V^*, there exists a simulator $S = (S_{\mathrm{pre}}, S_{\mathrm{main}})$ such that (i) S_{pre} is computationally unbounded and S_{main} is* PPT *and (ii) for every probabilistic t-time distinguisher D, there exists an $n_0 \in \mathbb{N}$ such that for every $n > n_0$, $z_V \in \{0,1\}^*$, and $z_D \in \{0,1\}^*$, it holds*

$$\left| \begin{array}{l} \Pr\left[D(x, z_D, \langle P(w), V^*(z_V)\rangle(x)) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n, \mathcal{W}_n)\right] \\ - \Pr\left[D(x_{i^*}, z_D, v) = 1 \;\middle|\; \begin{array}{l} \mathsf{st}_S \leftarrow S_{\mathrm{pre}}(1^n, z_V) \\ (x_i, z_{x,i}) \leftarrow (\mathcal{X}_n, \mathcal{Z}_n) \text{ for } \forall i \in [N_n] \\ (i^*, v) \leftarrow S_{\mathrm{main}}(\{x_i, z_{x,i}\}_{i\in[N_n]}, \mathsf{st}_S) \end{array} \right] \end{array} \right| \leq \epsilon(n) \enspace ,$$

where $N_n \coloneqq N(n, 1/\epsilon(n))$.

We note that although the simulator is given some extra information $z_{x,i}$ about each x_i in the above definition, we will only consider the setting where $z_{x,i}$ does not contain much information about a witness for x_i. In particular, the distribution $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ that we will consider has a related distribution $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$ over $(\{0,1\}^n \setminus L) \times \{0,1\}^*$ such that $(\mathcal{X}_n, \mathcal{Z}_n)$ and $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$ are computationally indistinguishable.

Next, we introduce special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge.

Definition 7 (special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge). *Let $\boldsymbol{L}$ be an NP language, N, t be polynomials, ϵ be an inverse polynomial, and $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n \in \mathbb{N}}$ be a sequence of joint distributions such that each $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ ranges over $(\boldsymbol{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$. Then, a 2-round delayed-input interactive argument (P, V) for $\boldsymbol{L}$ is said to be* special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge *if for every deterministic polynomial-time delayed-input verifier V^* and every probabilistic t-time distinguisher D, there exists a* PPT *simulator S and an $n_0 \in \mathbb{N}$ such that for every $n > n_0$, $z_V \in \{0,1\}^*$, and $z_D \in \{0,1\}^*$, it holds*

$$\Pr\left[D(x_{i^*}, z_D, v) = 1 \;\middle|\; \begin{array}{l} (x_i, z_{x,i}) \leftarrow (\mathcal{X}_n, \mathcal{Z}_n) \text{ for } \forall i \in [N_n] \\ (i^*, v) \leftarrow S(\{x_i, z_{x,i}\}_{i \in [N_n]}, z_V, z_D) \end{array} \right]$$
$$\geq \Pr\left[D(x, z_D, \langle P(w), V^*(z_V) \rangle(x)) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n, \mathcal{W}_n) \right] - \epsilon(n) \enspace ,$$

where $N_n := N(n, 1/\epsilon(n))$.

Remark 1 (Non-uniform indistinguishability). In both Definition 6 and Definition 7, the indistinguishability between a real proof and simulation holds against non-uniform distinguisher since the distinguisher takes its own auxiliary input z_D (which can contain z_V if necessary). Note that in Definition 7, the simulator also takes z_D since we consider the weak ZK setting.

3.3 Falsifiable Assumption and Black-Box Reduction

Falsifiable assumption. First, let us recall the definition of falsifiable assumptions from [29, 16].

Definition 8 (Falsifiable assumption). *A* falsifiable cryptographic assumption *consists of a* PPT *interactive Turing machine C and a constant $c \in [0, 1)$, where C is called the* challenger. *On security parameter n, the challenger $C(1^n)$ interacts with an interactive Turing machine $\mathcal{A}(1^n, z)$ for some $z \in \{0,1\}^*$ and C outputs a bit $b \in \{0,1\}$ at the end of the interaction; $\mathcal{A}$ is called the* adversary, *and when $b = 1$, it is said that $\mathcal{A}(1^n, z)$* wins *$C(1^n)$. The assumption associated with the tuple (C, c) states that for every* PPT *adversary $\mathcal{A}$ there exists a negligible function* negl *such that for every $n \in \mathbb{N}$ and $z \in \{0,1\}^*$, it holds* $\Pr\left[\langle \mathcal{A}(z), C \rangle(1^n) = 1\right] \leq c + \mathsf{negl}(n)$.

For any polynomial p and security parameter n, we say that an (possibly inefficient) adversary $\mathcal{A}$ *breaks* a falsifiable assumption (C, c) on n with *advantage*

$1/p(n)$ if there exists $z \in \{0,1\}^*$ such that it holds $\Pr\left[\langle \mathcal{A}(z), C\rangle(1^n) = 1\right] \geq c + 1/p(n)$. We say that an (possibly inefficient) adversary $\mathcal{A}$ breaks a falsifiable assumption (C, c) if there exists a polynomial p such that on infinitely many $n \in \mathbb{N}$, $\mathcal{A}$ breaks (C, c) with advantage $1/p(n)$.

Black-box reduction. Next, we introduce the definitions of black-box (BB) reductions. We consider BB reductions for adaptive soundness and BB reductions for strong WI. The former is defined as in [16, 33] and the latter is defined similarly to "oblivious" BB reductions for witness hiding [23].

Definition 9 (BB reduction for adaptive soundness). *Let (P, V) be a pair of interactive Turing machines that satisfies the correctness of a delayed-input 2-round interactive argument for an NP language $\boldsymbol{L}$. Then, a* PPT *oracle Turing machine R is said to be a* black-box reduction for showing the adaptive soundness of (P, V) based on a falsifiable assumption (C, c) *if there exists a polynomial p such that for every (possibly inefficient) interactive Turing machine P^* and every sufficiently large $n \in \mathbb{N}$, if there exists $z \in \{0,1\}^*$ such that*

$$\Pr\left[\begin{array}{l}\mathsf{out} = 1 \\ \wedge\ x \in \{0,1\}^n \setminus \boldsymbol{L}\end{array}\ \middle|\ \begin{array}{l} m_1 \leftarrow V(1^n);\ \ (x, m_2) \leftarrow P^*(1^n, z, m_1) \\ \mathsf{out} \leftarrow V(x, m_2)\end{array}\right] \geq \frac{1}{2}\ ,$$

then the machine $R^{P_z^*}$ breaks the assumption (C, c) on n with advantage $1/p(n)$ (where P_z^* is the same as P^* except that z is hardwired as its auxiliary input).

Definition 10 (Oblivious BB reduction for (delayed-input) strong WI). *Let (P, V) be a pair of interactive Turing machines that satisfies the correctness of 2-round interactive argument (resp., delayed-input interactive argument) for an NP language $\boldsymbol{L}$. Then, a* PPT *oracle Turing machine R is said to be an* oblivious black-box reduction for showing the strong WI (resp., delayed-input strong WI) of (P, V) based on a falsifiable assumption (C, c) *if for every polynomial p, there exists a polynomial p' such that for every (possibly inefficient) verifier (resp., delayed-input verifier) V^*, every sufficiently large $n \in \mathbb{N}$, every two joint distributions $\mathcal{D}_n^0 = (\mathcal{X}_n^0, \mathcal{W}_n^0), \mathcal{D}_n^1 = (\mathcal{X}_n^1, \mathcal{W}_n^1)$ such that each $(\mathcal{X}_n^b, \mathcal{W}_n^b)$ ranges over $\boldsymbol{R_L} \cap (\{0,1\}^n \times \{0,1\}^*)$, and every $z \in \{0,1\}^*$, if*

$$\Pr\left[\langle P(w), V^*(z)\rangle(x) = b\ \middle|\ \begin{array}{l} b \leftarrow \{0,1\} \\ (x, w) \leftarrow (\mathcal{X}_n^b, \mathcal{W}_n^b)\end{array}\right] \geq \frac{1}{2} + \frac{1}{p(n)}\ ,$$

then either (i) $R^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{p(n)})$ breaks the assumption (C, c) on n with advantage $1/p'(n)$ or (ii) $R^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{p(n)})$ distinguishes $\mathcal{X}_n^0$ and $\mathcal{X}_n^1$ with advantage $1/p'(n)$, i.e., it holds

$$\left|\begin{array}{l} \Pr\left[R^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{p(n)}, x) = 1\ \middle|\ x \leftarrow \mathcal{X}_n^0\right] \\ -\Pr\left[R^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{p(n)}, x) = 1\ \middle|\ x \leftarrow \mathcal{X}_n^1\right]\end{array}\right| \geq \frac{1}{p'(n)}\ ,$$

where V_z^* is the same as V^* except that z is hardwired as its auxiliary input.

Remark 2. As is [7, 33], we assume that given security parameter n, BB reductions make queries to the adversary with the same security parameter n. Also, we note that in Definition 9, the reduction R is given access to an adversary P^* that strongly breaks soundness (in the sense that the success probability is $1/2$ rather than non-negligible). Since we consider negative results (which essentially show the nonexistence of BB reductions), focusing on reductions that have access to such an adversary makes our results stronger.

Conventions. Note that in Definition 9 and Definition 10, BB reductions are given access to probabilistic interactive Turing machines. When an oracle machine R is given oracle access to a probabilistic interactive Turing machine $\mathcal{A}$, we follow the following conventions (see, e.g., [3, 17]), which are (to the best of our knowledge) general enough to capture the existing BB reductions.

- What R actually makes queries to is the next-message function of $\mathcal{A}$, i.e., a function $\mathcal{A}_r$ for some randomness r such that for any input x and a (possibly empty) list of messages $\boldsymbol{m}$, $\mathcal{A}_r(x, \boldsymbol{m})$ returns the message that $\mathcal{A}(x; r)$ will send after receiving messages $\boldsymbol{m}$ (or it returns the output of $\mathcal{A}$ if the interaction reaches the last round after $\mathcal{A}(x; r)$ receives $\boldsymbol{m}$).
- The randomness for $\mathcal{A}$ is set uniformly randomly, and in each query R can choose whether $\mathcal{A}$ should reuse the current randomness or it should use new (uniformly random) randomness.

3.4 Puncturable (CCA-Secure) Public-Key Encryption

Let us first recall the definition of CCA-secure public-key encryption [30, 32].

Definition 11. *A* CCA-secure public-key encryption scheme *(PKE) consists of three* PPT *algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *that satisfy the following.*

- ***Correctness.*** *For every* $n \in \mathbb{N}$ *and* $m \in \{0,1\}^n$, $\Pr\left[\mathsf{Dec}(\mathsf{sk}, c) = m \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n);\ c \leftarrow \mathsf{Enc}(\mathsf{pk}, m)\right] = 1$.
- ***CCA security.*** *For every pair of* PPT *Turing machines* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, there exists a negligible function* negl *such that for every* $n \in \mathbb{N}$ *and* $z \in \{0,1\}^*$,

$$\Pr\left[b = b' \;\middle|\; \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n) \\ (m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_1^{\mathsf{Dec}(\mathsf{sk}, \cdot)}(1^n, \mathsf{pk}, z) \ \textit{s.t.}\ |m_0| = |m_1| \\ b \leftarrow \{0,1\};\ c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b);\ b' \leftarrow \mathcal{A}_2^{\mathsf{Dec}'(\mathsf{sk}, \cdot)}(\mathsf{st}, c) \end{array} \right] \leq \frac{1}{2} + \mathsf{negl}(n),$$

  where the oracle $\mathsf{Dec}'(\mathsf{sk}, \cdot)$ *is the same as* $\mathsf{Dec}(\mathsf{sk}, \cdot)$ *except that it returns* $\perp$ *when* $\mathcal{A}_2$ *queries the challenge ciphertext* c *to it.*

Next, we introduce a new type of PKE schemes that we call puncturable public-key encryption.[11]

[11] Our definition of puncturable PKE is related to but is much simpler than the one that is proposed in [21].

Definition 12. *A public-key encryption scheme* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is called* puncturable *if there exist two* PPT *algorithms* $(\mathsf{PuncGen}, \mathsf{PuncDec})$ *that satisfy the following.*

- ***Correctness of punctured keys.*** *For every pair of* PPT *Turing machines* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, the outputs of the following two probabilistic experiments are computationally indistinguishable for every* $n \in \mathbb{N}$ *and* $z \in \{0,1\}^*$.
  - ***Experiment 1.***
    1. *Run* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$*,* $(m, \mathsf{st}) \leftarrow \mathcal{A}_1(1^n, \mathsf{pk}, z)$*,* $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m)$*,* $\mathsf{sk}_{\{c\}} \leftarrow \mathsf{PuncGen}(\mathsf{sk}, c)$*, and* $\mathsf{out} \leftarrow \mathcal{A}_2^{\mathsf{Dec}(\mathsf{sk},\cdot)}(\mathsf{st}, c, \mathsf{sk}_{\{c\}})$.
    2. *If* $\mathcal{A}_2$ *queried* c *to* Dec *in the previous step, the output of the experiment is* $\perp$*. Otherwise, the output is* out.
  - ***Experiment 2.***
    1. *Run* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$*,* $(m, \mathsf{st}) \leftarrow \mathcal{A}_1(1^n, \mathsf{pk}, z)$*,* $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m)$*,* $\mathsf{sk}_{\{c\}} \leftarrow \mathsf{PuncGen}(\mathsf{sk}, c)$*, and* $\mathsf{out} \leftarrow \mathcal{A}_2^{\mathsf{PuncDec}(\mathsf{sk}_{\{c\}},\cdot)}(\mathsf{st}, c, \mathsf{sk}_{\{c\}})$.
    2. *If* $\mathcal{A}_2$ *queried* c *to* $\mathsf{PuncDec}$ *in the previous step, the output of the experiment is* $\perp$*. Otherwise, the output is* out.
- ***Security of punctured keys.*** *For every pair of* PPT *Turing machines* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, there exists a negligible function* negl *such that for every* $n \in \mathbb{N}$ *and* $z \in \{0,1\}^*$,

$$\Pr\left[b = b' \;\middle|\; \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n) \\ (m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_1(1^n, \mathsf{pk}, z) \text{ s.t. } |m_0| = |m_1| \\ b \leftarrow \{0,1\};\ c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b) \\ \mathsf{sk}_{\{c\}} \leftarrow \mathsf{PuncGen}(\mathsf{sk}, c);\ b' \leftarrow \mathcal{A}_2(\mathsf{st}, c, \mathsf{sk}_{\{c\}}) \end{array} \right] \le \frac{1}{2} + \mathsf{negl}(n) \ .$$

It is easy to verify that the CCA-secure PKE of Dolev et al. [10] is puncturable. (Indeed, their proof of CCA security relies on the very fact that we can create a key with which we can emulate the decryption oracle without disturbing the security of the challenge ciphertext.) Thus, we have the following lemma.

Lemma 1. *Assume the existence of trapdoor permutations. Then, there exists a puncturable CCA-secure public-key encryption scheme.*

4 From 2-Round Delayed-Input Strong WI to 2-Round Special-Purpose Weak ZK

We show that 2-round delayed-input strong WI arguments satisfy a weak form of delayed-input weak ZK if their strong WI is proven by oblivious BB reductions.

Lemma 2. *Assume the existence of puncturable CCA-secure public-key encryption schemes. Then, there exists an NP language* $\boldsymbol{L}$ *such that if there exist*

- *a 2-round delayed-input interactive argument* (P, V) *for* $\boldsymbol{L}$ *and*
- *an oblivious black-box reduction* R_{SWI} *for showing the delayed-input strong WI of* (P, V) *based on a falsifiable assumption* (C, c)*,*

then either (i) the assumption (C, c) is false or (ii) (P, V) is special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge for every polynomial t and every inverse polynomial ϵ, where N is a polynomial and $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ is a sequence of efficient joint distributions such that each $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ ranges over $(\boldsymbol{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$. Furthermore, there exists a sequence of joint distributions $\overline{\mathcal{D}}_{xz} = \{(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)\}_{n\in\mathbb{N}}$ such that each $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$ ranges over $(\{0,1\}^n \setminus L) \times \{0,1\}^*$ and $\overline{\mathcal{D}}_{xz}$ is computationally indistinguishable from $\mathcal{D}_{xz} := \{(\mathcal{X}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$.

Proof. Let $\mathsf{PuncPKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{PuncGen}, \mathsf{PuncDec})$ be a puncturable CCA-secure PKE and $\mathbf{L}$ be the NP language that consists of all the public-key–ciphertext pairs of $\mathsf{PuncPKE}$ such that either 0 or 1 is encrypted (the public key is not necessarily honestly generated), i.e.,

$$\mathbf{L} := \left\{(\mathsf{pk}, \mathsf{ct}) \;\middle|\; \exists b \in \{0,1\}, r \in \{0,1\}^{\mathsf{poly}(n)} \text{ s.t. } \mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, b; r)\right\} \quad .$$

Assume, as stated in the statement of the lemma, the existence of a 2-round delayed-input interactive argument (P, V) and an oblivious black-box reduction R_{SWI} for showing the delayed-input strong WI of (P, V) based on a falsifiable assumption (C, c). For any inverse polynomial ϵ', let $Q_{\epsilon'}$ denote a polynomial such that for every delayed-input verifier V^*, every $n \in \mathbb{N}$, every two joint distributions $\mathcal{D}_n^0 = (\mathcal{X}_n^0, \mathcal{W}_n^0)$ and $\mathcal{D}_n^1 = (\mathcal{X}_n^1, \mathcal{W}_n^1)$ over $\mathbf{R_L} \cap (\{0,1\}^n \times \{0,1\}^*)$, and every $z \in \{0,1\}^*$, if it holds

$$\Pr\left[\langle P(w), V^*(z)\rangle(x) = b \;\middle|\; b \leftarrow \{0,1\};\; (x, w) \leftarrow (\mathcal{X}_n^b, \mathcal{W}_n^b)\right] \geq \frac{1}{2} + \epsilon'(n) \quad ,$$

then either (i) $R_{\text{SWI}}^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{1/\epsilon'(n)})$ breaks the assumption (C, c) on n with advantage $1/Q_{\epsilon'}(n)$ or (ii) $R_{\text{SWI}}^{V_z^*, \mathcal{D}_n^0, \mathcal{D}_n^1}(1^n, 1^{1/\epsilon'(n)})$ distinguishes $\mathcal{X}_n^0$ and $\mathcal{X}_n^1$ with advantage $1/Q_{\epsilon'}(n)$. (Such a polynomial is guaranteed to exist because of our assumption on R_{SWI}.) Fix any polynomial t and inverse polynomial ϵ.

At a high level, the proof proceeds as outlined in Section 2.2. Specifically, for any verifier and distinguisher against the weak ZK of (P, V), we first define a cheating verifier V_{SWI}^* against the strong WI of (P, V). Then, we proceed with case analysis about the behavior of $R_{\text{SWI}}^{V_{\text{SWI}}^*}$, where in the first case, we show that we can efficiently break the assumption (C, c) by using R_{SWI}, and in the second case, we show that we can obtain a simulator for weak ZK by using R_{SWI}. We note that in what follows, we use several constants that are chosen rather arbitrarily so that the proof works.

We first introduce distributions over $\mathbf{R_L}$ and a delayed-input verifier against the strong WI of (P, V). For any $n \in \mathbb{N}$, let Keys_n be the set of all the keys that can be output by $\mathsf{Gen}(1^n)$, i.e., $\mathsf{Keys}_n := \{(\mathsf{pk}, \mathsf{sk}) \mid \exists r \in \{0,1\}^* \text{s.t. } (\mathsf{pk}, \mathsf{sk}) = \mathsf{Gen}(1^n; r)\}$. Then, for any $n \in \mathbb{N}$ and any $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{Keys}_n$, let $\mathcal{D}_{\mathsf{pk}}^0$ and $\mathcal{D}_{\mathsf{pk}}^1$ be the distributions that are defined over $\mathbf{R_L}$ as follows: $\forall b \in \{0,1\}$, $\mathcal{D}_{\mathsf{pk}}^b := \{((\mathsf{pk}, \mathsf{ct}), (b, r)) \mid r \leftarrow \{0,1\}^{\mathsf{poly}(n)};\; \mathsf{ct} := \mathsf{Enc}(\mathsf{pk}, b; r)\}$ i.e., the first part of $\mathcal{D}_{\mathsf{pk}}^b$ outputs pk and a random encryption of b, and the second part

Algorithm 1 Delayed-input strong WI verifier $V^*_{\text{SWI}}[n, z, \mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$, where $z = z_V \,\|\, z_D$.

1. On input 1^n, invoke $V^*_{\text{WZK}}(1^n, z_V)$ and let it interact with the external prover. Let $x^\star = (\mathsf{pk}^\star, \mathsf{ct}^\star)$ denote the statement that is obtained in the last round of the interaction and $\mathsf{out}^\star$ denote the output of V^*_{WZK}. If $\mathsf{pk}^\star \neq \mathsf{pk}$, output a random bit and abort.
2. Sample a key key for a pseudorandom function PRF. In the following, whenever new randomness is required, it is obtained by applying $\mathsf{PRF}(\mathsf{key}, \cdot)$ on the transcript that is exchanged with the prover in the previous step. (The previous step does not require randomness since V^*_{WZK} is assumed to be deterministic.)
3. **(Approximation of honest prover's success probability.)** Obtain a $(\epsilon(n)/16, \mathsf{negl}(n))$-approximation $\tilde{p}$ of
$$p := \Pr\left[D_{\text{WZK}}(x, z_D, \langle P(w), V^*_{\text{WZK}}(z_V)\rangle(x)) = 1 \,\middle|\, \begin{array}{l}(\mathsf{pk}', \mathsf{sk}') \leftarrow \mathsf{Gen}(1^n)\\ b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}'}\end{array}\right] .$$
4. **(Approximation of external prover's success probability.)** Obtain a $(\epsilon(n)/16, \mathsf{negl}(n))$-approximation $\tilde{p}^\star$ of $p^\star := \Pr\left[D_{\text{WZK}}(x^\star, z_D, \mathsf{out}^\star) = 1\right]$.
5. Output a random bit and abort if $\tilde{p}^\star < \tilde{p} - \epsilon(n)/2$ (which suggests that the external prover with the given statement $x^\star$ is not likely to convince D_{WZK} with probability as high as an honest prover with a random statement). Otherwise, run $b \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct}^\star)$ and output b.

outputs b and the randomness of the encryption. We use $(\mathcal{X}^b_{\mathsf{pk}}, \mathcal{W}^b_{\mathsf{pk}})$ to denote the joint distributions such that $\mathcal{X}^b_{\mathsf{pk}}$ denotes the first part of $\mathcal{D}^b_{\mathsf{pk}}$ and $\mathcal{W}^b_{\mathsf{pk}}$ denotes the second part of $\mathcal{D}^b_{\mathsf{pk}}$. Next, for any $n \in \mathbb{N}$, any $z = z_V \,\|\, z_D \in \{0,1\}^*$, any $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{Keys}_n$, and any pair of a (deterministic) delayed-input verifier V^*_{WZK} and a (probabilistic) distinguisher D_{WZK} against the weak zero-knowledge property of (P, V), let $V^*_{\text{SWI}}[n, z, \mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$ be the delayed-input verifier described in Algorithm 1. Note that due to the correctness of $\mathsf{PuncPKE}$, our verifier $V^*_{\text{SWI}}[n, z, \mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$ distinguishes $\mathcal{D}^0_{\mathsf{pk}}$ and $\mathcal{D}^1_{\mathsf{pk}}$ with probability 1 when it interacts with a prover that passes the test in the last step of $V^*_{\text{SWI}}[n, z, \mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$. In the following, we usually write $V^*_{\text{SWI}}[n, z, \mathsf{pk}, \mathsf{sk}, V^*_{\text{WZK}}, D_{\text{WZK}}]$ as V^*_{SWI} for editorial simplicity.

We proceed with case analysis about the behavior of the strong WI reduction R_{SWI} in the setting where R_{SWI} is combined with our strong WI verifier V^*_{SWI}. Specifically, we consider two cases about the behavior of R_{SWI} in the setting where we use $R^{V^*_{\text{SWI}}}_{\text{SWI}}$ as a distinguisher against $\mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}$ for randomly chosen $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$. Toward this end, we first introduce the following notations about $(\mathsf{pk}, \mathsf{sk})$ of $\mathsf{PuncPKE}$. For any n, z, $(\mathsf{pk}, \mathsf{sk})$, V^*_{WZK}, and D_{WZK}:

– $(\mathsf{pk}, \mathsf{sk})$ is called *interesting (w.r.t. $(n, z, V^*_{\text{WZK}}, D_{\text{WZK}})$)* if
$$\Pr\left[\langle P(w), V^*_{\text{SWI}}\rangle(x) = b \,\middle|\, b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}}\right] \geq \frac{1}{2} + \frac{\epsilon(n)}{18} . \qquad (1)$$

Intuitively, $(\mathsf{pk},\mathsf{sk})$ is interesting if $V^*_{\mathrm{SWI}}[n,z,\mathsf{pk},\mathsf{sk},V^*_{\mathrm{WZK}},D_{\mathrm{WZK}}]$ breaks the strong WI of (P,V) w.r.t. $\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}$ with high advantage (which implies that R_{SWI} either breaks (C,c) or distinguishes $\mathcal{X}^0_{\mathsf{pk}}$ and $\mathcal{X}^1_{\mathsf{pk}}$ given V^*_{SWI}).

- $(\mathsf{pk},\mathsf{sk})$ is called *type-1 interesting* if it is interesting and in addition satisfies the following.

$$\Pr\left[\mathsf{INTERESTING\text{-}QUERY}\;\middle|\;\begin{array}{l} b\leftarrow\{0,1\};\ (x,w)\leftarrow\mathcal{D}^b_{\mathsf{pk}} \\ b'\leftarrow R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n,1^{36/\epsilon(n)},x)\end{array}\right]\le\frac{1}{4Q_{\epsilon/36}(n)}\ ,$$

where (i) $Q_{\epsilon/36}$ is the polynomial that is introduced at the beginning of the proof and (ii) $\mathsf{INTERESTING\text{-}QUERY}$ is the event that is defined as follows: through oracle queries to V^*_{SWI}, the reduction $R_{\mathrm{SWI}}(1^n,1^{36/\epsilon(n)},x)$ invokes an execution of (P,V) in which R_{SWI} forwards the statement x to V^*_{SWI} along with an accepting prover message (i.e., a message that passes the test in the last step of V^*_{SWI}). Note that by the construction of V^*_{SWI}, $\mathsf{INTERESTING\text{-}QUERY}$ implies that R_{SWI} produces a prover message that convinces D_{WZK} with high probability on the statement x—thus, intuitively, $(\mathsf{pk},\mathsf{sk})$ is type-1 interesting if R_{SWI} can either break (C,c) or distinguish $\mathcal{X}^0_{\mathsf{pk}}$ and $\mathcal{X}^1_{\mathsf{pk}}$ without producing such a prover message.
- $(\mathsf{pk},\mathsf{sk})$ is called *type-2 interesting* if it is interesting but is not type-1 interesting.

Now, we consider the following two cases.

- **Case 1.** There exist a deterministic polynomial-time delayed-input verifier V^*_{WZK} and a probabilistic t-time distinguisher D_{WZK} such that for infinitely many $n\in\mathbb{N}$ there exists $z\in\{0,1\}^*$ such that

$$\Pr\left[(\mathsf{pk},\mathsf{sk})\text{ is type-1 interesting}\mid(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^n)\right]\ge\frac{\epsilon(n)}{8}\ . \tag{2}$$

- **Case 2.** The condition of Case 1 does not hold.

We analyze each case below.

Analysis of Case 1. We show that R_{SWI} can be used to break the assumption (C,c). Fix any V^*_{WZK}, D_{WZK}, n, and z such that we have (2). Note that for any interesting $(\mathsf{pk},\mathsf{sk})$, we have (1) and therefore for any constant $k\ge 18$, either (i) $R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}$ breaks the assumption (C,c) with advantage $1/Q_{\epsilon/k}(n)$ or (ii) $R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}$ distinguishes $\mathcal{X}^0_{\mathsf{pk}}$ and $\mathcal{X}^1_{\mathsf{pk}}$ with advantage $1/Q_{\epsilon/k}(n)$.

We first show, roughly speaking, that with high probability over the sampling of $(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^n)$, we obtain a type-1 interesting $(\mathsf{pk},\mathsf{sk})$ such that $R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}$ breaks the assumption (C,c)—later, we use this to argue that we can break the assumption (C,c) by finding such a type-1 interesting $(\mathsf{pk},\mathsf{sk})$ via sampling. Formally, let us say that a type-1 interesting $(\mathsf{pk},\mathsf{sk})$ is *bad* if

$R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)})$ does not break the assumption (C, c) on n with advantage $1/2Q_{\epsilon/36}(n)$, i.e.,

$$\Pr\left[\langle R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^{36/\epsilon(n)}), C\rangle(1^n) = 1\right] \leq c + \frac{1}{2Q_{\epsilon/36}(n)} \ .$$

Then, what we show is that a bad type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ is sampled with probability at most $\epsilon(n)/16$ in the sampling of $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$. Assume for contradiction that we sample a bad type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ with probability greater than $\epsilon(n)/16$. Then, consider the following adversary $\mathcal{A}_{\mathrm{CCA}}$ against the CCA security of $\mathsf{PuncPKE}$.

1. On input $(1^n, \mathsf{pk}, z)$, the adversary $\mathcal{A}_{\mathrm{CCA}}$ sends $m_0 := 0$ and $m_1 := 1$ to the challenger as the challenge plaintexts.
2. On receiving the challenge ciphertext ct, the adversary $\mathcal{A}_{\mathrm{CCA}}$ first does the following to check whether or not the key pair $(\mathsf{pk}, \mathsf{sk})$ that the challenger has is likely to be bad type-1 interesting.
   (a) Obtain a $(1/4Q_{\epsilon/36}(n), \mathsf{negl}(n))$-approximation $\tilde{p}_1$ of
   $$p_1 := \Pr\left[\langle P(w), V^*_{\mathrm{SWI}}\rangle(x) = b \;\middle|\; b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}}\right] \ ,$$
   where during the approximation, the decryption oracle $\mathsf{Dec}(\mathsf{sk}, \cdot)$ is used to emulate V^*_{SWI} efficiently without knowing sk. (Since the definition of p_1 is independent of ct, the probability that ct needs to be queried to $\mathsf{Dec}(\mathsf{sk}, \cdot)$ is negligible.)
   (b) Obtain a $(1/4Q_{\epsilon/36}(n), \mathsf{negl}(n))$-approximation $\tilde{p}_2$ of
   $$p_2 := \Pr\left[\textsf{INTERESTING-QUERY} \;\middle|\; \begin{array}{l} b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}} \\ b' \leftarrow R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)}, x) \end{array}\right] \ ,$$
   where as above the decryption oracle $\mathsf{Dec}(\mathsf{sk}, \cdot)$ is used to emulate V^*_{SWI} during the approximation.
   (c) Obtain a $(1/4Q_{\epsilon/36}(n), \mathsf{negl}(n))$-approximation $\tilde{p}_3$ of
   $$p_3 := \Pr\left[\langle R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^{36/\epsilon(n)}), C\rangle(1^n) = 1\right] \ ,$$
   where as above the decryption oracle $\mathsf{Dec}(\mathsf{sk}, \cdot)$ is used to emulate V^*_{SWI} during the approximation.
   (d) If $\tilde{p}_1 < 1/2 + \epsilon(n)/18 - 1/4Q_{\epsilon/36}(n)$, $\tilde{p}_2 > 1/2Q_{\epsilon/36}(n)$, or $\tilde{p}_3 \geq c + 3/4Q_{\epsilon/36}(n)$ (which suggests that $(\mathsf{pk}, \mathsf{sk})$ is unlikely to be bad type-1 interesting), output a random bit and abort.
3. Finally, let $x^\star := (\mathsf{pk}, \mathsf{ct})$ and run $b^\star \leftarrow R_{\mathrm{SWI}}^{V^*_{\mathrm{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)}, x^\star)$, where as above the decryption oracle $\mathsf{Dec}(\mathsf{sk}, \cdot)$ is used to emulate V^*_{SWI}. If $\textsf{INTERESTING-QUERY}$ occurs during the execution of R_{SWI}, output a random bit. Otherwise, output $b^\star$.

We now analyze $\mathcal{A}_{\text{CCA}}$. Let ABORT be the event that $\mathcal{A}_{\text{CCA}}$ aborts, and APPROX-FAIL be the event that the approximation of any of $\tilde{p}_1, \tilde{p}_2, \tilde{p}_3$ fails, i.e., $\max(|p_1 - \tilde{p}_1|, |p_2 - \tilde{p}_2|, |p_3 - \tilde{p}_3|) > 1/4Q_{\epsilon/36}(n)$. From the union bound, we have $\Pr[\mathsf{APPROX\text{-}FAIL}] \leq \mathsf{negl}(n)$. Also, we have $\Pr[\neg\mathsf{ABORT}] \geq \epsilon(n)/16 - \mathsf{negl}(n)$ since $\mathcal{A}_{\text{CCA}}$ does not abort when pk is the public key of a bad type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ and APPROX-FAIL does not occur. Now, under the condition that neither APPROX-FAIL nor ABORT occurs, we have

$$p_1 \geq \tilde{p}_1 - \frac{1}{4Q_{\epsilon/36}(n)} \geq \frac{1}{2} + \frac{\epsilon(n)}{18} - \frac{1}{2Q_{\epsilon/36}(n)} \geq \frac{1}{2} + \frac{\epsilon(n)}{36} \ , \tag{3}$$

$$p_2 \leq \tilde{p}_2 + \frac{1}{4Q_{\epsilon/36}(n)} \leq \frac{3}{4Q_{\epsilon/36}(n)} \ , \text{ and} \tag{4}$$

$$p_3 \leq \tilde{p}_3 + \frac{1}{4Q_{\epsilon/36}(n)} < c + \frac{1}{Q_{\epsilon/36}(n)} \ , \tag{5}$$

where the last inequality in (3) follows since we can assume without loss of generality that $Q_{\epsilon/36}(n)$ is sufficiently large and satisfies $1/Q_{\epsilon/36}(n) \leq \epsilon(n)/18$. Note that when we have (3) and (5) (where the former means that V^*_{SWI} breaks the strong WI of (P, V) w.r.t. $\mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}$ with advantage $\epsilon(n)/36$ while the latter means that $R_{\text{SWI}}^{V^*_{\text{SWI}}, \mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}}$ does not break (C, c) with advantage $1/Q_{\epsilon/36}(n)$), it is guaranteed that $R_{\text{SWI}}^{V^*_{\text{SWI}}, \mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}}$ distinguishes $\mathcal{X}^0_{\mathsf{pk}}$ and $\mathcal{X}^1_{\mathsf{pk}}$ with advantage $1/Q_{\epsilon/36}(n)$ due to the definition of $Q_{\epsilon/36}$. Thus, by additionally using (4) and recalling the definitions of $\mathcal{X}^0_{\mathsf{pk}}$ and $\mathcal{X}^1_{\mathsf{pk}}$ (i.e., that $\mathcal{X}^b_{\mathsf{pk}}$ outputs pk and a random encryption of b), we conclude that $\mathcal{A}_{\text{CCA}}$ wins with advantage at least

$$\left(\frac{1}{Q_{\epsilon/36}(n)} - \Pr\left[\begin{array}{l}\mathsf{INTERESTING\text{-}QUERY} \text{ occurs}\\ \text{in Step 3 of } \mathcal{A}_{\text{CCA}}\end{array}\right]\right) \times \Pr[\neg\mathsf{ABORT}] - \Pr[\mathsf{APPROX\text{-}FAIL}]$$
$$\geq \frac{1}{4Q_{\epsilon/36}(n)} \times \left(\frac{\epsilon(n)}{16} - \mathsf{negl}(n)\right) - \mathsf{negl}(n) = \frac{1}{\mathsf{poly}(n)} \ .$$

Since this is a contradiction, we conclude that we sample a bad type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ with probability at most $\epsilon(n)/16$ in the sampling of $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$.

We are now ready to show that R_{SWI} can be used to break the assumption (C, c). Consider the following adversary $\mathcal{A}$ against (C, c).

1. Repeat the following to find a type-1 interesting $(\mathsf{pk}^\star, \mathsf{sk}^\star)$ that is likely to be useful to break (C, c).
   (a) Sample $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$.
   (b) Obtain $(1/8Q_{\epsilon/36}(n), \mathsf{negl}(n))$-approximations $\tilde{p}_1, \tilde{p}_2, \tilde{p}_3$ of p_1, p_2, p_3, where p_1, p_2, p_3 are defined as in $\mathcal{A}_{\text{CCA}}$ above and sk is used (instead of the decryption oracle) to emulate V^*_{SWI} efficiently during the approximations.
   (c) If $\tilde{p}_1 \geq 1/2 + \epsilon(n)/18 - 1/8Q_{\epsilon/36}(n)$, $\tilde{p}_2 \leq 3/8Q_{\epsilon/36}(n)$, and $\tilde{p}_3 \geq c + 3/8Q_{\epsilon/36}(n)$ (which suggests that $(\mathsf{pk}, \mathsf{sk})$ is likely to be "good" type-1 interesting), let $(\mathsf{pk}^\star, \mathsf{sk}^\star) := (\mathsf{pk}, \mathsf{sk})$ and exit the loop to go to the next step.

If $(\mathsf{pk}^\star, \mathsf{sk}^\star)$ cannot be found within $128Q_{\epsilon/36}(n)/\epsilon(n)$ attempts, abort.

2. Let $R_{\text{SWI}}^{V^*_{\text{SWI}}, \mathcal{D}^0_{\mathsf{pk}^\star}, \mathcal{D}^1_{\mathsf{pk}^\star}}(1^n, 1^{36/\epsilon(n)})$ interact with the challenger C.

We analyze $\mathcal{A}$ as follows. From (2) and what is shown in the previous paragraph, with probability at least $\epsilon(n)/8 - \epsilon(n)/16 = \epsilon(n)/16$ over the sampling of $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$, we obtain a type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ such that $R_{\text{SWI}}^{V^*_{\text{SWI}}, \mathcal{D}^0_{\mathsf{pk}}, \mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)})$ breaks the assumption (C, c) with advantage at least $1/2Q_{\epsilon/36}(n)$. Let us call such a type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ *good*, and observe when $\mathcal{A}$ samples a good type-1 interesting $(\mathsf{pk}, \mathsf{sk})$, it does not abort unless the approximation of any of $\tilde{p}_1, \tilde{p}_2, \tilde{p}_3$ fails. Also, observe that (i) by Markov's inequality, $\mathcal{A}$ samples a good type-1 interesting $(\mathsf{pk}, \mathsf{sk})$ within $128Q_{\epsilon/36}(n)/\epsilon(n)$ attempts with probability at least $1 - 1/8Q_{\epsilon/36}(n)$, and (ii) when $\mathcal{A}$ does not abort, $\mathcal{A}$ wins with probability at least $c + 3/8Q_{\epsilon/36}(n) - 1/8Q_{\epsilon/36}(n) = c + 1/4Q_{\epsilon/36}(n)$ unless the approximation of $\tilde{p}_3$ fails. Thus, $\mathcal{A}$ wins with probability at least

$$\begin{aligned} &\Pr[\mathcal{A} \text{ wins} \mid \mathcal{A} \text{ does not abort}] - \Pr[\mathcal{A} \text{ aborts}] \\ &\geq \left(c + \frac{1}{4Q_{\epsilon/36}(n)} - \mathsf{negl}(n)\right) - \left(\frac{1}{8Q_{\epsilon/36}(n)} + \mathsf{negl}(n)\right) = c + \frac{1}{\mathsf{poly}(n)} \enspace . \end{aligned}$$

We thus conclude that the assumption (C, c) is false in this case.

Analysis of Case 2. We show that R_{SWI} can be used to construct a simulator for the special-purpose distributional super-weak (ϵ, t)-zero-knowledge property of (P, V). For each $n \in \mathbb{N}$, let $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ be the following joint distributions.

$$(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n) := \left\{ ((\mathsf{pk}, \mathsf{ct}), (b, r), \mathsf{sk}_{\{\mathsf{ct}\}}) \;\middle|\; \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n) \\ b \leftarrow \{0,1\};\ r \leftarrow \{0,1\}^{\mathsf{poly}(n)} \\ \mathsf{ct} := \mathsf{Enc}(\mathsf{pk}, b; r) \\ \mathsf{sk}_{\{\mathsf{ct}\}} \leftarrow \mathsf{PuncGen}(\mathsf{sk}, \mathsf{ct}) \end{array} \right\} \enspace .$$

(Note that $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ indeed ranges over $(\mathbf{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$ as required.[12] Also, note that $(\mathcal{X}_n, \mathcal{W}_n)$ is identically distributed with $\{(x, w) \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n); b \leftarrow \{0,1\}; (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}}\}$.) Let N be the polynomial such that $N(n, 1/\epsilon(n)) := 320Q_{\epsilon/36}(n)/\epsilon(n)^2$.

For any deterministic polynomial-time delayed-input verifier V^*_{WZK} and a probabilistic t-time distinguisher D_{WZK}, we consider the simulator S described in Algorithm 2.

We now proceed with the analysis of S. Fix any V^*_{WZK} and D_{WZK}. Since it is assumed that the condition of Case 1 does not hold, we have that for every sufficiently large $n \in \mathbb{N}$ and every $z = z_V \| z_D \in \{0,1\}^*$,

$$\Pr[(\mathsf{pk}, \mathsf{sk}) \text{ is type-1 interesting} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)] < \frac{\epsilon(n)}{8} \enspace . \tag{6}$$

[12] We assume without loss of generality that on security parameter 1^n, Gen and Enc generate $(\mathsf{pk}, \mathsf{ct})$ such that $|(\mathsf{pk}, \mathsf{ct})| = n$.

Algorithm 2 Weak zero-knowledge simulator S.

```
Input: $\{x_i, z_{x,i}\}_{i\in[N_n]}$ and $z_V, z_D \in \{0,1\}^*$, where $N_n := N(n, 1/\epsilon(n))$ and each $(x_i, z_{x,i})$ is sampled from $(\mathcal{X}_n, \mathcal{Z}_n)$.
Hardwired information: the verifier $V^*_{\text{WZK}}$ and the distinguisher $D_{\text{WZK}}$.

1. Let $z := z_V \,\|\, z_D$. Then, for each $i \in [N_n]$, do the following.
   (a) Parse $(x_i, z_{x,i})$ as $((\mathsf{pk}, \mathsf{ct}), \mathsf{sk}_{\{\mathsf{ct}\}})$, and run $R_{\text{SWI}}^{V^*_{\text{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)}, x_i)$ as a
       distinguisher for $\mathcal{D}^0_{\mathsf{pk}}$ and $\mathcal{D}^1_{\mathsf{pk}}$ to see whether INTERESTING-QUERY occurs, where
       the punctured secret key $\mathsf{sk}_{\{\mathsf{ct}\}}$ is used to emulate $V^*_{\text{SWI}}$ efficiently for $R_{\text{SWI}}$ until
       INTERESTING-QUERY occurs. (Recall that INTERESTING-QUERY occurs if $R_{\text{SWI}}$ makes
       a query (to $V^*_{\text{SWI}}$) that contains $x_i$ and an accepting prover message.)
   (b) It INTERESTING-QUERY occurs, let $i^\star := i$, and let $\mathsf{out}^\star$ denote the output of $V^*_{\text{WZK}}$
       that is computed inside $V^*_{\text{SWI}}$ when the query that causes INTERESTING-QUERY is
       made; then, exit the loop and go to the next step.
2. If $(i^\star, \mathsf{out}^\star)$ is not defined in the above step, abort. Otherwise, output $(i^\star, \mathsf{out}^\star)$.
```

Fix any such n and $z = z_V \,\|\, z_D$. Let p be defined by

$$p := \Pr\left[D_{\text{WZK}}(x, z_D, \langle P(w), V^*_{\text{WZK}}(z_V)\rangle(x)) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n, \mathcal{W}_n)\right] \quad . \tag{7}$$

(Note that p is defined as in description of V^*_{SWI} in Algorithm 1.)

We first make a simplifying assumption. First, note that S runs the reduction R_{SWI} with our (probabilistic) verifier V^*_{SWI}. Following the conventions stated in Section 3.3, in general the reduction R_{SWI} can make V^*_{SWI} reuse the same randomness multiple times when it makes queries to V^*_{SWI}. However, since V^*_{SWI} obtains randomness by applying PRF on the transcript exchanged with the prover (where the prover message is actually given by R_{SWI}), we can safely think, by assuming without loss of generality that R_{SWI} never makes the same query twice to V^*_{SWI} while making V^*_{SWI} reuse the same randomness, as if V^*_{SWI} always uses new true randomness in each invocation during the execution of S. Second, note that S uses the punctured secret key $\mathsf{sk}_{\{\mathsf{ct}\}}$ to emulate V^*_{SWI} for R_{SWI}. We can however safely think as if S uses the real secret key sk to perfectly emulate V^*_{SWI} since the correctness of punctured keys of $\mathsf{PuncPKE}$ guarantees that the output of R_{SWI} (and hence that of S) is indistinguishable in these two cases. (Note that by the definition of INTERESTING-QUERY, decrypting ct is not required for the emulation of V^*_{SWI} unless INTERESTING-QUERY occurs.)

Next, we bound the probability that S aborts. Toward this end, it suffices to show that we have

$$\Pr\left[(\mathsf{pk}, \mathsf{sk}) \text{ is type-2 interesting} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)\right] \geq \frac{\epsilon(n)}{8} \quad . \tag{8}$$

Indeed, by combining (8) with the definition of type-2 interesting keys, we obtain

$$\Pr\left[\text{INTERESTING-QUERY} \;\middle|\; \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n) \\ b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}} \\ b' \leftarrow R_{\text{SWI}}^{V^*_{\text{SWI}},\mathcal{D}^0_{\mathsf{pk}},\mathcal{D}^1_{\mathsf{pk}}}(1^n, 1^{36/\epsilon(n)}, x) \end{array}\right] \geq \frac{\epsilon(n)}{32 Q_{\epsilon/36}(n)} \quad ,$$

and thus, by using Markov's inequality, we can bound the probability that S aborts as follows.

$$\Pr\left[S(\{x_i, z_{x,i}\}_{i\in[N_n]}, z_V, z_D) \text{ aborts} \,\middle|\, (x_i, z_{x,i}) \leftarrow (\mathcal{X}_n, \mathcal{Z}_n) \text{ for } \forall i \in [N_n]\right] \leq \frac{\epsilon(n)}{10}. \quad (9)$$

So, we focus on showing (8). Observe that from (7) and an average argument, it follows that with probability at least $\epsilon(n)/4$ over the choice of $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$, we obtain $(\mathsf{pk}, \mathsf{sk})$ such that

$$\Pr\left[D_{\text{WZK}}(x, z_D, \langle P(w), V^*_{\text{WZK}}(z_V)\rangle(x)) = 1 \,\middle|\, \begin{array}{l} b \leftarrow \{0,1\} \\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}} \end{array}\right] \geq p - \frac{\epsilon(n)}{4} \quad . \quad (10)$$

For any such $(\mathsf{pk}, \mathsf{sk})$, it follows from (10) and an average argument that with probability at least $\epsilon(n)/8$ over the choice of $b \leftarrow \{0,1\}$, $(x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}}$, and $\mathsf{out} \leftarrow \langle P(w), V^*_{\text{WZK}}(z_V)\rangle(x)$, we obtain out such that

$$\Pr\left[D_{\text{WZK}}(x, z_D, \mathsf{out}) = 1\right] \geq p - \frac{3\epsilon(n)}{8} \quad . \quad (11)$$

Now, for any $(\mathsf{pk}, \mathsf{sk})$ such that we have (10), we have

$$\begin{aligned} &\Pr\left[\langle P(w), V^*_{\text{SWI}}(z)\rangle(x) = b \,\middle|\, b \leftarrow \{0,1\};\ (x, w) \leftarrow \mathcal{D}^b_{\mathsf{pk}}\right] \\ &= \frac{1}{2} + \frac{1}{2}\Pr\left[V^*_{\text{SWI}} \text{ does not abort}\right] \geq \frac{1}{2} + \frac{1}{2}\left(\frac{\epsilon(n)}{8} - \mathsf{negl}(n)\right) \geq \frac{1}{2} + \frac{\epsilon(n)}{18} \quad , \quad (12) \end{aligned}$$

where to see the first inequality, observe that we have $\Pr\left[V^*_{\text{SWI}} \text{ does not abort}\right] \geq \epsilon(n)/8 - \mathsf{negl}(n)$ since if the output out of V^*_{WZK} that is computed in the first step of V^*_{SWI} satisfies (11), we have $\tilde{p}^\star \geq p^\star - \epsilon(n)/16 \geq p - \epsilon(n)/16 - 3\epsilon(n)/8 \geq \tilde{p} - \epsilon(n)/16 - 3\epsilon(n)/8 - \epsilon(n)/16 = \tilde{p} - \epsilon(n)/2$ in V^*_{SWI} unless the approximations of p and $p^\star$ fails (the second inequality follows from (11)). Thus, by (12) and the definition of interesting keys, any $(\mathsf{pk}, \mathsf{sk})$ such that we have (10) is interesting, so we have

$$\Pr\left[(\mathsf{pk}, \mathsf{sk}) \text{ is interesting} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)\right] \geq \frac{\epsilon(n)}{4} \quad . \quad (13)$$

Combining (6) and (13), we obtain (8).

Next, we analyze the behavior of S under the condition that it does not abort. Since S makes at most polynomially many queries to V^*_{SWI}, it follows from a union bound that with overwhelming probability, in each query the approximations of p and $p^\star$ by V^*_{SWI} are correct, i.e., $\max(|p - \tilde{p}|, |p^\star - \tilde{p}^\star|) \leq \epsilon(n)/16$. Thus, under the condition that S does not abort, with overwhelming probability the output $(i^\star, \mathsf{out}^\star)$ of $S(\{x_i, z_{x,i}\}_{i\in[N_n]}, z_V, z_D)$ satisfies

$$\Pr\left[D_{\text{WZK}}(x_{i^\star}, z_D, \mathsf{out}^\star) = 1\right] \geq \tilde{p} - \frac{\epsilon(n)}{2} - \frac{\epsilon(n)}{16} \geq p - \frac{5\epsilon(n)}{8} \quad . \quad (14)$$

Finally, by combining (9) and (14), we obtain

$$\begin{aligned}
&\Pr\left[D_{\text{WZK}}(x_{i^\star}, z_D, \mathsf{out}^\star) = 1 \;\middle|\; \begin{array}{l}(x_i, z_{x,i}) \leftarrow (\mathcal{X}_n, \mathcal{Z}_n) \text{ for } \forall i \in [N_n] \\ (i^\star, \mathsf{out}^\star) \leftarrow S(\{x_i, z_{x,i}\}_{i\in[N_n]}, z_V, z_D)\end{array}\right] \\
&\geq p - \frac{5\epsilon(n)}{8} - \frac{\epsilon(n)}{10} - \mathsf{negl}(n) \\
&\geq \Pr\left[D_{\text{WZK}}(x, z_D, \langle P(w), V^*_{\text{WZK}}(z_V)\rangle(x)) = 1 \mid (x, w) \leftarrow (\mathcal{X}_n, \mathcal{W}_n)\right] - \epsilon(n)
\end{aligned}$$

as required. Thus, (P, V) is special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge in this case.

Completing the proof of the first part of Lemma 2. Combining the analyses of Case 1 and Case 2, we conclude that for any $t, \epsilon, V^*_{\text{WZK}}, D_{\text{WZK}}$, either the assumption (C, c) is false or S is a good simulator for the delayed-input special-purpose $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge property of (P, V), where $\mathcal{D}_{xwz}$ and N are defined as above.

Proof of the furthermore part. We define $\overline{\mathcal{D}}_{xz} = \{(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)\}_{n\in\mathbb{N}}$ by

$$(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n) := \left\{((\mathsf{pk}, \mathsf{ct}), \mathsf{sk}_{\{\mathsf{ct}\}}) \;\middle|\; \begin{array}{l}(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n);\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, 2) \\ \mathsf{sk}_{\{\mathsf{ct}\}} \leftarrow \mathsf{PuncGen}(\mathsf{sk}, \mathsf{ct})\end{array}\right\} .$$

Due to the (perfect) correctness of $\mathsf{PuncPKE}$, each $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$ indeed ranges over $(\{0,1\}^n \setminus L) \times \{0,1\}^*$. Also, $\overline{\mathcal{D}}_{xz}$ is indeed computationally indistinguishable from $\mathcal{D}_{xz} = \{(\mathcal{X}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ because of the security of $\mathsf{PuncPKE}$. □

5 From Special-Purpose Weak ZK to Special-Purpose Pre-Processing ZK

We show that special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge implies special-purpose delayed-input $(\mathcal{D}_{xwz}, N')$-distributional pre-processing (t', ϵ')-zero-knowledge for some N', t', ϵ'.

Lemma 3. *Let (P, V) be a 2-round delayed-input interactive argument for an NP language $\boldsymbol{L}$ and $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ be a sequence of joint distributions such that each $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ ranges over $(\boldsymbol{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$. Then, if there exists a polynomial N such that (P, V) is special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional super-weak (t, ϵ)-zero-knowledge for every polynomial t and inverse polynomial ϵ, there also exists a polynomial N' such that (P, V) is special-purpose delayed-input $(\mathcal{D}_{xwz}, N')$-distributional pre-processing (t', ϵ')-zero-knowledge for every polynomial t' and inverse polynomial ϵ'.*

As mentioned in Section 2.1, we prove this lemma by slightly modifying the proof of [8, Theorem 9] (where it is shown that a certain version of weak ZK implies a certain version of ZK as in this lemma). For lack of space, we defer the proof to the full version of this paper.

6 BB Impossibility of 2-Round Special-Purpose Pre-Processing ZK

We give a BB impossibility result about special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional pre-processing (t, ϵ)-zero-knowledge.

Lemma 4. *Let $\boldsymbol{L}$ be an NP language and $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ be a sequence of efficient joint distributions such that (i) each $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ ranges over $(\boldsymbol{R_L} \times \{0,1\}^*) \cap (\{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*)$ and (ii) there exists a sequence of joint distributions $\overline{\mathcal{D}}_{xz} = \{(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)\}_{n\in\mathbb{N}}$ such that $\overline{\mathcal{D}}_{xz}$ is computationally indistinguishable from $\mathcal{D}_{xz} \coloneqq \{(\mathcal{X}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ and each $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$ ranges over $(\{0,1\}^n \setminus L) \times \{0,1\}^*$. Then, if there exists a 2-round delayed-input interactive argument (P, V) for $\boldsymbol{L}$ such that*

- *there exists a polynomial N such that (P, V) is special-purpose delayed-input $(\mathcal{D}_{xwz}, N)$-distributional pre-processing (t, ϵ)-zero-knowledge for every polynomial t and every inverse polynomial ϵ, and*
- *there exists a black-box reduction R for showing the adaptive soundness of (P, V) based on a falsifiable assumption (C, c),*

then, the assumption (C, c) is false.

As mentioned in Section 2.2, the proof of this lemma closely follows the proof of [7, Theorem 2]. For lack of space, we defer the proof to the full version of this paper.

7 Obtaining Main Results

We obtain our main results by using the lemmas given in the previous sections.

BB impossibility of 2-round delayed-input weak ZK. By using Lemma 3 and Lemma 4, we obtain the following black-box impossibility result about 2-round delayed-input weak ZK.

Theorem 1. *Assume the existence of one-way functions. Then, there exists an NP language $\boldsymbol{L}$ such that if there exist (i) a 2-round delayed-input interactive argument (P, V) for $\boldsymbol{L}$ that is delayed-input distributional weak (t, ϵ)-zero-knowledge for every polynomial t and inverse polynomial ϵ and (ii) a black-box reduction R for showing the adaptive soundness of (P, V) based on a falsifiable assumption (C, c), then the assumption (C, c) is false.*

Proof. Let PRG be any pseudorandom generator (which can be obtained from one-way functions [24]) and $\mathbf{L}$ be the NP language that is defined by $\mathbf{L} \coloneqq \{\mathsf{PRG}(s) \mid s \in \{0,1\}^*\}$, where we assume without loss of generality that PRG is length-doubling. For each $n \in \mathbb{N}$, consider the following joint distributions $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)$ and $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)$: $(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n) \coloneqq \{(\mathsf{PRG}(s), s, \perp) \mid s \leftarrow \{0,1\}^{n/2}\}$ and $(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n) \coloneqq \{(r, \perp) \mid r \leftarrow \{0,1\}^n \setminus \mathbf{L}\}$. It is easy to see that $\{(\mathcal{X}_n, \mathcal{Z}_n)\}_{n\in\mathbb{N}}$ and $\{(\overline{\mathcal{X}}_n, \overline{\mathcal{Z}}_n)\}_{n\in\mathbb{N}}$ are computationally indistinguishable, and delayed-input distributional weak (t, ϵ)-zero-knowledge implies special-purpose delayed-input $(\mathcal{D}_{xwz}, 1)$-distributional super-weak (t, ϵ)-zero-knowledge, where $\mathcal{D}_{xwz} = \{(\mathcal{X}_n, \mathcal{W}_n, \mathcal{Z}_n)\}_{n \in \mathbb{N}}$. Now, the lemma follows from Lemma 3 and Lemma 4. □

BB impossibility of 2-round delayed-input strong WI. By combining Lemma 1, Lemma 2, Lemma 3, and Lemma 4, we immediately obtain the following black-box impossibility result about 2-round delayed-input strong WI.

Theorem 2. *Assume the existence of trapdoor permutations. Then, there exists an NP language $\boldsymbol{L}$ such that if there exist (i) a 2-round delayed-input interactive argument (P, V) for $\boldsymbol{L}$, (ii) an oblivious black-box reduction R_{SWI} for showing the delayed-input strong WI of (P, V) based on a falsifiable assumption (C, c), and (iii) a black-box reduction R' for showing the adaptive soundness of (P, V) based on a falsifiable assumption (C', c'), then either the assumption (C, c) is false or the assumption (C', c') is false.*

BB impossibility of 2-round (non-delayed-input) strong WI. By adjusting the proof of Lemma 2, we obtain the following black-box impossibility result about 2-round (non-delayed-input) strong WI.

Theorem 3. *Assume the existence of CCA-secure public-key encryption schemes. Then, there exists an NP language $\boldsymbol{L}$ such that if there exist (i) a 2-round interactive argument (P, V) for $\boldsymbol{L}$ and (ii) an oblivious black-box reduction R_{SWI} for showing the strong WI of (P, V) based on a falsifiable assumption (C, c), then the assumption (C, c) is false.*

Since Theorem 3 can be proven by closely following the proof of Lemma 2, for lack of space, we defer the proof to the full version of this paper.

BB impossibility of 2-round publicly verifiable delayed-input strong WI. By adjusting the proof of Lemma 2, we obtain the following black-box impossibility result about 2-round delayed-input publicly verifiable strong WI.

Theorem 4. *Assume the existence of CCA-secure public-key encryption schemes. Then, there exists an NP language $\boldsymbol{L}$ such that if there exist (i) a 2-round delayed-input publicly verifiable interactive argument (P, V) for $\boldsymbol{L}$ and (ii) an oblivious black-box reduction R_{SWI} for showing the delayed-input strong WI of (P, V) based on a falsifiable assumption (C, c), then the assumption (C, c) is false.*

Since Theorem 4 can be proven very similarly to Theorem 3 (as mentioned in Section 2), we omit the proof.

References

1. Badrinarayanan, S., Fernando, R., Jain, A., Khurana, D., Sahai, A.: Statistical ZAP arguments. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 642–667. Springer, Heidelberg (May 2020)

2. Badrinarayanan, S., Garg, S., Ishai, Y., Sahai, A., Wadia, A.: Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part III. LNCS, vol. 10626, pp. 275–303. Springer, Heidelberg (Dec 2017)
3. Bellare, M., Micali, S., Ostrovsky, R.: The (true) complexity of statistical zero knowledge. In: 22nd ACM STOC. pp. 494–502. ACM Press (May 1990)
4. Bitansky, N., Canetti, R., Paneth, O., Rosen, A.: On the existence of extractable one-way functions. SIAM Journal on Computing 45(5), 1910–1952 (2016)
5. Bitansky, N., Khurana, D., Paneth, O.: Weak zero-knowledge beyond the black-box barrier. In: Charikar, M., Cohen, E. (eds.) 51st ACM STOC. pp. 1091–1102. ACM Press (Jun 2019)
6. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. Journal of Computer and System Sciences 37(2), 156–189 (1988)
7. Chung, K.M., Lui, E., Mahmoody, M., Pass, R.: Unprovable security of two-message zero knowledge. Cryptology ePrint Archive, Report 2012/711 (2012), https://eprint.iacr.org/2012/711
8. Chung, K.M., Lui, E., Pass, R.: From weak to strong zero-knowledge and applications. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 66–92. Springer, Heidelberg (Mar 2015)
9. Dachman-Soled, D., Jain, A., Kalai, Y.T., Lopez-Alt, A.: On the (in)security of the Fiat-Shamir paradigm, revisited. Cryptology ePrint Archive, Report 2012/706 (2012), https://eprint.iacr.org/2012/706
10. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM Journal on Computing 30(2), 391–437 (2000)
11. Dwork, C., Naor, M.: Zaps and their applications. SIAM Journal on Computing 36(6), 1513–1543 (2007), https://doi.org/10.1137/S0097539703426817
12. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. Journal of the ACM 50(6), 852–921 (2003), https://doi.org/10.1145/950620.950623
13. Feige, U., Lapidot, D., Shamir, A.: Multiple noninteractive zero knowledge proofs under general assumptions. SIAM Journal on Computing 29(1), 1–28 (1999), https://doi.org/10.1137/S0097539792230010
14. Feige, U., Shamir, A.: Witness indistinguishable and witness hiding protocols. In: 22nd ACM STOC. pp. 416–426. ACM Press (May 1990)
15. 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 (Jun 2013)
16. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) 43rd ACM STOC. pp. 99–108. ACM Press (Jun 2011)
17. Goldreich, O.: Foundations of Cryptography: Basic Tools, vol. 1. Cambridge University Press, Cambridge, UK (2001)
18. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. Journal of Cryptology 7(1), 1–32 (Dec 1994)
19. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing 18(1), 186–208 (1989)
20. Goyal, V., Jain, A., Jin, Z., Malavolta, G.: Statistical zaps and new oblivious transfer protocols. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 668–699. Springer, Heidelberg (May 2020)
21. Green, M.D., Miers, I.: Forward secure asynchronous messaging from puncturable encryption. In: 2015 IEEE Symposium on Security and Privacy. pp. 305–320. IEEE Computer Society Press (May 2015)

22. Groth, J., Ostrovsky, R., Sahai, A.: New techniques for noninteractive zero-knowledge. Journal of the ACM 59(3) (Jun 2012), https://doi.org/10.1145/2220357.2220358
23. Haitner, I., Rosen, A., Shaltiel, R.: On the (im)possibility of Arthur-Merlin witness hiding protocols. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 220–237. Springer, Heidelberg (Mar 2009)
24. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing 28(4), 1364–1396 (1999)
25. Jain, A., Kalai, Y.T., Khurana, D., Rothblum, R.: Distinguisher-dependent simulation in two rounds and its applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 158–189. Springer, Heidelberg (Aug 2017)
26. Kalai, Y.T., Khurana, D., Sahai, A.: Statistical witness indistinguishability (and more) in two messages. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 34–65. Springer, Heidelberg (Apr / May 2018)
27. Khurana, D., Sahai, A.: How to achieve non-malleability in one or two rounds. In: Umans, C. (ed.) 58th FOCS. pp. 564–575. IEEE Computer Society Press (Oct 2017)
28. Lombardi, A., Vaikuntanathan, V., Wichs, D.: Statistical ZAPR arguments from bilinear maps. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 620–641. Springer, Heidelberg (May 2020)
29. Naor, M.: On cryptographic assumptions and challenges (invited talk). In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 96–109. Springer, Heidelberg (Aug 2003)
30. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC. pp. 427–437. ACM Press (May 1990)
31. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 160–176. Springer, Heidelberg (May 2003)
32. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) CRYPTO'91. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (Aug 1992)
33. Wichs, D.: Barriers in cryptography with weak, correlated and leaky sources. In: Kleinberg, R.D. (ed.) ITCS 2013. pp. 111–126. ACM (Jan 2013)