

Composable Adaptive Secure Protocols without Setup under Polytime Assumptions

Carmit Hazay^{1*} and Muthuramakrishnan Venkatasubramanian^{2**}

¹ Bar-Ilan University, Israel

² University of Rochester, NY

Abstract. All previous constructions of general multiparty computation protocols that are secure against adaptive corruptions in the concurrent setting either require some form of setup or non-standard assumptions. In this paper we provide the first general construction of secure multi-party computation protocol *without* any setup that guarantees composable security in the presence of an *adaptive adversary* based on standard polynomial-time assumptions. We prove security under the notion of “UC with super-polynomial helpers” introduced by Canetti et al. (FOCS 2010), which is closed under universal composition and implies “super-polynomial-time simulation”. Moreover, our construction relies on the underlying cryptographic primitives in a black-box manner.

Next, we revisit the zero-one law for two-party secure functions evaluation initiated by the work of Maji, Prabhakaran and Rosulek (CRYPTO 2010). According to this law, every two-party functionality is either trivial (meaning, such functionalities can be reduced to any other functionality) or complete (meaning, any other functionality can be reduced to these functionalities) in the Universal Composability (UC) framework. As our second contribution, assuming the existence of a simulatable public-key encryption scheme, we establish a zero-one law in the adaptive setting. Our result implies that every two-party non-reactive functionality is either trivial or complete in the UC framework in the presence of adaptive, malicious adversaries.

Keywords: UC Security, Adaptive Secure Computation, Coin-Tossing, Black-box construction, Extractable Commitments, Zero-One Law

1 Introduction

Secure computation enables a set parties to mutually run a protocol that computes some function f on their private inputs, while preserving a number of security properties. Two of the most important properties are privacy and correctness. The former implies data confidentiality, namely, nothing leaks by the protocol execution but the computed

* Email: carmit.hazay@cs.biu.ac.il. Research supported by the Israel Ministry of Science and Technology (grant No. 3-10883) and by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office.

** Email: muthuv@cs.rochester.edu. Research supported by Google Faculty Research Grant and NSF Award CNS-1526377.

output. The later requirement implies that no corrupted party or parties can cause the output to deviate from the specified function. It is by now well known how to securely compute any efficient functionality [Yao86,GMW87,MR91,Bea91,Can01] in various models and under the stringent simulation-based definitions (following the ideal/real paradigm). Security is typically proven with respect to two adversarial models: the semi-honest model (where the adversary follows the instructions of the protocol but tries to learn more than it should from the protocol transcript), and the malicious model (where the adversary follows an arbitrary polynomial-time strategy), and feasibility results are known in the presence of both types of attacks. The initial model considered for secure computation was of a static adversary where the adversary controls a subset of the parties (who are called corrupted) before the protocol begins, and this subset cannot change. In a stronger corruption model the adversary is allowed to choose which parties to corrupt throughout the protocol execution, and as a function of its view; such an adversary is called adaptive.

These feasibility results rely in most cases on stand-alone security, where a *single* set of parties run a *single* execution of the protocol. Moreover, the security of most cryptographic protocols proven in the stand-alone setting does not remain intact if many instances of the protocol are executed concurrently [Lin03]. The strongest (but also the most realistic) setting for concurrent security is known by *Universally Composable* (UC) security [Can01]. This setting considers the execution of an unbounded number of concurrent protocols in an arbitrary and adversarially controlled network environment. Unfortunately, stand-alone secure protocols typically fail to remain secure in the UC setting. In fact, without assuming some *trusted help*, UC security is impossible to achieve for most tasks [CF01,CKL06,Lin03]. Consequently, UC secure protocols have been constructed under various *trusted setup* assumptions in a long series of works; see [BCNP04,CDPW06,KLP07,CPS07,LPV09,DMRV13] for few examples.

Concurrent security *without any setup*. In many situations, having a trusted setup might be hard or expensive. Designing protocols in the plain model that provide meaningful security in a concurrent setting is thus an important challenge. In this regard, a relaxation of UC security allows the adversary in an ideal execution to run in *super-polynomial time*; this notion is referred to as super-polynomial security (or SPS) [Pas03]. On a high-level, this security notion guarantees that any attack carried out by an adversary running in polynomial time can be mounted in the ideal execution with super-polynomial resources. In many scenarios, such a guarantee is meaningful and indeed several past works have designed protocols guaranteeing this relaxed UC security against static adversaries [Pas03,BS05,LPV09] and adaptive adversaries [BS05,DMRV13,Ven14]. While initial works relied on sub-exponential hardness assumptions, more recent works in the static setting have been constructed based on standard polynomial-time hardness assumptions.

The work of [CLP10], put forth some basic desiderata regarding security notions in a concurrent setting. One of them requires supporting *modular analysis*: Namely, there should be a way to deduce security properties of the overall protocol from the security properties of its components. Quite surprisingly, it was shown in [CDPW06] that most protocols in the UC framework that consider both trusted setups and relaxed models of security, in fact, do not support this.

Towards remedying the drawbacks of SPS security, Prabhakaran and Sahai [PS04] put forth the notion of Angel-based UC security that provides guarantees analogous to SPS security while at the same time supporting modular analysis. In this model, both the adversary and the simulator have access to an oracle, referred to as an “angel” that provides judicious use of super-polynomial resources. In the same work and subsequent effort [MMY06] the authors provided constructions under this security notion relying on non-standard hardness assumptions. Recently, Canetti, Lin and Pass [CLP10] provided the first constructions in this model relying on standard polynomial time assumptions. Moreover, to emphasize the modular analysis requirement, they recast the notion of Angel-based security in the extended UC (EUC) framework of [CDPW06] calling it UC with super-polynomial helpers. While prior approaches relied on non-interactive helpers that were stateless, this work designed a helper that was highly interactive and stateful. Since this work, several follow up works [LP12a, GLP⁺15, Kiy14] have improved both the round complexity and the computational assumptions. The most recent work due to Kiyoshima [Kiy14] provides a $\tilde{O}(\log^2 n)$ -round protocol to securely realize any functionality in this framework based on semi-honest oblivious transfer protocols where the underlying primitives are used in a black-box manner. In this line of research, the work of Canetti, Lin and Pass [CLP13] distinguishes itself by designing protocols that guarantee a stronger notion of security. More precisely, they extend the angel-based security so that protocols developed in this extended framework additionally preserve security of other protocols running the system (i.e. cause minimal “side-effect”). They refer to such protocols “environment friendly” protocols. However, as observed in the same work, this strong notion inherently requires non-black-box simulation techniques. Moreover, the constructions presented in [CLP13] are non-black-box as well.

While considerable progress has been made in constructing protocols secure against static adversaries, very little is known regarding adaptive adversaries. Specifically, the work of Barak and Sahai [BS05] and subsequent works [DMRV13, Ven14] show how to achieve SPS security under non-standard assumptions. Besides these works, every other protocol that guarantees any meaningful security against adaptive adversaries in a concurrent setting has required setup. The main question left open by previous work regarding adaptive security is:

*Can we realize general functionalities with SPS security in the plain model under standard polynomial time assumptions? and,
Can we show adaptively secure angel-based (or EUC-security) under standard hardness assumptions where the underlying primitives are used in a black-box manner?*

We stress that even the works that provide SPS security require non-standard or sub-exponential hardness assumptions and are *non-black-box*, that is, the constructions rely on the underlying assumptions in non-black-box way. A more ambitious goal would be to construct “environment-friendly” protocols [CLP13] and we leave it as future work.

1.1 Our Results

In this work we resolve both these questions completely and provide the first realizations of general functionalities under EUC security against malicious, adaptive adver-

saries (See [CDPW06,CLP10] for a formal definition). More formally, we prove the following theorem:

Theorem 1.1 *Assume the existence of a simulatable public-key encryption scheme. Then there exists a sub-exponential time computable (interactive) helper machine \mathcal{H} such that for any “well formed” polynomial-time functionality \mathcal{F} , there exists a protocol that realizes \mathcal{F} with \mathcal{H} -EUC security, in the plain model secure against malicious, adaptive adversaries. Furthermore, the protocol makes only black-box use of the underlying encryption scheme.*

We recall here that simulatable public-key encryption (PKE), introduced by Damgard and Nielsen [DN00], allows to obliviously sample the public key/ciphertext without the knowledge of the corresponding secret key/plaintext.

As far as we know, this is the first construction based on polynomial-time hardness assumptions of a secure multi-party computation that achieves any non-trivial notion of concurrent security against adaptive adversaries without any trusted-set up (in the plain model) and without assuming an honest majority. Furthermore, the construction supports *modular analysis* and relies on the underlying scheme in a black-box way. In essence, our protocol provides the *strongest* possible security guarantees in the plain model.

A zero-one law for adaptive security In [PR08], Prabhakaran and Rosulek initiated the study of the “cryptographic complexity” of two-party secure computation tasks in the UC framework. Loosely speaking, in their framework a functionality \mathcal{F} *UC-reduces* to another functionality \mathcal{G} if there is a UC secure protocol for \mathcal{F} in the \mathcal{G} -hybrid, i.e., using ideal access to \mathcal{G} . Under this notion of a reduction in the presence of static adversaries, Maji et al in [MPR10] established a *zero-one* law for two-party (non-reactive) functionalities which states that every functionality is either trivial or complete. In this work, we extend their result to the adaptive setting to obtain the following theorem.

Theorem 1.2 (Informal) *All non-reactive functionalities are either trivial or complete under UC-reductions in the presence of adaptive adversaries.*

1.2 Previous Techniques

All previous approaches for Angel-based UC secure protocols relied on a particular “adaptive hardness” assumption which amounts to guaranteeing security in the presence of an adversary that has adaptive access to a helper function. Indeed, as pursued in the original approaches by [PS04,MMY06], complexity leveraging allows for designing such primitives. A major breakthrough was made by Canetti, Lin and Pass [CLP10] that showed that a helper function could be based on standard assumptions. The main technical tool introduced in this work is a new notion of a commitment scheme that is secure against an adaptive chosen commitment attack (CCA security). On a high-level, a tag-based commitment scheme, which are schemes that have additionally a tag as a common input, is said to be CCA-secure if a commitment made with tag id is hiding even if the receiver has access to a (possibly, super-polynomial time) oracle that is capable of “breaking” commitments made using any tag $id' \neq id$. In the original work, they

constructed a $O(n^\epsilon)$ -round CCA-secure commitment scheme based on one-way functions (OWFs) [CLP10]. Since then, several followup works have improved this result, culminating in the work of Kiyoshima [Kiy14] who gave a $\tilde{O}(\log^2 n)$ -round construction of a CCA-secure commitment scheme based on OWFs while relying on the underlying OWF in a black-box way.³ We remark here that Angel-based security based on standard polynomial-time assumptions have been constructed only in the static setting. Moreover, all constructions in this line of work, first construct a CCA-secure commitment scheme and then realize a complete UC functionality, such as the commitment or oblivious-transfer functionality using a “decommitment” oracle as the helper functionality.

When we consider the adaptive setting, we begin with the observation that any cryptographic primitive in use must be secure in the presence of adaptive corruptions. Saying differently, we require a simulation that can produce random coins consistent with any honest party during the execution as soon as it is adaptively corrupted. A first attempt would be to enhance a CCA-secure commitment scheme to the adaptive setting. This means there must be a mechanism to equivocate the commitment scheme. It is in fact crucial in all works using CCA-secure commitments that the helper functionality be able to break the commitment and obtain the unique value (if any) that the commitment can be decommitted to. However, equivocal commitments by definition can have commitments that do not have unique decommitments. In essence, standard CCA-secure commitment schemes are necessarily statistically binding (and all previous constructions indeed are statistically binding). Hence, it would be impossible to use any of those schemes in the adaptive setting.

Note that previous works [DMRV13, Ven14] get around this issue by relying on some sort of setup, namely, a mechanism by which the commitments will be statistically binding in the real world for adversaries, yet can be equivocated in the ideal world by the simulator. The notion of an adaptive instance-dependent scheme [LZ11] provides exactly such a primitive. Loosely speaking, such commitment schemes take additionally as input an NP-statement and provides the following guarantee: If the statement is true, the commitment can be equivocated using the witness, whereas if the statement is false then the commitment is statistically binding. Moreover, it admits adaptive corruptions where a simulator can produce random coins for an honest committer revealing a simulated commitment to any value. The work of [BS05] relies on complexity leveraging in order to generate statements that a simulator, in super-polynomial time can break but an adversary, in polynomial time, cannot break. On the other hand, the works of [DMRV13, Ven14] rely on—the so called *UC puzzle*—that provides similar advantage for the simulation while relying on milder assumptions.

A second issue arises in the adaptive setting where any commitment scheme that tolerates concurrent executions (even with fixed roles) and is equivocal, implies some sort of selective opening security. Indeed, the result of Ostrovsky et al. [ORSV13] proves that it is impossible, in general, to construct concurrent commitments secure w.r.t. selective opening attacks. Getting around this lower bound is harder. Previous results [DMRV13, Ven14] get around this lower bound by first constructing a “weaker” com-

³ We further note that Goyal et al. [GLP⁺15] gave a $\tilde{O}(\log n)$ -round CCA-secure commitment scheme but makes use of the OWF in a non-black-box way.

mitment scheme in a limited concurrent environment. Namely, they construct an equivocal non-malleable commitment scheme that can simulate any man-in-the-middle adversary receiving “left” commitments made to independent and identically distributed values (via some *a priori* fixed distribution), and is acting as a committer in many “right” interactions. This allows to get around the [ORSV13] lower bound, as Ostrovsky et al. lower bound holds only if the simulator *does not* know the distribution of the commitments received by the adversary. In any case, all previous works fail to achieve the stronger Angel-based UC security, where the helper function is provided to the adversary and the simulator in the real and ideal world respectively are the same.

Given these bottlenecks, it seems unlikely to use a commitment scheme with such a property. In this work, we introduce a new primitive that will allow to both provide the adaptive hardness property as well as admit adaptive corruptions. This primitive is coin-tossing and will additionally require to satisfy an *adaptive hardness guarantee* that we define in the next section. We chose coin-tossing as a primitive as it does not require any inputs from the parties and the output is independent of any “global” inputs of the parties participating in the coin-tossing. Roughly speaking, if a party is adaptively corrupted it is possible to sample a random string as the output and equivocate the interaction to output this string. On the other hand, a commitment scheme will not allow such a mechanism as corrupting a sender requires equivocating the interaction to a particular value (that could potentially depend on a global input).

2 Our Main Tool: CCA-Secure Coin-Tossing

The main technical tool used in our construction is a new notion of a coin-tossing protocol that is secure against adaptive chosen coins attack (CCA security). Cryptographic primitives with an adaptive hardness property has been studied extensively in the case of the encryption schemes (chosen ciphertext attack security), and more recently in the case of commitments [CLP10, KMO14, Kiy14, GLP⁺15]. We define here an analogous notion for coin-tossing protocols for the stronger case of adaptive corruptions.

A natural approach is to say that a coin-tossing protocol is CCA-secure if the coin-tossing scheme retains its simulatability even if a “Receiver” has access to a “biasing” oracle \mathcal{O} that has the power to bias the protocol outcome of the coin-tossing to any chosen value. Unfortunately, we do not know how to realize such a notion and will instead, consider a weaker “indistinguishability”-based notion (as opposed to simulation based notion) that will be sufficient for our application.

A motivating example. We motivate our definition by discussing what security properties are desirable for coin-tossing protocols (in general). Consider a public-key cryptosystem that additionally has a property that a public-key can be obliviously sampled using random coins without knowledge of the secret-key (eg, dense cryptosystems, simulatable public-key encryption schemes). Furthermore, semantic security holds for a key sampled using the oblivious strategy. Consider a protocol where the parties after engaging in a coin-toss protocol sample a public-key using the outcome of the coin-toss. In such a scenario we would like the coin-tossing scheme to ensure that the semantic-security continues to hold if parties encrypt messages using the public-key.

The natural “simulatable” definition requires the coin-toss to be “simulatable”. If we instantiate a simulatable coin-toss protocol in our motivating application, semantic

security of ciphertexts constructed using the public-key sampled from the coin-toss outcome indeed holds via a simple security reduction. Suppose there exists an adversary that distinguishes an encryption of 0 from 1 when encrypted under a public-key sampled using the coin-toss. We can use the simulator to construct an adversary that violates the security of the underlying encryption scheme. Consider a simulator that receives as a challenge a uniformly sampled string and a ciphertext generated with the associated public-key. The simulator can internally simulate the coin-tossing to be this sampled string and thereby use the adversary to break the security of the encryption scheme.

A weaker alternative to simulatability is an information-theoretic based definition where the requirement would be that the entropy of the outcome is sufficiently high. However, such a definition will not suffice in our motivating example.⁴ This is because we will not be able to “efficiently” reduce a cheating adversary to the violating the security game of the underlying cryptosystem.

Instead, we take a more direct approach where the security for the coin-toss is defined so that it will be useful in our motivating example. First, we generalize the security game of the underlying encryption scheme in our motivating example to any indistinguishability based primitive. We model such a primitive via a (possibly) interactive challenger \mathcal{C} that receives as input a random string o and a private bit b . We say that an adversary interacting with \mathcal{C} succeeds if when interacting on a randomly chosen o and bit b , the adversary can guess b with probability better than a $\frac{1}{2}$. Let π be a (two-party) coin-toss protocol. Our motivating example can be formulated using the following experiment EXP_b with an adversary \mathcal{A} :

- \mathcal{A} interacts with an honest party using π to generate o .
- Next, it interacts with a challenger \mathcal{C} on input o and bit b .

We compare this experiment with a stand-alone experiment STA_b where an adversary \mathcal{B} simply interacts with \mathcal{C} on input b and o where o is uniformly sampled. Our security definition of the coin-tossing protocol must preserve the following security property against a challenger \mathcal{C} : *if the stand-alone game is hard to distinguish, i.e. STA_0 from STA_1 , then the experiments EXP_0 from EXP_1 must also be hard to distinguish.* More formally, our definition will (explicitly) give a reduction from any adversary that \mathcal{A} distinguishes EXP_b to a stand-alone adversary \mathcal{B} that can distinguish STA_b . Finally, in a CCA-setting, we generalize this definition by requiring that if there exists any oracle adversary $\mathcal{A}^\mathcal{O}$ with access to a biasing oracle \mathcal{O} that can distinguish EXP_0 from EXP_1 , then there exists a stand-alone adversary \mathcal{B} (without access to any oracle) that can distinguish STA_b from STA_1 .

Towards formalizing this notion and incorporating adaptive corruptions, we first consider a tag-based coin-tossing protocol between two parties, an *Initiator* I and a *Receiver* R with $l(n)$ -bit identities and $m(n)$ -bit outcomes. A biasing oracle \mathcal{O} interacts with an adversary \mathcal{A} as follows: \mathcal{O} participates with \mathcal{A} in many sessions using the protocol where the oracle controls the initiator, using identities of length $l(n)$ that are chosen adaptively by \mathcal{A} . At the beginning of each session, the adversary produces a coin

⁴ Unless the cryptosystems have additional properties. For instance, consider dual-mode encryption schemes where there are keys sampled via a high-entropy string and could potentially be statistically hiding.

outcome $c \in \{0, 1\}^{m(n)}$ to the oracle where at the end of this session, if the initiator that is initially controlled by the oracle is not (adaptively) corrupted by the adversary, then the outcome of the interaction must result in the *chosen coin* c . If at any point during the interaction the initiator is corrupted, then the oracle simply provides the random-tape of I that is consistent with the partial transcript of the interaction.

We compare an experiment EXP_b with oracle PPT adversary $\mathcal{A}^{\mathcal{O}}$ and a stand-alone experiment STA_b with adversary \mathcal{B} . In the man-in-the-middle experiment, an adversary with oracle access to \mathcal{O} interacts with a honest receiver R on identity id to generate an output $o \in \{0, 1\}^n$ where n is the security parameter. Then it interacts with a challenger \mathcal{C} on common input (n, o, id) and private input b for \mathcal{C} . The adversary is allowed to corrupt the receiver R , challenger \mathcal{C} and any of the interactions with \mathcal{O} . If the adversary \mathcal{A} corrupts either \mathcal{C} or I then the output of the experiment is set to \perp . If for some identity id' on which \mathcal{A} queries \mathcal{O} , it holds that $\text{id}' = \text{id}$, then the output of the experiment is set to \perp . Otherwise, the output of the experiment is set to be the output of the adversary.

In the stand-alone experiment STA_b , we consider a PPT adversary \mathcal{B} that interacts with \mathcal{C} on common input (n, o) and private input b for \mathcal{C} where o is uniformly sampled from $\{0, 1\}^n$. The output of the experiment is set to be the output of \mathcal{B} . Observe that in the stand-alone experiment \mathcal{B} does not get to corrupt \mathcal{C} .

Informally, a tag-based coin-tossing scheme $\langle I, R \rangle$ is said to be CCA-secure against a challenger \mathcal{C} , if there exists a biasing oracle \mathcal{O} for $\langle I, R \rangle$ such that for every oracle PPT adversary \mathcal{A} and distinguisher \mathcal{D} such that \mathcal{D} distinguishes EXP_0 and EXP_1 with \mathcal{A} , then there exist a (stand-alone) PPT \mathcal{B} and distinguisher \mathcal{D}' such that \mathcal{D}' distinguishes STA_0 and STA_1 with \mathcal{B} .

In addition to this security requirement we will additionally consider the following definition of CCA-security which simply requires that any adversary with oracle access to a biasing oracle \mathcal{O} can be simulated by a stand-alone PPT machine. In this case, we simply say $\langle I, R \rangle$ is CCA-secure w.r.t \mathcal{O} .

Quite surprisingly, we show how to realize such a primitive by relying on a CCA-secure commitment that is secure only against static adversaries. The idea here is that while CCA-secure commitments cannot admit adaptive corruptions, the basic security game ensures that an *unopened* commitment remains hiding in the presence of an adversary having access to a decommitment oracle. We combine such a commitment scheme with the technique of Hazay and Venkatasubramanian from [HV15] who showed how to construct an adaptive UC commitment scheme, starting from a public-key encryption scheme (with an oblivious ciphertext generation property) in the CRS model. On a high-level, the protocol can be abstracted as providing a transformation from an extractable (only) commitment scheme (that has an oblivious generation property) to a full adaptively secure UC-commitment. At first, it would be tempting to simply replace the invocations of extractable commitments with a CCA-secure commitment scheme as we only require extraction from these commitments and not equivocation in the simulation. However, this intuition fails in an adaptive setting when considering the fact that we additionally require that the commitment scheme has an oblivious generation property and it is unclear how to construct such an extractable scheme (based on rewinding) to have this property. Nevertheless, we show how to carefully use CCA-secure commitments in the same protocol to obtain a CCA-secure coin-tossing scheme. Next, we show that

given a CCA-secure coin-tossing protocol with a biasing oracle \mathcal{O} it is possible to realize the ideal commitment functionality using a helper functionality. Again, we use another variant of the same protocol from [HV15] to accomplish this transformation. Our constructions and proofs of security are highly modular and quite simple. Moreover, all our transformations rely on the underlying primitives in a black-box manner.

Finally, we show that the black-box construction of an $O(n^\epsilon)$ -round CCA-secure commitment scheme from Lin and Pass [LP12a] will satisfy the required property to be instantiated in our protocol for the CCA-secure coin-tossing scheme.

We remark here that while the focus of the present work is to achieve *plain* angel-based security, we could achieve the stronger “environment-friendly” property if we instead rely on a *strongly unprovable* CCA-secure commitment scheme [CLP13] to construct our CCA-secure coin-tossing scheme. We leave this as future work.

2.1 A Formal Definition of CCA-Secure Coin-Tossing

We begin with the simpler security requirement of CCA-security w.r.t biasing oracles.

Definition 1 (CCA-secure coin-tossing) *Let $\langle I, R \rangle$ be a tag-based coin-tossing scheme with $l(n)$ -bit identities, $m(n)$ -bit outcomes and \mathcal{O} a biasing oracle for it. We say that $\langle I, R \rangle$ is robust CCA-secure w.r.t. \mathcal{O} , if for every PPT adversary \mathcal{A} there exists a simulator \mathcal{S} such that the following distributions are indistinguishable.*

$$(i) \{\mathcal{A}^{\mathcal{O}}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*} \quad (ii) \{\mathcal{S}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

2.2 CCA-Security w.r.t Challengers

Let the random variable $\text{EXP}_b(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z)$ denote the output of the following experiment:

1. On common input 1^n and auxiliary input z , $\mathcal{A}^{\mathcal{O}}$ chooses an identity $\text{id} \in \{0,1\}^{l(n)}$ and first interacts with a honest receiver R using $\langle I, R \rangle$. Let o be the outcome of the execution.
2. Next, it interacts with \mathcal{C} with common input (n, o) and private input b for \mathcal{C} .

Finally, the experiment outputs the view of the adversary \mathcal{A} in the experiment and the output is set to \perp unless \mathcal{A} corrupts either \mathcal{C} or I or any of the identities chosen for the interactions of \mathcal{A} with \mathcal{O} is equal to id . Let the random variable $\text{STA}_b(\mathcal{B}, \mathcal{C}, n, z)$ denote the output of \mathcal{B} in an interaction between \mathcal{B} and \mathcal{C} with common input (n, o) where o is uniformly sampled from $\{0,1\}^n$, private input b for \mathcal{C} and auxiliary input c with \mathcal{B} .

Definition 2 (CCA-secure coin-tossing) *Let $\langle I, R \rangle$ be a tag-based coin-tossing scheme with $l(n)$ -bit identities, $m(n)$ -bit outcomes and \mathcal{O} a biasing oracle for it. We say that $\langle I, R \rangle$ is CCA-secure w.r.t. \mathcal{O} against a pair of challengers $(\mathcal{C}_0, \mathcal{C}_1)$, if for every PPT adversary \mathcal{A} and distinguisher \mathcal{D} , if \mathcal{D} distinguishes the following ensembles with non-negligible probability:*

$$\{\text{EXP}_0(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}_0, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}, \quad \{\text{EXP}_1(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}_1, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

then there exists a stand-alone adversary (that does not have access to \mathcal{O}) \mathcal{B} and distinguisher \mathcal{D}' such that \mathcal{D}' distinguishes the following ensembles with non-negligible probability:

$$(i)\{\text{STA}_0(\mathcal{B}, \mathcal{C}, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}, \quad (ii)\{\text{STA}_1(\mathcal{B}, \mathcal{C}, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

We highlight that in a real experiment, o is the result of the outcome of a coin-tossing between the adversary acting as the receiver and an honest initiator. However, the game between \mathcal{B} and \mathcal{C}_b is instantiated with a randomly chosen o . In essence, the definition says that if a challenge presented by \mathcal{C}_0 and \mathcal{C}_1 is hard to distinguish for a randomly sampled o , then it will be hard to distinguish even if o was sampled according to $\langle I, R \rangle$ with an adversarial receiver R who has access to oracle \mathcal{O} .

3 Preliminaries

We assume familiarity with basic notions of Turing machines, probabilistic-polynomial time computation and standard security notions of computational indistinguishability, public-key encryption and commitment schemes.

3.1 Simulatable PKE

Definition 3 (Simulatable public-key encryption scheme) A ℓ -bit simulatable encryption scheme consists of an encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ augmented with $(\text{oGen}, \text{oRndEnc}, \text{rGen}, \text{rRndEnc})$. Here, oGen and oRndEnc are the oblivious sampling algorithms for public keys and ciphertexts, and rGen and rRndEnc are the respective inverting algorithms, rGen (resp. rRndEnc) takes r_G (resp. (PK, r_E, m)) as the trapdoor information. We require that, for all messages $m \in \{0, 1\}^\ell$, the following distributions are computationally indistinguishable:

$$\{\text{rGen}(\text{PK}), \text{rRndEnc}(\text{PK}, c), \text{PK}, c \mid (\text{PK}, \text{SK}) = \text{Gen}(1^n; r_G), c = \text{Enc}_{\text{PK}}(m; r_E)\} \\ \text{and } \{\hat{r}_G, \hat{r}_E, \hat{\text{PK}}, \hat{c} \mid (\hat{\text{PK}}, \perp) = \text{oGen}(1^n; \hat{r}_G), \hat{c} = \text{oRndEnc}_{\hat{\text{PK}}}(1^n; \hat{r}_E)\}$$

It follows from above that a simulatable encryption scheme is also semantically secure.

3.2 CCA-Secure Commitment Schemes

The following is taken verbatim from [CLP10]. Roughly speaking, a commitment scheme is CCA (chosen-commitment-attack) secure if the commitment scheme retains its hiding property even if the receiver has access to a “decommitment oracle”. Let $\langle C, R \rangle$ be a tag-based commitment scheme with $l(n)$ -bit identities. A decommitment oracle \mathcal{O} of $\langle C, R \rangle$ acts as follows in interaction with an adversary \mathcal{A} : it participates with \mathcal{A} in many sessions of the commit phase of $\langle C, R \rangle$ as an honest receiver, using identities of length n , chosen adaptively by \mathcal{A} . At the end of each session, if the session is accepting and valid, it reveals a decommitment of that session to \mathcal{A} . Otherwise, it sends \perp . Note that when a session has multiple decommitments, the decommitment oracle only

returns one of them. Hence, there might exist many valid decommitment oracles. We remark that we will rely on a slightly weaker oracle, referred to as “committed-value” oracle in [LP12a] that simply extracts the committed value instead of providing the decommitment information. This relaxation is required for the black-box construction in [LP12a] and we will rely on the same definition.

Loosely speaking, a tag-based commitment scheme $\langle C, R \rangle$ is said to be CCA-secure, if there exists a committed-value oracle \mathcal{O} for $\langle C, R \rangle$, such that the hiding property of the commitment holds even with respect to adversaries with access to \mathcal{O} . More precisely, let $\mathcal{A}^{\mathcal{O}}$ denote the adversary \mathcal{A} with access to the oracle \mathcal{O} . Let $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)$, where $b \in \{0, 1\}$, denote the output of the following probabilistic experiment: on common input 1^n and auxiliary input z , $\mathcal{A}^{\mathcal{O}}$ (adaptively) chooses a pair of challenge values $(v_0, v_1) \in \{0, 1\}$, the values to be committed to, and an identity $\text{id} \in \{0, 1\}^{l(n)}$, and receives a commitment to v_b using identity id . Finally, the experiment outputs the output y of $\mathcal{A}^{\mathcal{O}}$, the output y is replaced by \perp if during the execution \mathcal{A} sends \mathcal{O} any commitment using identity id (that is, any execution where the adversary queries the committed-value oracle on a commitment using the same identity as the commitment it receives, is considered invalid).

Definition 4 (CCA-secure commitments) *Let $\langle C, R \rangle$ be a tag-based commitment scheme with $l(n)$ -bit identities, and \mathcal{O} a committed-value oracle for it. We say that $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O} , if for every PPT \mathcal{A} , the following ensembles are computationally indistinguishable:*

$$(i) \{\text{IND}_0(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)\}_{n \in \mathbb{N}}, \quad (ii) \{\text{IND}_1(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)\}_{n \in \mathbb{N}}$$

We say that $\langle C, R \rangle$ is CCA-secure if there exists a committed-value oracle \mathcal{O}' , such that, $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O}' .

We extend this definition to include adversaries that can *adaptively* corrupt the committer C in the left interaction and any of the receivers in the interactions with the committed-value oracle. We present this definition in Appendix A. We stress here that the security definition only requires the standard static guarantee of hiding even in the presence of adaptive corruptions. Finally, we will also require a strengthening of the CCA-security commitment scheme called k -robustness [CLP10] that preserves the security of arbitrary k -round protocols w.r.t any adversary that has access to the committed-value oracle and its adaptive analogue (For a more precise definition, we refer the reader to the full version).

4 Black-Box Adaptive UC Secure Protocols with Super-Polynomial Helpers

We consider the model of UC with super-polynomial helpers introduced in [PS04, CLP10]. Informally speaking, in this UC model, both the adversary and the environment in the real and ideal worlds have access to a super-polynomial time functionality that assists the parties. For more details, we refer the reader to [CLP10]. In the original work of [CLP10] as well as subsequent works, only static adversaries were considered. In this work, we consider the stronger adaptive adversary and obtain the following theorem in this model.

Theorem 4.1 *Assume the existence of a simulatable public-key encryption scheme. Then, for every $\epsilon > 0$ there exists a super-polynomial time helper functionality \mathcal{H} , such that for every well-formed functionality \mathcal{F} , there exists a $\tilde{O}(d_{\mathcal{F}}n^\epsilon)$ -round protocol Π that \mathcal{H} -EUC emulates \mathcal{F} where $d_{\mathcal{F}}$ is the depth of the circuit implementing the functionality \mathcal{F} . Furthermore, the protocol uses the underlying encryption scheme in a black-box way.*

We will rely in our proof the following two lemmas.

Lemma 4.1 *Assume the existence of a simulatable public-key encryption scheme and a T_{COIN} -round CCA-secure coin-tossing protocol. Then, there exists a super-polynomial time helper functionality \mathcal{H} , such that there exists a $O(T_{\text{COIN}})$ -round protocol Π that \mathcal{H} -EUC emulates \mathcal{F}_{COM} against malicious adaptive adversaries. Furthermore, the protocol uses the underlying encryption scheme in a black-box way.*

Lemma 4.2 *Assume the existence of one-way functions, the for every $\epsilon > 0$ there exists a $O(n^\epsilon)$ -round CCA-secure coin-tossing scheme against malicious adaptive adversaries. Furthermore, the protocol uses the underlying primitives in a black-box way.*

First, we prove the theorem assuming the lemmas hold and then prove the lemmas in the following sections. Towards this, we first describe our helper functionality \mathcal{H} . The biasing oracle for the CCA-secure coin-tossing scheme provided in Lemma 4.2 will serve as \mathcal{H} . This in turn relies on Lin and Pass construction from [LP12a] of a $\tilde{O}(n^\epsilon)$ -round black-box construction of a CCA-secure commitment scheme based on one-way functions. Since one-way functions can be constructed from a simulatable public-key encryption scheme in a black-box way, combining [LP12a] with Lemmas 4.1 and 4.2 we have a $O(n^\epsilon)$ -round protocol that \mathcal{H} -EUC that emulates \mathcal{F}_{COM} . We conclude the proof of the theorem by combining the following three results:

1. The work of Choi et al. [CDMW09] provides a $O(T_{\text{OT}})$ -round construction that realizes \mathcal{F}_{OT} in the \mathcal{F}_{COM} -hybrid assuming the existence of a T_{OT} -round stand-alone adaptively-secure semi-honest oblivious-transfer protocol where the underlying protocol is used in a black-box way.
2. The work of Damgard and Nielsen [DN00] provides a black-box construction of a $O(1)$ -round stand-alone adaptively-secure semi-honest oblivious-transfer protocol assuming the existence of simulatable public-key encryption schemes.
3. The work of Ishai et al. [IPS08] provides a $O(d_{\mathcal{F}})$ -round protocol that realizes any well-formed functionality \mathcal{F} in the \mathcal{F}_{OT} -hybrid, where $d_{\mathcal{F}}$ is the depth of the circuit implementing functionality \mathcal{F} .

We rely on the $O(n^\epsilon)$ construction of CCA-secure commitment of Lin and Pass [LP12a] instead of the more round efficient construction of Kiyoshima [Kiy14] because we additionally need to prove that the commitment is secure in the presence of adaptive adversaries and we are able to achieve this only for the [LP12a] construction. We leave it as future work to improve it with respect to the [Kiy14] construction.

5 CCA-Secure Coin-Tossing from CCA-Secure Commitments

In this section, we provide our construction of CCA-secure coin-tossing protocol. The two primitives we will require are CCA-secure commitments and one-way functions. Recall that, standard CCA-secure commitments require that a value committed to, using a tag id , remains hidden even to an adversary who has access to a “decommitment oracle”. We will additionally require that if we consider an adversary that can adaptively corrupt receivers in its interactions with the decommitment oracle, the value committed to the adversary is hidden as long as the committer in this interaction is not corrupted. It is easy to show that standard CCA-secure commitments in the static setting satisfy this property. We discuss this at the end of this section.

Let $\langle C, R \rangle$ be a CCA-secure commitment scheme and Com be a statistically-binding commitment scheme with pseudorandom commitments. The 2-round commitment scheme of Naor [Nao91] based on one-way function satisfies this notion. Next, we prove the scheme from Figure 1 is CCA-secure and CCA-secure against challengers.

Theorem 5.1 *Suppose, $\langle C, R \rangle$ is a 0-robust CCA secure commitment scheme in the presence of adaptive adversaries. Then there exists an oracle helper \mathcal{O} such that $\langle I, R \rangle$ is a CCA-secure coin tossing protocol w.r.t \mathcal{O} .*

Proof. To demonstrate our scheme is CCA-secure, we construct a biasing oracle \mathcal{O} and show that given any PPT adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} such that:

$$\{\mathcal{A}^{\mathcal{O}}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\mathcal{S}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

We provide the description of our biasing oracle \mathcal{O} in Figure 2. On a high-level, this oracle follows the equivocation strategy analogous to the simulation in [HV15]. In slight more detail, this protocol that is a variant of the protocol in [HV15] allows for the initiator to equivocate m in Stage 3 if for a chosen set S at the beginning of the execution, the outcome of the coin-toss in Stage 2 can be biased to yield S . Our oracle \mathcal{O} will be able to accomplish this by breaking the commitment made by the receiver R in Stage 2 using $\langle C, R \rangle$ in exponential time.

Next, given an adversary \mathcal{A} , we construct a simulator \mathcal{S} . We do this in two steps:

Step 1: Suppose \mathcal{O}' is the oracle w.r.t which $\langle C, R \rangle$ is 0-robust. From the description of our oracle \mathcal{O} , it follows that every query to \mathcal{O} can be simulated by a PPT algorithm with access to \mathcal{O}' . Recall that the only super-polynomial computation made by \mathcal{O} is breaking a commitment made using $\langle C, R \rangle$, which can be done using \mathcal{O}' .⁵

Therefore, given any adversary \mathcal{A} , there exists another oracle adversary $\hat{\mathcal{A}}$ such that the following distributions are identically distributed:

$$\{\mathcal{A}^{\mathcal{O}}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\hat{\mathcal{A}}^{\mathcal{O}'}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

Step 2: Relying on the 0-robustness CCA-security of the $\langle C, R \rangle$ commitment scheme, it follows that given $\hat{\mathcal{A}}$, there exists a simulator \mathcal{S} such that the following distributions are indistinguishable. $\{\hat{\mathcal{A}}^{\mathcal{O}'}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\mathcal{S}(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$

⁵ We remark here that typical CCA-secure commitment schemes are statistically binding and such schemes can be easily broken in exponential time. However, the CCA-secure commitment of [LP12a] is not statistically binding. Yet, as shown in [LP12a] it is “strongly” computationally binding which will suffice.

Protocol $\pi_{\text{COIN}} = \langle I, R \rangle$.

Let 1^n be the common input to the initiator I and receiver R and the identity of the interaction $\text{id} \in \{0, 1\}^{l(n)}$.

Stage 1: Commit Phase: The receiver sends the first message σ of the Naor's commitment scheme. The initiator first picks a random bit m and chooses a random n -degree polynomial $p(\cdot)$ over a field $\mathbb{F}[x]$ such that $p(0) = m$. Namely, it randomly chooses $a_i \leftarrow \mathbb{F}$ for all $i \in [n]$ and sets $a_0 = m$, and defines the polynomial $p(x) = a_0 + a_1x + \dots + a_nx^n$. The initiator then creates a commitment to m as follows. For every $i \in [3n + 1]$, it first picks $b_i \leftarrow \{0, 1\}$ at random and then computes:

$$c_i^{b_i} = \text{Com}_\sigma(p(i); t_i) \text{ and } c_i^{1-b_i} = r_i$$

where $r_i, t_i \leftarrow \{0, 1\}^n$. The initiator sends $(c_0^0, c_0^1), \dots, (c_{3n+1}^0, c_{3n+1}^1)$ to the receiver.

Stage 2: Cut-and-Choose Phase: The initiator and receiver interact in a coin-tossing protocol to obtain the cut-and-choose set that is carried out as follows.

1. The receiver chooses a random σ_0 and commits to the initiator using $\langle C, R \rangle$ using identity id .
2. The initiator picks $\sigma_1 \leftarrow \{0, 1\}^N$ at random and sends it in the clear to the receiver.
3. The receiver decommits r_{σ_0} .

Both the initiator and the receiver compute $\sigma = \sigma_0 \oplus \sigma_1$ and use σ as the random string to sample a random subset $\Gamma \subset [3n + 1]$ of size n . The initiator provides the decommitments for $\{c_i^{b_i}\}_{i \in \Gamma}$ by sending the sequence $\{b_i, p(i), t_i\}_{i \in \Gamma}$. The receiver verifies that all the decommitments are correct and aborts otherwise.

Stage 3: Coin-Toss Phase: In the first two stages, the initiator essentially commits to the string m . Next they continue with the coin-tossing protocol.

1. The receiver commits to m' using $\langle C, R \rangle$.
2. This is followed by the initiator revealing its input m as follows: Let $\Gamma' = [3n + 1] - \Gamma$. The initiator decommits $\{c_i^{b_i}\}_{i \in \Gamma'}$ to their respective messages. The receiver checks if the decommitments are correct and aborts otherwise. Using the n polynomial evaluations revealed relative to $i \in \Gamma$ and any additional polynomial evaluation that was revealed relative to Γ' , the receiver reconstructs the polynomial $p(\cdot)$ (via polynomial interpolation of $n + 1$ points). Next, the receiver verifies whether $p(0) = m$, and that for every $i \in [3n + 1]$ the point $p(i)$ is the decrypted value within $c_i^{m_i}$.
3. The receiver decommits m' .

Finally, the outcome of the coin-tossing is $m \oplus m'$. More formally, $\text{out}(\tau)$ where τ is the transcript of this protocol is set to $m \oplus m'$.

Fig. 1. Our CCA-secure coin-tossing protocol $\langle I, R \rangle$.

The statement of the theorem now follows using a standard hybrid argument. ■

Next, we proceed to show the stronger security-preserving property of our scheme.

Description of biasing oracle \mathcal{O} : For the protocol described in Figure 1, our biasing oracle \mathcal{O} on input c proceeds as follows:

- In Stage 1, picks a random subset $\tilde{T} \subset [3n + 1]$ of size n and two random n -degree polynomials $p_0(\cdot)$ and $p_1(\cdot)$ such that p_0 and p_1 agree on all points $i \in \tilde{T}$ and $p_0(0) = 0$ and $p_1(0) = 1$.

- For every $i \in \tilde{T}$ the simulator proceeds as the honest sender would with polynomial $p_0(\cdot)$. Namely, it first picks $b_i \leftarrow \{0, 1\}$ at random and then sets

$$c_i^{b_i} = \text{Com}(p_0(i); t_i) \text{ and } c_i^{1-b_i} = r_i$$

where $t_i \leftarrow \{0, 1\}^n$ (we recall that $p_0(i) = p_1(i)$ for all $i \in \tilde{T}$).

- For every $i \in \tilde{T}' = [3n + 1] - \tilde{T}$, the simulator picks $b_i \leftarrow \{0, 1\}$ at random and then sets

$$c_i^{b_i} = \text{Com}(p_0(i); t_i^0) \text{ and } c_i^{1-b_i} = \text{Com}(p_1(i); t_i^1)$$

where $t_i^0, t_i^1 \leftarrow \{0, 1\}^n$ are chosen uniformly at random.

Finally, it sends $(c_0^0, c_0^1), \dots, (c_{3n+1}^0, c_{3n+1}^1)$ as the Stage 1 message of the initiator.

- In Stage 2, it breaks the commitment made using $\langle C, R \rangle$ and obtains the decommitted value σ_0 . Next, it sets σ_1 so that $\sigma = \sigma_1 \oplus \sigma_0$ yields the set \tilde{T} as the outcome in Stage 2.

- In Stage 3, it breaks the commitment made using $\langle C, R \rangle$ and obtains m' . Then it decommits to $m = c \oplus m'$ using the following strategy. Recall first as the initiator it needs to reveal points on a polynomial $p(\cdot)$ and pairs $\{(b_i, t_i)\}_{i \in [3n+1]}$ such that $p(0) = m$ and $c_i^{b_i} = \text{Com}(p(i); t_i)$. Let $\hat{b}_i = b_i \oplus m$ for all $i \in \tilde{T}'$, then \mathcal{S} reveals $p_m(\cdot)$, $\{\hat{b}_i, t_i^{\hat{b}_i}, r_i = c_i^{1-m}\}_{i \in \tilde{T}'}$.

Fig. 2. Biasing oracle \mathcal{O} .

Theorem 5.2 *Suppose, $\langle C, R \rangle$ is a k -robust CCA-secure commitment scheme in the presence of adaptive adversaries. Then for every pair of k -message PPT $(\mathcal{C}_0, \mathcal{C}_1)$, $\langle I, R \rangle$ is a CCA-secure coin-tossing scheme w.r.t. the biasing oracle \mathcal{O} against a pair of challengers $(\mathcal{C}_0, \mathcal{C}_1)$.*

Proof. Assume for contradiction there exist an adversary \mathcal{A} , sequence $\{z_n\}_{n \in \mathbb{N}}$ and distinguisher \mathcal{D} such that \mathcal{D} distinguishes the following ensembles

$$\{\text{EXP}_0(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z_n)\}_{n \in \mathbb{N}}, \quad \{\text{EXP}_1(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z_n)\}_{n \in \mathbb{N}}$$

with non-negligible probability. Namely, it distinguishes with probability $p(n)$ for some polynomial $p(\cdot)$ and infinitely many n 's. We need to construct a machine \mathcal{B} and distinguisher \mathcal{D}' that will distinguish STA_0 from STA_1 . Let \mathcal{O}' be the committed-value oracle guaranteed by the 0-robust CCA-security of $\langle C, R \rangle$ in the presence of an adaptive adversary. We will accomplish our goal of constructing \mathcal{B} in two steps.

Step 1: First we construct a simulator $\tilde{\mathcal{S}}$ such that the following distributions are distinguishable with non-negligible probability.

$$(1)\{\text{STA}_0(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}, (2)\{\text{STA}_1(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$$

Step 2: Since \mathcal{C} interacts in at most k -messages, we obtain the required \mathcal{B} directly by relying on the k -robustness of the CCA-security of $\langle C, R \rangle$ in the presence of an adaptive adversary.

Step 1: Constructing $\tilde{\mathcal{S}}^{\mathcal{O}'}$. Fix an n for which \mathcal{D} distinguishes the two ensembles with probability $p = p(n)$. Recall that in the EXP experiment, \mathcal{A} first interacts with an external R and then interacts with \mathcal{C}_b .

In a random instance of the EXP_b experiment, let T be the random variable representing the partial transcript up until the end of Stage 1 in \mathcal{A} 's interaction with external R . Now, we consider a slightly modified experiment $\widetilde{\text{EXP}}_b^T$ which proceeds identically to the EXP_b experiment starting from the prefix T .

Now, using an averaging argument, we can conclude that with probability at least $p/2$ over partial transcript $\tau_n \leftarrow T$ it holds that \mathcal{D} distinguishes the following two ensembles with probability at least $p/2$.

$$\{\widetilde{\text{EXP}}_0^{\tau_n}(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}_0, n, z_n)\}_{n \in \mathbb{N}}, \{\widetilde{\text{EXP}}_0^{\tau_n}(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}_1, n, z_n)\}_{n \in \mathbb{N}}$$

Now, we are ready to construct $\tilde{\mathcal{S}}$. The high-level approach is as follows:

- First, we show that, except with non-negligible probability over random executions starting from τ_n , there is a fixed value m_n that the adversary will decommit to in the Stage 3 of its interaction with R . We will rely on an information theoretic lemma from [HV15] for this. We state this step in the Claim 5.1 below.

Claim 5.1 *There exists a string m_n such that, starting from partial transcript τ_n , the probability that \mathcal{A} successfully decommits to a message different from m_n in Stage 3 is negligible.*

On a high-level the idea is that given the transcript until end of Stage 1, there is a unique set S that needs to be the outcome of Stage 2 in order for the an initiator to equivocate in Stage 3. We can show that if an adversarial initiator can equivocate with non-negligible probability to bias the coin-toss in Stage 2 to yield this unique set S , then it violates the CCA-security of the commitment made using $\langle C, R \rangle$ in Stage 2. We provide a formal proof of the claim at the end of this section.

- Next, for a fixed transcript τ_n , we will give τ_n, m_n and partial view of \mathcal{A} in the execution as the non-uniform advice. Our simulator $\tilde{\mathcal{S}}$ will start an execution with \mathcal{A} from the partial view with transcript τ_n and will use m_n to bias the outcome of the coin-toss to o by setting $m' = m_n \oplus o$ in Stage 3 of the execution. Now, we observe that, if o is uniformly distributed, then m' chosen by $\tilde{\mathcal{S}}$ will also be (non-negligibly) close to the uniform distribution given m_n and hence the view of \mathcal{S} output with \mathcal{C}_b will be statistically close to the distribution of \mathcal{A} when interacting with \mathcal{C}_b starting from τ_n . This means that if \mathcal{D} distinguishes the view of \mathcal{A} starting from τ_n in both the experiments, then it will also distinguish the output of $\tilde{\mathcal{S}}^{\mathcal{O}'}$ in the two experiments.

We now construct our simulation $\tilde{\mathcal{S}}$. On input $(1^n, o, (z, \tau_n, m_n, r_n))$, $\tilde{\mathcal{S}}^{\mathcal{O}'}$ internally emulates an execution of $\mathcal{A}(1^n, z; r)$ in the real experiment starting from the partial transcript τ_n . On the left, $\tilde{\mathcal{S}}^{\mathcal{O}'}$ needs to provide messages for the initiator I such that the outcome is o while simultaneously answering all oracle queries to \mathcal{O} . This it accomplishes by committing to $m' = o \oplus m_n$ in Stage 3. Then if the adversary reveals anything other than the m_n , it simply aborts

Answering \mathcal{O} queries. In any interaction, the oracle \mathcal{O} first receives a coin c . In the internal emulation $\tilde{\mathcal{S}}^{\mathcal{O}'}$ obtains c and needs to emulate \mathcal{O} . It carries out the actions exactly as \mathcal{O} with the exception that instead of breaking the commitments made using $\langle C, R \rangle$ (as \mathcal{O} does) $\tilde{\mathcal{S}}^{\mathcal{O}'}$ simply forwards it to \mathcal{O}' which breaks them for \mathcal{S} .

It follows from the construction and Claim 5.1 that the following distributions are statistically close:

- $\{\widetilde{\text{EXP}}_b^{\tau_n}(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z_n)\}_{n \in \mathbb{N}}$ and
- $\{\text{STA}_b(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$

and therefore \mathcal{D} distinguishes the distribution $\text{STA}_0(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))$ from $\text{STA}_1(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))$ with probability at least $p/2 - \nu(n) > p/4$ for all sufficiently large n 's.

Step 2: Constructing a stand-alone \mathcal{B} . In Step 1, we constructed a machine $\tilde{\mathcal{S}}^{\mathcal{O}'}$ that with access to \mathcal{O}' can violate the game. Now to get a stand-alone \mathcal{B} , we simply invoke k -robustness security of $\langle C, R \rangle$ with $\tilde{\mathcal{S}}^{\mathcal{O}'}$ and obtain \mathcal{B} . More precisely, using the robustness we have that the following distributions are computationally indistinguishable:

- $\{\text{STA}_b(\tilde{\mathcal{S}}^{\mathcal{O}'}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$
- $\{\text{STA}_b(\mathcal{B}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$

and therefore \mathcal{D} distinguishes the distribution $\text{STA}_0(\mathcal{B}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))$ from the distribution $\text{STA}_1(\mathcal{B}, \mathcal{C}, n, (z_n, \tau_n, m_n, r_n))$ with probability at least $p/4 - \nu(n) > p/8$ for all sufficiently large n 's and this completes the proof of the theorem.

To conclude the proof of Theorem 5.2, it only remains to prove Claim 5.1.

Proof of Claim 5.1 Assume for contradiction, the adversary \mathcal{A} equivocates with non-negligible probability starting from τ_n . We now show that $\mathcal{A}^{\mathcal{O}'}$ violates the CCA-security of $\langle C, R \rangle$ w.r.t \mathcal{O}' , namely, it violates the hiding property of the commitment made using $\langle C, R \rangle$ in Stage 2.

As stated above, we use an information theoretic lemma from [HV15]. On a high-level, the lemma states that for the adversary to be able to equivocate in Stage 3, there exists a unique set S that it must bias the outcome of the coin-toss in Stage 2 so that the resulting set is S . On a high-level, we can rely on this lemma, as a malicious initiator that equivocates must bias the outcome to a particular set S and using the set S . Then, we can construct an adversary $\tilde{\mathcal{A}}^{\mathcal{O}'}$ that violates the CCA-game for $\langle C, R \rangle$ by simply detecting this set S in the outcome of Stage 2.

More formally, given τ_n , and a partial view of \mathcal{A} , let us assume that \mathcal{A} equivocates with probability $\frac{1}{q(n)}$ for some polynomial $q(\cdot)$ and infinitely many n .

Before we recall the information theoretic lemma from [HV15], we first explain how our protocol is an instance of the protocol in their work. In [HV15], they construct an adaptively secure UC-commitment in the CRS hybrid where the protocol proceeds as follows:

1. In Stage 1, the committer using the same strategy as the initiator in our Stage 1 commits to a string m , where instead of using Com_σ , it uses an encryption scheme with oblivious ciphertext generation property (where the public-key for this scheme is placed in the CRS).
2. In Stage 2, the committer and receiver execute a coin-toss where the receiver makes the first move just as in $\langle I, R \rangle$ with the exception that the receiver in their protocol uses again an encryption scheme (with the public-key in the CRS) instead of a commitment scheme to commit to σ_0 .
3. In the decommitment phase of their protocol, the committer reveals its commitment just as the initiator does in Step 2 of Stage 3 in our protocol.

We remark that in essence, the protocol in [HV15] is used as a subprotocol in our work here where the initiator commits to a string m and then reveals it. The only property they need of the encryption scheme is that it is statistically binding⁶ and has the oblivious generation property. In our protocol, the Naor commitment scheme has both these properties. (See our next protocol for such a variant)

Claim 5.2 Restatement of Claim 5.5 [HV15] *Let τ be a fixed partial transcript up until end of Stage 1. Then, except with negligible probability, there exists no two transcripts $\text{trans}_1, \text{trans}_2$ that satisfy the following conditions:*

1. trans_1 and trans_2 are complete and accepting transcripts of π_{COM} with τ being their prefix.
2. There exists two distinct sets S_1, S_2 such that S_1 and S_2 are the respective outcomes of the coin-tossing phase within trans_1 and trans_2 .
3. There exist valid decommitments to two distinct strings in trans_1 and trans_2 .

Since the commitment made by our Initiator can be viewed as an instance of their protocol, we can conclude that there exists a unique set S that should be the outcome of the coin-toss in Stage 2 for a malicious initiator to equivocate m . Since \mathcal{A} equivocates with probability $\frac{1}{q(n)}$ it holds, there is a set S such that with the probability negligible close⁷ to $\frac{1}{q(n)}$, starting from τ_n , the outcome of Stage 2 is S . To construct an adversary $\hat{\mathcal{A}}$ that violates the CCA-security of the underling $\langle C, R \rangle$ scheme, we simply incorporate \mathcal{A} and use as auxiliary input τ_n, S and the partial view of \mathcal{A} . Next, it forwards the $\langle C, R \rangle$ interaction in Stage 2 to an external committer. All queries to the helper oracle \mathcal{O} by \mathcal{A} can be simulated using \mathcal{H} and $\hat{\mathcal{A}}$ simply uses \mathcal{H} to emulate \mathcal{O} . Then it halts the execution right after the adversary in the internal emulation reveals σ_1 . Now, $\hat{\mathcal{A}}$ simply

⁶ Any commitment can either be revealed a encryption of a unique string x or as a ciphertext that was obviously generated. In particular, it cannot be revealed as a an encryption of two distinct messages x and x' .

⁷ The probability is not identically equal to $\frac{1}{q(n)}$ since the commitment scheme is only statistically binding and not perfectly binding.

outputs $\sigma_0 = \sigma \oplus \sigma_1$ where σ is the string that maps to the set S . This violates the CCA game as with probability close to $\frac{1}{q(n)}$, $\hat{\mathcal{A}}$ identifies the message committed using $\langle C, R \rangle$. \square \blacksquare

6 Realizing \mathcal{F}_{COM} Using CCA-Secure Coin-Tossing

In this section, we provide our black-box construction of \mathcal{H} -EUC secure protocol Π_{COM} . Our protocol is a variant of the protocol described in [HV15] where it is shown how to realize \mathcal{F}_{COM} in the CRS model assuming only public-key encryption that admits oblivious-ciphertext generation with adaptive UC-security. While the [HV15] protocol assumes that every pair of parties share an independently generated CRS, in this work we assume no setup, but will require the stronger simulatable public-key encryption scheme. Assume that $\langle I, R \rangle$ is a CCA-secure coin-tossing scheme and that the public-key encryption scheme (Gen, Enc, Dec) is augmented with algorithms (oGen, oRndEnc, rGen, rRndEnc) which implies a simulatable public-key encryption scheme. Then we start with a formal description of our protocol.

Consider a helper functionality \mathcal{H} that “biases” the coin-toss in an interaction using $\langle I, R \rangle$ in the same way as the biasing oracle \mathcal{O} does, subject to the condition that player P_i in a protocol instance sid can only query the functionality on interactions that use identity (P_i, sid) . More precisely, every party P_i can simultaneously engage with \mathcal{H} in multiple sessions of $\langle I, R \rangle$ as an initiator using identity P_i where the functionality simply forwards all the messages internally to the biasing oracle \mathcal{O} , and ensures that the result of the coin-tossing is biased to a prescribed outcome at the end of each session. See Figure 3 for a formal description of the functionality. We note here that since \mathcal{O} can be implemented in super-polynomial time, this functionality can also be implemented in super-polynomial time.

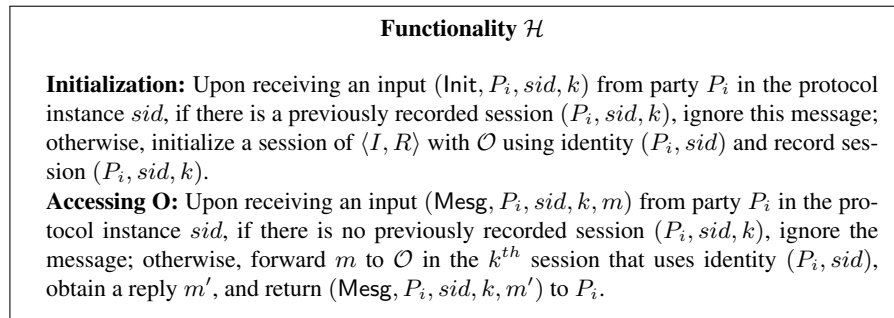


Fig. 3. The Helper Functionality \mathcal{H} (i.e. angel).

We proceed with a proof sketch of Lemma 4.1 before we proceed to a formal proof.

Proof overview: Recalling that an adversary can adaptively corrupt both parties, for the overview, we present the hardest cases for simulation, which is static corruption of one party and adaptive corruption of the other party.

<p>Protocol π_{COM}.</p> <p>Sender's Input: A message $m \in \{0, 1\}$ and a security parameter 1^n.</p> <p>Commitment Phase:</p> <p>Stage 1: Key Generations Phase: The sender and receiver engage in a protocol using $\langle I, R \rangle$ where the receiver acts as the initiator I and the sender acts as R. Let $\text{PK} = \text{oGen}(\text{out}(\tau_{S \rightarrow R}))$ where $\tau_{S \rightarrow R}$ is the transcript of the interaction.</p> <p>Stage 2: Input Encoding Phase: The sender chooses a random n-degree polynomial $p(\cdot)$ over a field $\mathbb{F}[x]$ such that $p(0) = m$. Namely, it randomly chooses $a_i \leftarrow \mathbb{F}$ for all $i \in [n]$ and sets $a_0 = m$, and defines the polynomial $p(x) = a_0 + a_1x + \dots + a_nx^n$. The sender then creates a commitment to m as follows. For every $i \in [3n + 1]$, it first pick $b_i \leftarrow \{0, 1\}$ at random and then computes:</p> $c_i^{b_i} = \text{Enc}_{\text{PK}}(p_0(i); t_i) \text{ and } c_i^{1-b_i} = \text{oRndEnc}(\text{PK}, r_i)$ <p>where $r_i, t_i \leftarrow \{0, 1\}^n$. The sender sends $(c_0^0, c_0^1), \dots, (c_{3n+1}^0, c_{3n+1}^1)$ to the receiver.</p> <p>Stage 3: Cut-and-choose Phase: The sender and receiver engage in a protocol using $\langle I, R \rangle$ where the sender acts as the initiator I and the receiver acts as R. Define a subset $\Gamma \subset [3n + 1]$ of size n using the outcome $\text{out}(\tau_{R \rightarrow S})$ where $\tau_{R \rightarrow S}$ is the transcript of the interaction. The sender provides the plaintexts encrypted in $\{c_i^{b_i}\}_{i \in \Gamma}$ by sending the sequence $\{b_i, p(i), t_i\}_{i \in \Gamma}$. The receiver verifies that all the decryptions are correct and aborts otherwise.</p> <p>Decommitment Phase: Let $\Gamma' = [3n + 1] - \Gamma$. The sender reveals its input m and all the plaintexts encrypted in $\{c_i^{b_i}\}_{i \in \Gamma'}$. The receiver checks if all the decryptions are correct and aborts otherwise. Using the n polynomial evaluations revealed relative to $i \in \Gamma$ and any additional polynomial evaluation that was revealed relative to Γ', the receiver reconstructs the polynomial $p(\cdot)$ (via polynomial interpolation of $n + 1$ points). Next, the receiver verifies whether $p(0) = m$, and that for every $i \in [3n + 1]$ the point $p(i)$ is the decrypted value within $c_i^{m_i}$.</p>

Fig. 4. Protocol Π_{COM} that realizes \mathcal{F}_{COM} using a CCA-secure coin-tossing protocol $\langle I, R \rangle$

Simulating static corruption of receiver and post-execution corruption of sender. To simulate the messages for a honest sender, the simulator generates random shares for 0 and 1 that agree on a randomly chosen n subset $\tilde{\Gamma}$ (chosen in advance). It then encrypts these shares in Stage 2 where for each index it randomly positions the shares for 0 and 1. Next, in Stage 3, the simulator biases $\tau_{R \rightarrow S}$ using the helper \mathcal{H} so that the subset generated using $\text{out}(\tau_{R \rightarrow S})$ is exactly $\tilde{\Gamma}$. As these shares are common for a sharing of 0 and 1, revealing them in the commit phase will go undetected. Later in the decommit phase, it can chose to reveal shares of 0 or 1 depending on the real message m (to show that the unopened shares were obviously generated will be done by exploiting the invertible sampling algorithm for the simulatable encryption scheme). The core argument in proving indistinguishability of simulation will be to reduce the hiding property of Stage 2 to the semantic-security of the underlying encryption scheme on a public-key generated using Gen, i.e., the CPA-security of the encryption scheme, where we will rely on the CCA-security game w.r.t challengers for our coin-tossing protocol

to achieve this. We discuss this reduction on a high-level below. Before that we remark that the adversary will not be able to use the helper oracle \mathcal{H} to bias the outcome of the coin-toss in Stage 1 because the helper oracle will not provide access to the biasing oracle on sessions where the party querying the helper is not the responder R of that coin-tossing session.

Reduction: The challengers $(\mathcal{C}_0, \mathcal{C}_1)$ for our CCA-game, on input a string o will set $\text{PK} = \text{rGen}(o)$. For a predetermined message t , \mathcal{C}_0 will output a ciphertext that is an honest encryption of t using Enc and \mathcal{C}_1 will obviously generate a ciphertext using oRndEnc . It will follow from the security guarantees of the simulatable public-key encryption that for a randomly chosen o , no (stand-alone) adversary can distinguish the outputs of \mathcal{C}_0 or \mathcal{C}_1 even given o (i.e. $\text{STA}_0 \approx \text{STA}_1$).

Now given an adversary \mathcal{A} controlling the receiver in our coin-tossing scheme $\langle I, R \rangle$ we consider a sequence of hybrid experiments where we replace the encryptions in Stage 2 from the honest sender's strategy to the simulated strategy. Namely, obviously generated ciphertexts $c_j^{1-b_j}$ will be generated using the encryption algorithm. More precisely, we consider a sequence of hybrids $H^0 = \mathbf{REAL}, H^1, \dots, H^{3n+1}$ where in the H^i we generate $c_j^{1-b_j}$ for $j = 1, \dots, i$ in Stage 2 according to the simulator's strategy (i.e. encryption of valid messages as opposed to being obviously generated). Next we show that H^{i-1} and H^i are indistinguishable. The only difference between the two hybrids is in how $c_i^{1-b_i}$ is generated. More precisely, in H^{i-1} , $c_i^{1-b_i}$ is generated using oRndEnc and in H^i it is generated using Enc . We now reduce the indistinguishability of the hybrids to the semantic-security of the encryption scheme via the CCA-game of $\langle I, R \rangle$. Towards this, we consider the pair of challengers $(\mathcal{C}_0, \mathcal{C}_1)$ described above for which the stand-alone game is hard.

Next, consider an oracle adversary $\tilde{\mathcal{A}}$ that internally incorporates \mathcal{A} and the environment and proceeds as follows: $\tilde{\mathcal{A}}$ forwards every oracle query made by \mathcal{A} to its oracle and forwards the interaction using $\langle I, R \rangle$ in Stage 1 externally to an honest receiver. $\tilde{\mathcal{A}}$ then stalls the internal emulation upon having the interaction within $\langle I, R \rangle$ complete, and outputs the view of \mathcal{A} and the outcome of the coin-toss o from the internal emulation, in the external interaction. Then it interacts with \mathcal{C} that on input o produces a ciphertext. Internally, $\tilde{\mathcal{A}}$ feeds the ciphertext in place of $c_i^{1-b_i}$ in Stage 2. The rest of the encryptions are honestly generated according to the strategy in H^i .

It now follows that if the message t is chosen according to the strategy in H^i , then we have that $\text{hyb}^{i-1} = \text{EXP}_1(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z)_{n \in \mathbb{N}, z \in \{0,1\}^*}$ and $\text{hyb}^i = \text{EXP}_0(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z)_{n \in \mathbb{N}, z \in \{0,1\}^*}$ where hyb^{i-1} and hyb^i are the views of the adversary \mathcal{A} in the hybrids H^{i-1} and H^i . Therefore, if hyb^{i-1} and hyb^i are distinguishable by the CCA-security of $\langle I, R \rangle$ we have that there exists a stand-alone PPT algorithm B that distinguishes the interaction with \mathcal{C}_0 and \mathcal{C}_1 for a randomly sampled coin-toss outcome o . Recalling that $\text{STA}_0 \approx \text{STA}_1$ by the hiding property of obviously generated ciphertexts in the underlying encryption scheme and thus we arrive at a contradiction. Therefore, hyb^{i-1} and hyb^i must be indistinguishable.

To complete this case, we need to handle post-execution corruption of the sender. This can be achieved exactly as in the decommitment phase which reveals all the randomness used in the commitment phase.

Simulating malicious senders. For a honest receiver, the simulator first biases the outcome of the coin-toss in Stage 1, so that PK is a public-key for which it knows the corresponding secret-key. This will allow the simulator to decrypt the ciphertexts provided by the adversary in Stage 2. However, this does not ensure extraction as an adversarial sender can equivocate just as the simulator for honest senders. Showing that there is a unique value that can be extracted requires showing that a corrupted sender cannot successfully predict exactly the n indexes Γ from $\{1, \dots, 3n + 1\}$ that will be chosen in the coin-tossing protocol. Using an information-theoretic argument from [HV15], we know that after an encoding phase, for any adversary to break binding (i.e. equivocate) it must ensure that the coin-tossing phase results in a particular set Γ . We can reduce the binding property of our scheme to the CCA-security of underlying coin-tossing scheme. Roughly speaking it suffices for our CCA-security to guarantee that the coin-toss outcome has high-entropy which can be shown by using a suitable challenger. Finally, to obtain extraction, we rely on a strategy from [HV15], that can determine the message using the decryptions from Stage 1 and the coin-toss outcome in Stage 3. Finally, to address post-execution corruption of the receiver we observe that it suffices to generate the messages for the receiver honestly and upon corruption simply provide the random coins of this honest receiver.

Formal proof of Correctness of UC-Commitment Protocol: Let \mathcal{A} be a PPT adversary that attacks Protocol Π_{COM} described in Figure 6 and recall that simulator \mathcal{S} interacts with the ideal functionality \mathcal{F}_{COM} and with the environment \mathcal{Z} . Then \mathcal{S} starts by invoking a copy of \mathcal{A} and running a simulated (internal) interaction of \mathcal{A} with the environment \mathcal{Z} and parties running the protocol. We fix the following notation. First, the session and sub-session identifiers are respectively denoted by sid and $ssid$. Next, the committing party is denoted P_i and the receiving party P_j . \mathcal{S} proceeds as follows:

Simulating the communication with \mathcal{Z} : Every message that \mathcal{S} receives from \mathcal{Z} it internally feeds to \mathcal{A} and every output written by \mathcal{A} is relayed back to \mathcal{Z} .

Simulating the commitment phase when the receiver is statically corrupted: In this case \mathcal{S} uses the honest sender's algorithm in Stage 1 and in Stage 2 proceeds as follows. Upon receiving message $(sid, \text{Sen}, \text{Rec})$ from \mathcal{F}_{COM} , the simulator picks a random subset $\tilde{\gamma} \subset [3n + 1]$ of size n and two random n -degree polynomials $p_0(\cdot)$ and $p_1(\cdot)$ such that p_0 and p_1 agree on all points $i \in \tilde{\Gamma}$ and $p_0(0) = 0$ and $p_1(0) = 1$.

- For every $i \in \tilde{\Gamma}$ the simulator proceeds as the honest sender would with polynomial $p_0(\cdot)$. Namely, it first picks $b_i \leftarrow \{0, 1\}$ at random and then sets the following pairs, $c_i^{b_i} = \text{Enc}_{\text{PK}}(p_0(i); t_i)$ and $c_i^{1-b_i} = \text{oRndEnc}(\text{PK}, r_i)$ where $r_i, t_i \leftarrow \{0, 1\}^n$ (we recall that $p_0(i) = p_1(i)$ for all $i \in \tilde{\Gamma}$).
- For every $i \in \tilde{\Gamma}' = [3n + 1] - \tilde{\Gamma}$ the simulator picks $b_i \leftarrow \{0, 1\}$ at random and then uses the points on both polynomials $p_0(\cdot)$ and $p_1(\cdot)$ to calculate the following pairs, namely $c_i^{b_i} = \text{Enc}_{\text{PK}}(p_0(i); t_i^0)$ and $c_i^{1-b_i} = \text{Enc}_{\text{PK}}(p_1(i), t_i^1)$ where $t_i^0, t_i^1 \leftarrow \{0, 1\}^n$ are chosen uniformly at random.

Finally, the simulator sends the pairs $(c_0^0, c_0^1), \dots, (c_{3n+1}^0, c_{3n+1}^1)$ to the receiver.

Next, in Stage 3, the simulator biases the coin-tossing result so that the set Γ that is chosen in this phase is identical to $\tilde{\Gamma}$. More precisely, produces coins c that will yield $\tilde{\Gamma}$

in Stage 3 and sends c to \mathcal{H} . Next, it forwards the messages the simulator receives from \mathcal{A} controlling R in this interaction using $\langle I, R \rangle$ to \mathcal{H} . Recall that the helper function will bias the outcome of this interaction to c (as the identity of this interaction is not equal to any identity made by the \mathcal{A}). Finally, the simulator reveals the plaintexts in all the ciphertexts within $\{c_i^{b_i}\}_{i \in \tilde{\Gamma}}$.

Simulating the decommitment phase where the receiver is statically corrupted:

Upon receiving a message $(\text{reveal}, \text{sid}, m)$ from \mathcal{F}_{COM} , \mathcal{S} generates a simulated decommitment message as follows. Recall first that the simulator needs to reveal points on a polynomial $p(\cdot)$ and pairs $\{(b_i, t_i)\}_{i \in [3n+1]}$ such that $p(0) = m$ and $c_i^{b_i} = \text{Enc}_{\text{PK}}(p(i); t_i)$. Let $\hat{b}_i = b_i \oplus m$ for all $i \in \tilde{\Gamma}'$, then \mathcal{S} reveals $p_m(\cdot)$, $\{\hat{b}_i, t_i^{\hat{b}_i}, r_i = \text{rRndEnc}(\text{PK}, t_i^{1-m}, p_{1-m}(i))\}_{i \in \tilde{\Gamma}'}$.

Simulating the commit phase when the sender is statically corrupted: Simulating the sender involves extracting the committed value as follows. In Stage 1, \mathcal{S} first samples (PK, SK) using the Gen algorithm with randomness r_G . Then it runs rGen on r_G to obtain c which it forwards to the helper \mathcal{H} . Then, it forwards the messages the simulator receives from \mathcal{A} controlling R in this interaction using $\langle I, R \rangle$ to \mathcal{H} . Recall that the helper function will bias the outcome of this interaction to c . This means that the public-key obtained from the coin-tossing is PK.

The simulation next uses the honest receiver's algorithm in Stages 2 and 3. Let Γ be the set obtained from the outcome of the coin-tossing phase. To extract the input, \mathcal{S} chooses an arbitrary index $j \in [3n+1] - \Gamma$ and reconstructs two polynomials $q(\cdot)$ and $\tilde{q}(\cdot)$ such that

$$\begin{aligned} q(i) &= \tilde{q}(i) = \beta_i^{b_i} \quad \forall i \in \Gamma \\ q(j) &= \beta_j^0 \quad \text{and} \quad \tilde{q}(j) = \beta_j^1 \quad \text{and} \quad q(0), \tilde{q}(0) \in \{0, 1\}. \end{aligned}$$

It then verifies whether for all $i \in [3n+1]$, $q(i) \in \{\beta_i^0, \beta_i^1\}$ and $\tilde{q}(i) \in \{\beta_i^0, \beta_i^1\}$. The following cases arise:

- Case 1:** Both $q(\cdot)$ and $\tilde{q}(\cdot)$ satisfy the condition and $\tilde{q}(0) \neq q(0)$. Then \mathcal{S} halts returning fail. Below we prove that the simulator outputs fail with negligible probability.
- Case 2:** At most one of $q(\cdot)$ and $\tilde{q}(\cdot)$ satisfy the condition or $\tilde{q}(0) = q(0)$. \mathcal{S} sends $(\text{commit}, \text{sid}, q(0))$ to the \mathcal{F}_{COM} functionality and stores the committed bit $q(0)$. Otherwise, \mathcal{S} sends a default value.
- Case 3:** Neither $q(\cdot)$ or $\tilde{q}(\cdot)$ satisfy the condition. \mathcal{S} sends a default value to the ideal functionality and need not store the committed bit since it will never be decommitted correctly.

Simulating adaptive corruptions: We remark that we only provide the description of the simulator for static corruption. If any honest party is adaptively corrupted during the simulation, since the simulation is straight-line and admits post-execution corruption, it can directly generate coins even in the middle of the execution.

Below we analyze each of the scenarios above, and show that no environment \mathcal{Z} interacting with \mathcal{S} in the ideal-world is distinguishable from that with \mathcal{A} in the real-world in each of the cases.

Analysis of receiver corruptions: Our proof follows a sequence of hybrids from the real world execution to the ideal world execution.

Hybrid H_0 : H_0 is identical to the real world execution.

Hybrid H_1 : The hybrid experiment H_1 proceeds identically to H_0 with the exception that a set $\tilde{\Gamma}$ of size n is chosen at random and the coin-tossing interaction using $\langle I, R \rangle$ in Stage 3 is biased so that the outcome yields $\tilde{\Gamma}$. Hybrids H_0 and H_1 are identically distributed except when the oracle \mathcal{O} fails. Since this happens only with negligible probability, the outputs of the two experiments are statistically close.

Hybrid H_2 : We gradually change the ciphertexts generated in Stage 2 from the real committer to the simulation. Indistinguishability of experiment H_1 and H_2 will rely on the security of the encryption scheme. simulatable public-key encryption scheme, we will require to bias the PK chosen in Stage 1 to a challenge public-key obtained from the challenger for the encryption security game. We will be able to do this by relying on the security game of our CCA-secure coin-tossing protocol. More formally, consider a sequence of hybrids H_1^0, \dots, H_1^{3n+1} where in the H_1^i we generate $c_j^{1-b_j}$ for $j = 1, \dots, i$ according to the simulator's strategy (i.e. encryption of valid messages as opposed to being obliviously generated). Now we show that H_1^{i-1} and H_1^i are indistinguishable. The only difference between the two hybrids is in how $c_i^{1-b_i}$ is generated. More precisely, in H_1^{i-1} , $c_i^{1-b_i}$ is generated using oRndEnc and in H_1^i it is generated using Enc . We now reduce the indistinguishability of the hybrids to the semantic-security of the encryption scheme via the CCA-game of $\langle I, R \rangle$.

Towards this, we give a challenger \mathcal{C} for which the stand-alone game is hard. On a high-level this game will be the semantic-security of the underlying simulatable public-key encryption scheme where the public-key is sampled using rGen on the coin-toss o .

Reduction: More formally, given a message t , define $\mathcal{C}(o, b)$ as the strategy that sets $\text{PK} = \text{rGen}(o)$ and outputs a ciphertext that was honest encryption of t using Enc when $b = 0$ and obliviously generated using oRndEnc when $b = 1$.

Next consider an oracle adversary $\tilde{\mathcal{A}}$ that internally incorporates \mathcal{A} and the environment and proceeds as follows: $\tilde{\mathcal{A}}$ forwards every oracle query made by \mathcal{A} to its oracle and forwards the interaction using $\langle I, R \rangle$ in Stage 1 externally to an honest receiver. Let o be the outcome of the interaction in the internal emulation. an encryption of a message using Enc or generates one obliviously.

Then it interacts with \mathcal{C} that on input o produces a ciphertext. Internally, $\tilde{\mathcal{A}}$ feeds the ciphertext in place of $c_i^{1-b_i}$ in Stage 2.

It now follows that if the message t is chosen according to the strategy in H_1^i , then

$$\begin{aligned} \text{hyb}_1^{i-1} &= \text{EXP}_1(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z)_{n \in \mathbb{N}, z \in \{0,1\}^*} \\ \text{hyb}_1^i &= \text{EXP}_0(\langle I, R \rangle, \mathcal{O}, \mathcal{A}, \mathcal{C}, n, z)_{n \in \mathbb{N}, z \in \{0,1\}^*} \end{aligned}$$

where hyb_1^{i-1} and hyb_1^i are the views of the adversary \mathcal{A} in the hybrids H_1^{i-1} and H_1^i . Therefore, if hyb_1^{i-1} and hyb_1^i are distinguishable by the CCA-security of $\langle I, R \rangle$ we have that there exists a stand-alone PPT algorithm B that distinguish the interaction with \mathcal{C}_0 and \mathcal{C}_1 for a randomly sampled coin-toss outcome o . This

violates the semantic-security of the encryption scheme and thus we arrive at a contradiction. Therefore, hyb_1^{i-1} and hyb_1^i must be indistinguishable.

Hybrid H_3 : In this hybrid, we follow H_2 except that we use the simulation strategy to decommit to the message m received from the \mathcal{F}_{COM} -functionality. Since in H_2 the commitment phase has been setup to be equivocated, this follows directly. Again using the CCA-security of $\langle I, R \rangle$ just as we used to argue indistinguishability for hybrids H_1 and H_2 , we can reduce the indistinguishability of H_2 and H_3 to the security of the underlying simulatable public-key encryption scheme.

Finally, we conclude by observing the H_3 is identical to the ideal world experiment.

Analysis of sender corruptions: Our proof follows a sequence of hybrids from the real world execution to the ideal world execution.

Hybrid H_0 : H_0 is identical to the real world execution.

Hybrid H_1 : This experiment proceeds identical to H_0 with the exception that we forward the interaction using $\langle I, R \rangle$ in Stage 1 to the oracle \mathcal{H} . More precisely, we pick (PK, SK) using the Gen algorithm with randomness r_G . Then rGen is invoked on r_G to obtain c which it forwards to the helper \mathcal{H} . Recall that \mathcal{H} will bias the coin-toss outcome to c and the resulting public-key agreed upon will be PK. Indistinguishability of H_1 and H_0 can be reduced directly to the indistinguishability of real and obviously generated public-keys of the simulatable public-key encryption scheme using the CCA-security of $\langle I, R \rangle$.

Hybrid H_2 : H_2 is the same as H_1 with the exception that the value committed to by the adversary is extracted using the simulator's strategy and forwarded to \mathcal{F}_{COM} . The only difference between the hybrids H_1 and H_2 is that in H_2 we extract a value for the commitment from the adversarial sender. This means that to argue indistinguishability it suffices to show that the value extracted is correct (i.e. the scheme is binding). We argue this by relying on the information-theoretic lemma proved in [HV15]. In more detail, this lemma shows that at the end of Stage 2, it is possible to define a set Γ such that for any adversarial sender to equivocate it needs to bias the outcome of the coin toss in Stage 3 to result in this set Γ . This coin-toss is decided using our protocol $\langle I, R \rangle$ where the adversarial sender controls the initiator and by relying on CCA-security we argue next that there exists no adversary that can bias the outcome to result in a particular set with non-negligible probability.

Suppose for contradiction there exists an adversary \mathcal{A} that can bias the outcome to Γ in H_1 with non-negligible probability. We now construct an adversary \mathcal{A}' that incorporates \mathcal{A} and internally emulates the hybrid experiment H_2 with the exception that it forwards the interaction of \mathcal{A} in Stage 3 to an external honest receiver. Now, consider a pair of challengers $(\mathcal{C}_0, \mathcal{C}_1)$ for the CCA-security game where \mathcal{C}_0 outputs 1 if the outcome o results in Γ and 0 otherwise and \mathcal{C}_1 outputs 0 irrespective of the outcome. By our assumption on \mathcal{A} , this means that EXP_0 and EXP_1 with the adversary \mathcal{A}' are distinguishable because the adversary biases the coin-toss to result in Γ with non-negligible probability. However, since a uniformly sampled coin will result in Γ with at most negligible probability we have that STA_0 and STA_1 are indistinguishable which is a contradiction. Therefore, we have that the value extracted by our simulator is correct except with non-negligible probability

and this concludes the proof. Finally, we conclude by observing the H_2 is identical to the ideal world experiment.

7 Application: A Zero-One Law for Adaptive Security

We extend the result of [MPR10] and establish a zero-one law under adaptive UC-reduction. More formally, we show that all (non-reactive)⁸ functionalities fall into two categories: *trivial* functionalities, those which can be UC-reduced to any other functionality; and *complete* functionalities, to which any other functionality can be UC-reduced.

Theorem 7.1 *Assume the existence of simulatable public-key encryption scheme. Then every two-party non-reactive functionality is either trivial or complete in the UC framework in the presence of adaptive, malicious adversaries.*

Proof. An important step in proving the zero-one law in [MPR10] was to identify all non-trivial functionalities into one of four categories (i.e. functionalities):

1. \mathcal{F}_{XOR} : This functionality enables simultaneous exchange of information, such as the XOR function.
2. \mathcal{F}_{CC} : This functionality enables to selectively hide one party's input from the other, typically characterized as a cut-and-choose functionality.
3. \mathcal{F}_{OT} : This functionality enables OT of inputs from one party to another.
4. \mathcal{F}_{COM} : This functionality allows information in internal memory to be hidden between rounds, an instance of which is the commitment functionality.

Specifically, it was shown in [MPR10] that every non-trivial functionality \mathcal{F} can realize one of the above four functionalities with information-theoretic security. We are able to demonstrate the zero-one law by proving the following key lemma.

Lemma 7.1 (Informal) *Assume the existence of simulatable public-key encryption scheme. Then \mathcal{F}_{COM} can be realized in the $\mathcal{F}_{\text{COIN}}$ -hybrid model in the presence of adaptive, malicious adversaries, using black-box access to the encryption scheme.*

As mentioned before, in order to demonstrate the zero-one law it suffices to show that the four categories of non-trivial functionalities are complete, where it suffices to only consider \mathcal{F}_{OT} , \mathcal{F}_{XOR} and \mathcal{F}_{CC} when considering non-reactive functionalities. Recalling that the previous results [IPS08, CDMW09] establish completeness of \mathcal{F}_{OT} and \mathcal{F}_{COM} , where the latter result additionally requires the existence of stand-alone adaptively secure semi-honest oblivious-transfer protocol, it is thus left to show that the remaining two categories \mathcal{F}_{CC} and \mathcal{F}_{XOR} are complete. We note first that combining our lemma with the result of [CDMW09] establishes that \mathcal{F}_{XOR} is complete. We remark here that simulatable PKE schemes are sufficient to construct adaptive semi-honest OT which is required in the transformation of [CDMW09]. In order to show that \mathcal{F}_{CC} is complete, we recall that in [MPR10], \mathcal{F}_{CC} is reduced to another functionality called the $\mathcal{F}_{\text{EXTCOM}}$ -functionality for the static corruptions case. Roughly speaking this functionality is a mild variant of the \mathcal{F}_{COM} functionality that admits straight-line extraction without straight-line equivocation. For more details, we refer the reader to the full version

⁸ Such functionalities are computed in a single round of communication with the functionality.

or [MPR10]. We argue that the same protocol also realizes $\mathcal{F}_{\text{EXTCOM}}$ in the presence of adaptive corruptions. On a high-level, we are able to accomplish this since $\mathcal{F}_{\text{EXTCOM}}$ does not require equivocation. To complete the picture, we show how to construct a variant of the $\mathcal{F}_{\text{COIN}}$ functionality in the $\mathcal{F}_{\text{EXTCOM}}$ -hybrid and argue that this variant suffices to establish that $\mathcal{F}_{\text{EXTCOM}}$ is complete even for the adaptive case. Our constructions make use of the underlying primitives only in a black-box manner.

References

- BCNP04. Boaz Barak, Ran Canetti, Jesper Buus Nielsen, and Rafael Pass. Focs. pages 186–195, 2004.
- Bea91. Donald Beaver. Foundations of secure interactive computing. In *CRYPTO*, pages 377–391, 1991.
- BS05. Boaz Barak and Amit Sahai. How to play almost any mental game over the net - concurrent composition via super-polynomial simulation. In *FOCS*, pages 543–552, 2005.
- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.
- CDMW09. Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Simple, black-box constructions of adaptively secure protocols. In *TCC*, pages 387–402, 2009.
- CDPW06. Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. *IACR Cryptology ePrint Archive*, 2006:432, 2006.
- CF01. Ran Canetti and Marc Fischlin. Universally composable commitments. In *CRYPTO*, pages 19–40, 2001.
- CKL06. Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. On the limitations of universally composable two-party computation without set-up assumptions. *J. Cryptology*, 19(2):135–167, 2006.
- CLP10. Ran Canetti, Huijia Lin, and Rafael Pass. Adaptive hardness and composable security in the plain model from standard assumptions. In *FOCS*, pages 541–550, 2010.
- CLP13. Ran Canetti, Huijia Lin, and Rafael Pass. From unprovability to environmentally friendly protocols. In *FOCS*, pages 70–79, 2013.
- CPS07. Ran Canetti, Rafael Pass, and Abhi Shelat. Cryptography from sunspots: How to use an imperfect reference string. In *FOCS*, pages 249–259, 2007.
- DDN03. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Review*, 45(4):727–784, 2003.
- DMRV13. Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Muthuramakrishnan Venkatasubramanian. Adaptive and concurrent secure computation from new adaptive, non-malleable commitments. In *ASIACRYPT*, pages 316–336, 2013.
- DN00. Ivan Damgård and Jesper Buus Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. In *CRYPTO*, pages 432–450, 2000.
- GLP⁺15. Vipul Goyal, Huijia Lin, Omkant Pandey, Rafael Pass, and Amit Sahai. Round-efficient concurrently composable secure computation via a robust extraction lemma. In *TCC*, pages 260–289, 2015.
- GMW87. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.

- HV15. Carmit Hazay and Muthuramakrishnan Venkatasubramanian. On black-box complexity of universally composable security in the CRS model. In *ASIACRYPT*, pages 183–209, 2015.
- IPS08. Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In *CRYPTO*, pages 572–591, 2008.
- Kiy14. Susumu Kiyoshima. Round-efficient black-box construction of composable multi-party computation. In *CRYPTO*, pages 351–368, 2014.
- KLP07. Yael Tauman Kalai, Yehuda Lindell, and Manoj Prabhakaran. Concurrent composition of secure protocols in the timing model. *J. Cryptology*, 20(4):431–492, 2007.
- KMO14. Susumu Kiyoshima, Yoshifumi Manabe, and Tatsuaki Okamoto. Constant-round black-box construction of composable multi-party computation protocol. In *TCC*, pages 343–367, 2014.
- Lin03. Yehuda Lindell. General composition and universal composability in secure multi-party computation. In *FOCS*, pages 394–403, 2003.
- LP12a. Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In *CRYPTO*, pages 461–478, 2012.
- LP12b. Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up (full version). Available at <https://www.cs.ucsb.edu/~rachel.lin>, 2012.
- LPV08. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.
- LPV09. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *STOC*, pages 179–188, 2009.
- LZ11. Yehuda Lindell and Hila Zarosim. Adaptive zero-knowledge proofs and adaptively secure oblivious transfer. *J. Cryptology*, 24(4):761–799, 2011.
- MMY06. Tal Malkin, Ryan Moriarty, and Nikolai Yakovenko. Generalized environmental security from number theoretic assumptions. In *TCC*, pages 343–359, 2006.
- MPR10. Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A zero-one law for cryptographic complexity with respect to computational UC security. In *CRYPTO*, pages 595–612, 2010.
- MR91. Silvio Micali and Phillip Rogaway. Secure computation (abstract). In *CRYPTO*, pages 392–404, 1991.
- Nao91. Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- ORSV13. Rafail Ostrovsky, Vanishree Rao, Alessandra Scafuro, and Ivan Visconti. Revisiting lower and upper bounds for selective decommitments. In *TCC*, pages 559–578, 2013.
- Pas03. Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.
- PR08. Manoj Prabhakaran and Mike Rosulek. Cryptographic complexity of multi-party computation problems: Classifications and separations. In *CRYPTO*, pages 262–279, 2008.
- PS04. Manoj Prabhakaran and Amit Sahai. New notions of security: achieving universal composability without trusted setup. In *STOC*, pages 242–251, 2004.
- RK99. Ransom Richardson and Joe Kilian. On the concurrent composition of zero-knowledge proofs. In *EUROCRYPT*, pages 415–431, 1999.
- Ven14. Muthuramakrishnan Venkatasubramanian. On adaptively secure protocols. In *SCN*, pages 455–475, 2014.
- Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.

A Adaptive Extension to CCA-Secure Commitments

In our work, we need to consider the CCA-Security game in the presence of an adaptive adversary \mathcal{A} . This definition as we will see does not require full-fledged adaptive security and in particular our definition will not imply equivocability of the commitments.

We recall the CCA-security game for the commitments as introduced in [CLP10]. Roughly speaking, a commitment scheme is CCA-secure if the commitment scheme retains its hiding property even if the receiver has access to a “decommitment oracle”. The experiment considers an oracle adversary \mathcal{A} with oracle access to a helper function \mathcal{H} and interacts as the receiver with an honest committer C . In our adaptive setting, we will require two additional properties: (1) The adversary will be allowed to corrupt the external committer C . However, security is required to hold, i.e. hiding property of the left commitment, only if the committer is not corrupted, and (2) In the interaction between the adversary and the helper oracle, where it interacts as the committer, the adversary will be allowed to corrupt the receiver. In this case, the helper oracle is required to provide random coins for the receiver consistent with the transcript.

The second property does not require any explicit change in the definition of the security as it only alters the semantics of the interaction between \mathcal{A} and \mathcal{H} . The first property however needs to be incorporated in the definition which we do next.

Modifying the IND_b random variable in the definition. In the standard definition $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)$ represents the output of the $\mathcal{A}^\mathcal{O}$ in an experiment where it interacts with an honest committer with input $b \in \{0, 1\}^n$. This output is set to \perp , if the identity of the execution with C is the same as the identity of any interaction of \mathcal{A} with \mathcal{O} . We define a new random variable $\overline{\text{IND}}_b(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)$ which is equal to $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)$ only if $\mathcal{A}^\mathcal{O}$ does not corrupt the honest committer C in the execution. Otherwise it is set to \perp .

Definition 5 (CCA-secure commitments with adaptive adversary) *Let $\langle C, R \rangle$ be a tag-based commitment scheme with $l(n)$ -bit identities, and \mathcal{O} a committed-value oracle for it. We say that $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O} in the presence of an adaptive adversary, if for every PPT \mathcal{A} , the following ensembles are computationally indistinguishable: $\{\overline{\text{IND}}_0(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)\}_{n \in \mathbb{N}} \approx \{\overline{\text{IND}}_1(\langle C, R \rangle, \mathcal{O}, \mathcal{A}, n, z)\}_{n \in \mathbb{N}}$. We say that $\langle C, R \rangle$ is CCA-secure if there exists a committed-value oracle \mathcal{O}' , such that, $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O}' .*

Theorem A.1 *Assume the existence of one-way functions. Then, for every $\epsilon > 0$, there exists a $O(n^\epsilon)$, there exists a $O(n^\epsilon)$ -round commitment scheme that is CCA-secure w.r.t. the committed-value oracle in the presence of an adaptive adversary and only relies on black-box access to one-way functions (where n is the security parameter).*

Proof Sketch: Lin and Pass [LP12a] gave a black-box construction of a $O(n^\epsilon)$ -round CCA-secure commitment scheme $\langle C, R \rangle$. We rely on the same construction for our stronger definition of security in the presence of an adaptive adversary. We provide a high-level proof sketch of its correctness. We begin with a short overview of their proof.

In the proof of standard security of the scheme provided in [LP12a], the idea is to reduce the indistinguishability of the IND_b experiments to the stand-alone hiding property of a different commitment scheme $\langle \tilde{C}, \tilde{R} \rangle$ (that is a slight variant of $\langle C, R \rangle$).

The main part of the proof is to show that given an oracle adversary for $\langle C, R \rangle$ there exists a stand-alone malicious receiver R^* (that does not have access to the oracle) for $\langle \tilde{C}, \tilde{R} \rangle$. On a high-level, R^* will internally incorporate \mathcal{A} and emulate the committed-value oracle for \mathcal{A} while forwarding the left interaction externally to \tilde{C} (which it can do as it is a variant that has a “similar” structure). To emulate the oracle, R^* needs to extract the value committed value which it will accomplish by rewinding the right interactions. Two issues arise:

- Since the left interaction is forwarded to an external committer, R^* needs to be able to rewind the right interactions without rewinding the left. The main idea here that is reminiscent of previous work [DDN03,LPV08] is to identify the so-called *safe-points* where this can be done. In slight more detail, when rewinding from a safe-point the only thing the adversary can do in the left interaction is to request “complete” (3-round witness-indistinguishable) proofs and such a request will be accommodated by the variant $\langle \tilde{C}, \tilde{R} \rangle$.
- There are unbounded-many right interactions and will result in R^* recursively rewinding interactions to extract the committed value in the interactions. In [LP12a], they achieve this by provided several points to rewind from and rely on the [RK99] to ensure that expected running time of the rewindings in each level is polynomial and the recursive depth is at most a constant.

Next, we argue why the same protocol satisfies our stronger definition of security. We begin with the observation that if the adversary \mathcal{A} does not corrupt the left or right interactions, then our definition reduces to the standard CCA-security. We will prove security identically to [LP12a] by reducing it to the stand-alone hiding property of $\langle \tilde{C}, \tilde{R} \rangle$. We will employ the exact rewinding strategy as in [LP12a] for R^* with the following exception: Our definition of *safe-point* will have one additional requirement: A *safe-point* for our scheme is any *safe-point* according to [LP12a] with the added requirement that the adversary corrupts neither the committer in the left-interaction or the receiver of the right interaction (associated with the safe-point) before the 3-round witness indistinguishable (WI) proof associated with the safe-point completes.⁹

We remark that our definition of *safe-point* can modularly replace the definition in [LP12a] and the entire proof goes through. This is because the definition affects only the run-time analysis of the reduction. For the run-time analysis to go through the only requirements are that there are sufficiently many *safe-point*’s and when rewound from a safe-point, it continues to be a *safe-point* with at least the same probability (See Step 1 in Sub-Claim 2 of [LP12b]).¹⁰ The first property holds because, a right receiver needs to be rewound only if \mathcal{A} completes the entire right session without corrupting the right receiver or the left committer. In this event there will be as many safe points according to the definition of [LP12a] as there according to ours. The second property holds because a rewinding will be cancelled only if the point is not safe. This concludes the proof.

⁹ In [LP12b] the definition of a safe-point is parameterized with the depth of the recursion and our additional requirement naturally extends to the definition.

¹⁰ The second property is to ensure that the expected running time of the rewinding is polynomial in each recursive depth. In more detail, if a point in the transcript is safe with probability p , then if the property holds the expected number of times before which a rewinding is successful is at most $O(\frac{1}{p})$.