Stronger Definitions for
Anonymous Electronic Cash

Mårten Trolin
marten@nada.kth.se

Royal Institute of Technology (KTH), Stockholm, Sweden

Abstract In this paper we investigate definitions of security for previously propesed schemes and point out that they can be strengthend so that the bank does need to be trusted to the same extent. We give experiment-based definitions for our stronger notion. Finally we propose a scheme secure under our definitions in the common reference string (CRS) model under the assumption that trapdoor permutations exist.

1   Introduction

Electronic payments is an interesting cryptographic task. The concept was introduced by Chaum et al. [10]. As was common at the time, the claimed security properties were not defined in a precise way. Many schemes for anonymous electronic followed [7,12,26,23,20,19,18,8],

In the recent years, several papers have focused on giving precise security definitions for tasks such as group signatures [2,3] and ring signatures [4]. In this paper we investigate the security properties of schemes for electronic cash. We point out that the definitions current schemes do not rule out a corrupt bank cheating a user. The scenario is that the bank claims a user has withdrawn, but the user denies this. We argue that the protocol should include a mechanism to solve such an issue. We suggest stronger security definitions to remedy this. We also and construct a scheme using general methods, which is secure in the common reference string (CRS) model. The scheme is not intended for practical use, but it is rather a proof of concept.

2   Notation and Definitions

We write $[a, b]$ to denote the set $\{x \in \mathbb{Z} \mid a \leq x \leq b\}$, and we let $[b] = [1, b]$. We say that an element is chosen "randomly" instead of the more cumbersome "independently and uniformly at random". By $r \leftarrow_R S$ we mean that r is chosen randomly in S. Throughout the paper, κ denotes the security parameter. A function $\varepsilon : \mathbb{N} \to \mathbb{R}^+$ is said to be negligible if for each $c > 0$ there exists a $\kappa_0 \in \mathbb{N}$ such that $\varepsilon(\kappa) < \kappa^{-c}$ for $\kappa_0 < \kappa \in \mathbb{N}$. We say that a function $f : \mathbb{N} \to \mathbb{R}^+$ is non-negligible whenever it is not negligible.

We write $\emptyset$ to denote both the empty set and the empty string, and we let $\perp$ we a special symbol. All adversaries in this paper are modeled as polynomial

time Turing machines with *non-uniform* auxiliary advice string. We denote the set of such adversaries by PPT*.

Informally a family of functions is called a family of *trapdoor permutations* if the permutation is hard to invert unless a trapdoor is known, in which case it can be efficiently inverted. The precise definition is given in Appendix A.

3 Security Model

Here we give definitions for a scheme for electronic cash as well as for its security.

While some security properties are obvious and dealt with from the very first scheme, others are more subtle. Naturally a scheme must not allow for a user to forge coins, and a double-spender must be detected. The schemes [8,27] require that a corrupt bank cannot accuse an honest user of double-spending, whereas this requirement is not explicit in, e.g., [18,17]. However, to our knowledge, no scheme discusses the possibility of a corrupt bank falsely claiming that an honest user has withdrawn a coin, or rejecting a deposition from a merchant of a legally spent coin. The tendency seem to be to, apart from anonymity, protect the interests of the bank rather than of the user.

We give definitions requiring that the bank be able to prove withdrawals. Thus after executing the withdrawal protocol the output of the bank should be a proof of withdrawal and the output of the user should be a valid coin. However, the user should not be able to abort prematurely so that the bank cannot construct the proof of withdrawal. While there exist protocols that release information piece by piece, so called fair exchange [5], they are either based on gradual release of information and thus not very practical, or require the presence of a third trusted party. Since we would like definitions that can be instantiated with a practical protocol, we use a different approach. After the execution, the bank receives a withdrawal proof, and the user receives a coin secret key. The honest bank would send the proof to the user, who can use it as a coin. Should the bank fail to do this, the user can challenge the transaction and require the bank to prove that a coin has indeed been withdrawn. Since the proof can be used a coin, the scheme is fair also from the point of view of the user.

We require that spent coins be publicly verifiable to avoid the possibility of the bank rejecting a deposition and to ensure a merchant cannot deny having received a payment. In particular the bank can verify a spent coin. Therefore there is no need for an interactive deposition protocol, since the merchant simply hands the spent coin to the bank.

The merchants do not have a secret key in our setting. Instead the receiving merchant's identity mid is encoded in the spent coin together with a transaction ID tid. Thus the merchant's consent is not necessary in order to spend a coin. We use this approach to make the definitions cleaner. In practice, a user would require some sort of contract before handing the merchant a coin, but we feel this is best handled outside of the protocol. In standard banking systems it is indeed possible to wire money without the recipient's approval, although this is

not very sensible to do. Since there is no secret for the merchant and the resulting spent coin is publicly verifiable, the spending protocol is non-interactive.

We assume the existence of a PKI, i.e., given a public key there exists a method to obtain the identity of the holder of the key. We also assume the existence of secure and authenticated communication. We do not explicitly define a protocol to register a user. If the protocol requires some secret information to be passed from the bank to the user, this can be done in the withdrawal protocol. Therefore there is no loss of generality in this.

Here we discuss payment schemes containing all basic properties, but there are many possible extensions. Examples of such alternate definitions include the presence of a trusted third party which can identify coins, even when they have not been double-spent. Such schemes are called *fair*. Another extension is the possibility to transfer a coin between users in several steps before it is deposited at the bank, and divisible coins. We leave it as an open problem to adjust the definitions to handle also such cases.

3.1   Participants

The participants are the bank $\mathcal{B}$, merchants $\mathcal{M}_i$, and users $\mathcal{U}_i$.

3.2   Algorithms and Protocols

We now define the algorithms and the protocols that a scheme for electronic cash consists of. For non-interactive algorithms the definitions are straight-forward. We define two-party protocols as a pair of algorithms, where each participant executes one algorithm. The algorithms take as input a message, a state, and the private input of the party. On startup of a protocol the initiating party executes the algorithm with $\emptyset$ as message. Each algorithm outputs a pair (msg, state), where msg is handed as message to the other party's algorithm, and state is passed as input by the executing party the next round. When a party's algorithm outputs $p = \bot$ the protocol is finished, and the final value of each party's state is parsed as private output. The *transcript* of a protocol is defined as the list of messages exchanged.

The below algorithm illustrates the execution between two parties using algorithms A and B with private input sk_A, sk_B, respectively. The parameter msg_A is given as input to the algorithm and the private output is returned.

$\mathsf{state}_A \leftarrow \mathsf{sk}_A$
$\mathsf{state}_B \leftarrow \mathsf{sk}_B$
while $(\mathsf{msg}_A \neq \emptyset) \wedge (\mathsf{msg}_B \neq \emptyset)$ **do**
$\quad (\mathsf{msg}_B, \mathsf{state}_A) \leftarrow A(\mathsf{msg}_A, \mathsf{state}_A)$
$\quad (\mathsf{msg}_A, \mathsf{state}_B) \leftarrow B(\mathsf{msg}_B, \mathsf{state}_B)$
end while
return $(\mathsf{state}_A, \mathsf{state}_B)$

Note that the only protocol we use involve exactly two parties.

Any algorithm not explicitly stated to be deterministic is assumed to be randomized.

There is an algorithm for creating a bank key pair and a user key pair. After the user has generated its key, it is inserted into the PKI and hence tied to the user's identity. Each merchant has an identity $\mathsf{mid} \in \{0,1\}^{\kappa/2}$, but no secret key.

Algorithm Header 1 (Bank Key Generation BKg).
INTERFACE: $\mathsf{BKg}(1^\kappa)$, where κ is the security parameter.
OUTPUT: $(\mathsf{bpk}, \mathsf{bsk})$, where bpk is a bank public key and bsk is a bank secret key.

Algorithm Header 2 (User Key Generation UKg).
INTERFACE: $\mathsf{UKg}(1^\kappa)$, where κ is the security parameter.
OUTPUT: $(\mathsf{upk}, \mathsf{usk})$, where upk is a user public key and usk is a user secret key.

Protocol Header 1 (Coin Withdrawal, $(\mathsf{UWithdraw}, \mathsf{BWithdraw})$).
PARTIES: Bank $\mathcal{B}$, User $\mathcal{U}$.
PRIVATE INPUT OF $\mathcal{B}$: Bank public key bpk, bank secret key bsk.
PRIVATE INPUT OF $\mathcal{U}$: Bank public key bpk, user public key upk, user secret key usk.
PRIVATE OUTPUT OF $\mathcal{B}$: Coin coin.
PRIVATE OUTPUT OF $\mathcal{U}$: Coin secret key csk.

$(\mathsf{UWithdraw}, \mathsf{BWithdraw})$ is the protocol used when a user withdraws a coin. The private input of the user is a user private key usk, a user public key upk, and a bank public key bpk. The private input of the bank is a bank secret key bsk and a bank public key bpk. The user's private output is a coin secret key csk used when spending the coin, whereas the bank's output is interpreted as a withdrawn coin (coin). Normally the bank would hand the withdrawn coin to the user, but we do not address this in the protocol.

Algorithm Header 3 (Coin Spending Spend).
INTERFACE: $\mathsf{Spend}(\mathsf{coin}, \mathsf{upk}, \mathsf{usk}, \mathsf{csk}, \mathsf{mid}, \mathsf{tid}, \mathsf{bpk})$, where coin is a coin, upk is a user public key, usk is a user secret key, csk is a coin secret key, mid is a merchant identity, $\mathsf{tid} \in \{0,1\}^{\kappa/2}$ is a transaction identity, and bpk is a bank public key.
OUTPUT: $\mathsf{spentcoin}$, where $\mathsf{spentcoin}$ is a (publicly verifiable) spent coin.

$\mathsf{VfDoubleSpent}(\mathsf{spentcoin}_1, \mathsf{spentcoin}_2, \mathsf{bpk})$ returns the public key upk of the double-spender if $\mathsf{spentcoin}_1$ and $\mathsf{spentcoin}_2$ are transcripts of two spendings of the same coin. Otherwise it returns $\bot$.

Algorithm Header 4 (Identifying a Double-Spender, $\mathsf{VfDoubleSpent}$).
INTERFACE: $\mathsf{VfDoubleSpent}(\mathsf{spentcoin}_1, \mathsf{spentcoin}_2, \mathsf{bpk})$, where the parameters $\mathsf{spentcoin}_1$ and $\mathsf{spentcoin}_2$ are two spent coins and bpk is a bank public key.
OUTPUT: (upk), a (possibly empty) user public key.

In addition to the above, there are two non-interactive algorithms which verify the validity of proofs produced durning withdrawal and spending, VfCoin(coin, upk, bpk), VfSpentCoin(spentcoin, mid, bpk), The algorithms output 1 if the proof is valid with regards to the additional input parameters and 0 otherwise.

Algorithm Header 5 (Verifying a Withdrawal, VfCoin).
INTERFACE: VfCoin(coin, upk, bpk), where coin is a withdrawn coin, upk a user public key, and bpk a bank public key.
OUTPUT: $b \in \{0, 1\}$.

Algorithm Header 6 (Verifying a Spent Coin, VfSpentCoin).
INTERFACE: VfSpentCoin(spentcoin, tid, mid, bpk), where spentcoin is a spent coin, tid $\in \{0, 1\}^{\kappa/2}$ is a transaction identity, mid is a merchant identity, and bpk is a bank public key.
OUTPUT: $b \in \{0, 1\}$.

3.3 Correctness

By correctness we mean that the scheme works as expected when all participants are honest. Proving correctness is often straight-forward, and this property is sometimes not stated explicitly. Here we define correctness for a scheme as defined above.

Experiment 1 (Correctness, $\mathbf{Exp}_{\mathcal{EC}}^{\text{correct}}(\kappa)$).
  $(\text{bpk}, \text{bsk}) \leftarrow \text{BKg}(1^\kappa)$
  $(\text{upk}, \text{usk}) \leftarrow \text{UKg}(1^\kappa)$
  $(\text{coin}, \text{csk}) \leftarrow \text{Withdraw}(\text{bpk}, \text{bsk}, \text{usk})$
  if VfCoin(coin, upk, bpk) = 0 **then**
    return 0
  end if
  spentcoin $\leftarrow$ Spend(usk, coin, csk, mid, tid, bpk)
  if VfSpentCoin(spentcoin, mid, bpk) = 0 **then**
    return 0
  end if
  return 1

Definition 1 (Correctness). *A scheme for electronic cash $\mathcal{EC}$ is correct if* $\Pr[\mathbf{Exp}_{\mathcal{EC}}^{\text{correct}}(\kappa) = 0]$ *is negligible as a function of κ.*

Detection of double-spenders is not included in the definition of correctness. This may seem strange at first, but as we will see later the definition of unforgeability implies that double-spenders are detected. Another way to see it is that correctness only stipulates how the protocol works with honest parties, and an honest party does not double-spend.

3.4 Security

We describe four experiments, or games, to define security for a scheme for electronic cash.

In each experiment the adversary has access to a number of oracles defined below. They operate on the following global parameters.

- $\mathfrak{U}$ contain all public keys inserted into the PKI.
- $\mathfrak{C}$ contain the public keys of corrupt users. Obviously $\mathfrak{C}$ is a subset of $\mathfrak{U}$.
- $(\mathsf{upk}_i, \mathsf{usk}_i)$ is the public and private key of the ith honest user.
- l is the number of coins withdrawn from the bank using the withdrawal oracle.
- $CoinKeySet_i$ is the set of coin keys for user upk_i produced when the withdrawal oracle is used. For a new user it is initiated as the empty set.

$\mathsf{HonestUKg}(1^\kappa)$ calls $\mathsf{UKg}(1^\kappa)$ to generate a key pair $(\mathsf{upk}, \mathsf{usk})$. The public key upk is inserted into the PKI. The key pair $(\mathsf{upk}, \mathsf{usk})$ is stored in the key list. The public key upk is returned.

$\mathsf{CorruptUKg}(1^\kappa)$ calls $\mathsf{UKg}(1^\kappa)$ to generate a key pair $(\mathsf{upk}, \mathsf{usk})$. The public key upk is inserted into the PKI and is inserted into the set $\mathfrak{C}$. The key pair $(\mathsf{upk}, \mathsf{usk})$ is returned.

$\mathsf{HonestUWithdraw}(i, j, \mathsf{msg})$ lets the adversary interact in the withdrawal protocol with user i in session j. More precisely, if session j has not been instantiated, i.e., state_j is not defined, then $\mathsf{state}_j \leftarrow (\mathsf{upk}_i, \mathsf{usk}_i, \mathsf{bpk})$. Then a call is made to $\mathsf{UWithdraw}(\mathsf{msg}, \mathsf{state}_j)$ with output $(\mathsf{msg}', \mathsf{state}_j)$. The message msg' is returned, and state_j is stored. After the session has finished, state_j is parsed as a coin secret key csk_j and stored. The key set for user i is updated $\mathsf{CSK}_i \leftarrow \mathsf{CSK}_i \cup \{\mathsf{csk}_j\}$.

$\mathsf{HonestBWithdraw}(j, \mathsf{msg})$ lets the adversary interact in the withdrawal protocol with the bank in session j. More precisely, if session j has not been instantiated, i.e., state_j is not defined, then $\mathsf{state}_j \leftarrow (\mathsf{bpk}, \mathsf{bsk})$. Then a call is made to $\mathsf{BWithdraw}(\mathsf{msg}, \mathsf{state}_j)$ with output $(\mathsf{msg}', \mathsf{state}_j)$. The message msg' is returned, and state_j is stored. After the session has finished, state_j is parsed as a coin coin and returned. Each time a coin is returned the counter l is incremented.

$\mathsf{HonestSpend}(i, j, \mathsf{coin}, \mathsf{mid}, \mathsf{tid})$ first checks if (i, coin) has been stored or $\mathsf{csk}_j \notin \mathsf{CSK}_i$ and returns $\perp$ if this is the case. Otherwise it stores (i, coin), calls $\mathsf{Spend}(\mathsf{coin}, \mathsf{upk}_i, \mathsf{usk}_i, \mathsf{csk}_j, \mathsf{mid}, \mathsf{tid}, \mathsf{bpk})$, and returns the output.

It can be noted that the construction of the $\mathsf{HonestSpend}$ oracle prevents the adversary from using it to double-spend a coin.

In addition to the above we use ordinary dot notation when the adversary is given oracle access to protocol algorithms, e.g., $\mathsf{BWithdraw}(\cdot, (\mathsf{bpk}, \mathsf{bsk}))$.

We let $\mathfrak{Q}_O$ be the set of queries to oracle O and we let $\mathfrak{R}_O$ be the set of responses.

Unforgeability The property of *unforgeability* informally says that one cannot create valid coins by other means than withdrawing them from the bank. A little more precisely it says that if a coalition of users spend more than they have legally withdrawn, then one of them will get caught as a double-spender. Recall the l is the number of withdrawn coins using the withdrawal oracle. Unforgeability corresponds to the property *balance* of [8].

Experiment 2 (Unforgeability, $\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{unforge}}(\kappa)$).

$(\mathsf{bpk}, \mathsf{bsk}) \leftarrow \mathsf{BKg}(1^\kappa)$
$(\mathsf{spentcoin}_1, \ldots, \mathsf{spentcoin}_k) \leftarrow A^{\mathsf{CorruptUKg}(1^\kappa), \mathsf{HonestBWithdraw}(\cdot, \cdot)}(\mathsf{bpk})$
if $k \leq l$ **then**
   return 0
end if
if $\exists i \in [k] : \mathsf{VfCoin}(\mathsf{spentcoin}_i, \mathsf{bpk}) = 0$ **then**
   return 0
end if
if $\exists (i, j) \in [k]^2 : \mathsf{VfDoubleSpent}(\mathsf{spentcoin}_i, \mathsf{spentcoin}_j, \mathsf{bpk}) \in \mathfrak{C}$ **then**
   return 0
end if
return 1

The advantage of the adversary is defined as

$$\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{unforge}}(\kappa) = \Pr[\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{unforge}}(\kappa) = 1] \ .$$

Definition 2 (Unforgeability). *A scheme for electronic cash $\mathcal{EC}$ has* unforgeability *if for any $A \in \mathrm{PPT}$ the advantage $\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{unforge}}(\kappa)$ is negligible as a function of κ.*

Non-Frameability A potential problem could be that a coalition of users and possibly the bank could accuse an honest user of double-spending. We say the a scheme has *non-frameability* if it is infeasible to frame an honest user in such a way.

Note that in the experiment below, the adversary cannot use the $\mathsf{HonestSpend}$ oracle to create a double-spending, due to the fact that the construction of the oracle prevents two queries on the same coin.

Non-frameability corresponds to *weak exculpability* of [8]. The strong variant would guarantee that a user double-spending t times cannot be accused of $t + 1$ double-spendings, but we choose not to include such a game in our defintion.

Experiment 3 (Non-Frameability, $\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{non-frame}}(\kappa)$).

$(\mathsf{bpk}, \mathsf{state}) \leftarrow A(\mathsf{setup}, 1^\kappa)$
$((\mathsf{spentcoin}_1, \mathsf{tid}_1, \mathsf{mid}_1), (\mathsf{spentcoin}_2, \mathsf{tid}_2, \mathsf{mid}_2)) \leftarrow$
   $A^{\mathsf{HonestUKg}(1^\kappa), \mathsf{HonestUWithdraw}(\cdot, \cdot, \cdot), \mathsf{HonestSpend}(\cdot, \cdot, \cdot, \cdot)}(\mathsf{guess}, \mathsf{state})$
if $\mathsf{VfSpentCoin}(\mathsf{spentcoin}_1, \mathsf{tid}_1, \mathsf{mid}_1, \mathsf{bpk}) = 0$ **then**
   return 0

end if
if VfSpentCoin(spentcoin$_2$, tid$_2$, mid$_2$, bpk) = 0 **then**
   return 0
end if
if VfDoubleSpent(spentcoin$_1$, spentcoin$_2$, bpk) $\in \mathfrak{U} \setminus \mathfrak{C}$) **then**
   return 1
else
   return 0
end if

The advantage of the adversary is defined as

$$\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{non-frame}}(\kappa) = \Pr[\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{non-frame}}(\kappa) = 1] \ .$$

Definition 3 (Non-Frameability). *A scheme for electronic cash $\mathcal{EC}$ has non-frameability if for any $A \in$ PPT the advantage $\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{non-frame}}(\kappa)$ is negligible as a function of κ.*

Anonymity A scheme for electronic cash is *anonymous* if it is infeasible for any player, including the bank, to decide the identity of a spender. We define anonymity in a very strong sense, namely that not even the spender herself can recognize her coins (except, of course, by remembering the coins she has spent).

We let the adversary corrupt the bank fully, which includes possibly creating the bank keys in a dishonest way. This differs from the definitions of [16,3], where the group manager keys are supposed to be generated honestly.

Experiment 4 (Anonymity, $\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{anon}-b}(\kappa)$).
  (bpk, state) $\leftarrow A(\mathsf{setup}, 1^\kappa)$
  $(i_0, i_1, \mathsf{mid}, \mathsf{tid}) \leftarrow$
    $A^{\mathsf{HonestUKg}(1^\kappa), \mathsf{HonestUWithdraw}(\cdot,\cdot,\cdot), \mathsf{HonestSpend}(\cdot,\cdot,\cdot,\cdot)}(\mathsf{choose}, \mathsf{state})$
  spentcoin $\leftarrow \mathsf{Spend}(\mathsf{coin}_{i_b}, \mathsf{usk}_{i_b}, \mathsf{csk}_{i_b}, \mathsf{mid}, \mathsf{tid}, \mathsf{bpk})$
  $d \leftarrow A^{\mathsf{HonestUKg}(1^\kappa), \mathsf{HonestUWithdraw}(\cdot,\cdot,\cdot), \mathsf{HonestSpend}(\cdot,\cdot,\cdot,\cdot)}(\mathsf{guess}, \mathsf{state}, \mathsf{spentcoin})$
  if $(\{\mathsf{coin}_{i_0}, i_0, \mathsf{mid}, \mathsf{tid}), (\mathsf{coin}_{i_1}, i_1, \mathsf{mid}, \mathsf{tid})\} \cap \mathfrak{Q}_{\mathsf{HonestSpend}} \neq \emptyset$ **then**
   return 0
  end if
  if $d = b$ **then**
   return 1
  else
   return 0
  end if

The advantage of the adversary is defined as

$$\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{anon}}(\kappa) = |\Pr[\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{anon}-0}(\kappa) = 1] - \Pr[\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{anon}-1}(\kappa) = 1]| \ .$$

Definition 4 (Anonymity). *A scheme for electronic cash $\mathcal{EC}$ has* anonymity *if for any $A \in$ PPT the advantage $\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{anon}}(\kappa)$ is negligible as a function of κ.*

Exculpability Exculpability states that the bank should not be able to forge proofs of withdrawal.

Experiment 5 (Exculpability, $\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{exculp}}(\kappa)$).

> $(\mathsf{bpk}, \mathsf{state}) \leftarrow A(\mathsf{setup}, 1^\kappa)$
> $(\mathsf{coin}, i, \mathsf{mid}, \mathsf{tid}) \leftarrow$
> $A^{\mathsf{HonestUKg}(1^\kappa),\mathsf{HonestUWithdraw}(\cdot,\cdot,\cdot),\mathsf{HonestSpend}(\cdot,\cdot,\cdot,\cdot,\cdot)}(\mathsf{guess}, \mathsf{state})$
> **if** $\mathsf{VfCoin}(\mathsf{coin}, \mathsf{upk}_i, \mathsf{bpk})$ **then**
>   **return** 0
> **end if**
> **if** $\exists \mathsf{csk} \in \mathsf{CSK}_i : \mathsf{Spend}(\mathsf{coin}, \mathsf{upk}_i, \mathsf{usk}_i, \mathsf{csk}, \mathsf{mid}, \mathsf{tid}, \mathsf{bpk}) \neq \perp)$ **then**
>   **return** 0
> **end if**
> **return** 1

The advantage of the adversary is defined as

$$\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{exculp}}(\kappa) = \Pr[\mathbf{Exp}_{\mathcal{EC},A}^{\mathsf{exculp}}(\kappa) = 1] \ .$$

Definition 5 (Exculpability). *A scheme for electronic cash $\mathcal{EC}$ has* exculpability *if for any $A \in \mathrm{PPT}$ the advantage $\mathbf{Adv}_{\mathcal{EC},A}^{\mathsf{exculp}}(\kappa)$ is negligible as a function of κ.*

Finally we make the following definition.

Definition 6 (Secure Scheme for Electronic Cash). *A scheme for electronic cash is* secure *if it has unforgeability, non-frameability, anonymity, and exculpability.*

4 A Construction

In this section we describe a secure scheme for electronic cash based on general methods. We first define the primitives, then we give the algorithms, and finally we prove that our scheme is secure according to our definitions.

4.1 Primitives

Our construction uses a signature scheme, a commitment scheme, and simulation sound non-interactive zero-knowledge proofs of knowledge (NIZK). Here we briefly describe these notions, and refer to Appendix A for precise definitions of these well-known concepts.

Digital Signatures A signature scheme $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sig}, \mathsf{Vf})$ is *correct* if for $(\mathsf{pk}, \mathsf{sk})$ generated by Kg and any message m it holds that $\mathsf{Vf}_{\mathsf{pk}}(m, \mathsf{Sig}_{\mathsf{sk}}(m)) = 1$. $\mathcal{SS}$ is secure against chosen-message attacks, *CMA-secure* [15], if it is infeasible to produce valid message-signature pair for *any* message, even if the adversary has access to a signing oracle $\mathsf{Sig}_{\mathsf{sk}}(\cdot)$.

Commitment Schemes A commitment scheme $\mathcal{COM} = (\mathsf{Commit}, \mathsf{Reveal})$ for messages of length κ is secure if a commitment is both hiding and binding, i.e., that the adversary gains any useful information about the committed value from $(c, r)\mathsf{Commit}(m)$, and that given (c, r) it is infeasible to find $(m', r') \neq (c, r)$ such that $\mathsf{Reveal}(c, m', r') = 1$.

Non-interactive Proofs of Knowledge We use non-interactive zero-knowledge proofs of knowledge, or NIZKs, in our construction. Given a language $L \in \mathbf{NP}$ with witness relation R and $x \in L$, a NIZK (P, V) enables a prover P to prove to a verifier V that she knows a witness w such that $(x, w) \in R$.

A proof system is said to be zero-knowledge if there exists a simulator which produces proofs indistinguishable from real proofs, and the condition for it to be called non-interactive should be obvious. A NIZK is complete if for any $(x, w) \in R$ it holds that $V(x, P(x, w)) = 1$ and sound if for any algorithm A the probability that $V(x, \pi) = 1$ and $x \notin L$ is negligible, where $(x, \pi) \leftarrow A(\xi)$. A NIZK is a proof of knowledge (NIZK-PK) if there exists an extractor which, if allowed to choose the CRS, can extract a witness.

In the experiments we give the adversary access to oracles which sometimes produce simulated proofs. Potentially this could help the adversary in producing false proofs. The stronger notion of simulation sound NIZKs requires that no adversary can break the soundness even if given access to a simulator.

It has been shown [22,24] that a simulation sound NIZK exists for any **NP**-relation if trapdoor permutations exist. Combined with the standard method of [1] it can be turned into a NIZK-PK under the assumption that dense encryption schemes exist [1,25]. We detail the construction in Appendix A.

We need NIZKs for languages on the form $L = \{x \in \Im(f)\}$. Here the obvious witness relation is $R = \{(x, w) : f(w) = x\}$. For such a relation we use the notation $\mathrm{NIZK}((: \omega) : f(\omega) = x)$ to denote a NIZK. We use Greek letters to denote variables in the witness, i.e., known only to the prover, and Latin letters for variables known both to the prover and to the verifier. We the verifier denote by Vf. It will be clear from context for which relation the proof is.

4.2 The Protocol

Here we give the definitions for algorithms and protocols that form a scheme for electronic cash in the CRS-model. We begin by giving an informal description. In order to identify double-spenders, we use Ferguson's [12] trick of letting each coin contain a line $y = ax + \mathsf{upk}$, such that the coordinate of its intersection with the y-axis coincides with the identity upk of the owner. When spending the coin, one point on the line is revealed. Thus one spending of a coin gives no information about its owner, whereas the identity can be computed from two spendings of the same coin.

When withdrawing a coin, the user randomly selects the slope a. It computes a commitment $\tilde{a}$ if a and a commitment $\hat{a}$ with associated randomness r_a and $r_{\tilde{a}}$.

$$a \longrightarrow \tilde{a} \longrightarrow \hat{a}$$
$$\searrow \qquad \searrow$$
$$r_a \qquad r_{\tilde{a}}$$

A two-step commitment $\hat{b}$ of the user public key upk is also computed. Then $\hat{a}, \hat{b}$ is sent to the bank and signed, and when a coin is spent, $\tilde{a}$ is revealed together with a proof of knowledge that it is correctly formed. Intuitively this gives anonymity, since $\tilde{a}, \tilde{b}$ cannot be linked to $\hat{a}, \hat{b}$. It also assures that double-spenders are detected, since it is infeasible to open $\hat{a}, \hat{b}$ in more than one way. It can be noted that the signing mechanism is similar to the blind signature scheme found by Fischlin [13].

We let $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sig}, \mathsf{Vf})$ be a CMA-secure signature scheme and we let $\mathcal{COM} = (\mathsf{Commit}, \mathsf{Reveal})$ be a binding and hiding commitment scheme. Such signature schemes and commitment schemes exist if one-way functions exist [21,14], and thus certainly if trapdoor permutations exist.

We use NIZKs for two different relations in the withdrawal and the spending protocols. The NIZKs work in the common reference string model. We will denote the reference string ξ. Each NIZK needs its own CRS, so we divide ξ into two parts so that $\xi = \xi_0 || \xi_1$ such that both ξ_0 and ξ_1 are long enough. We let $S(\mathsf{setup}, 1^\kappa)$ create both ξ_0 and ξ_1, store the two secrets in simstate, and return $\xi_0 || \xi_1$. Now S can simulate and extract proofs for both relations.

We start with the key generation algorithms. Key generation for the bank consists of generating a key for the signature scheme. The key for the user is created by drawing a value at random and computing a commitment to the value. The private key is the random value and coin tosses used in the commitment, and the public key is the commitment. As noted above, we do not include user registration at the bank as part of the protocol.

Definition 7 ($\mathsf{BKg}(1^\kappa)$).

  $(\mathsf{bpk}, \mathsf{bsk}) \leftarrow \mathsf{Kg}(1^\kappa)$

Definition 8 ($\mathsf{UKg}(1^\kappa)$).

  $t \leftarrow_R \{0,1\}^\kappa$
  $(\mathsf{upk}, r_t) \leftarrow \mathsf{Commit}(t)$
  $\mathsf{usk} \leftarrow (t, r_t)$
  return $(\mathsf{upk}, \mathsf{usk})$

Merchant registration is straight-forward. Since our protocol does not use a merchant secret key, registration simply consists of handing the merchant identity to the bank, which registers the merchant.

The coin withdrawal protocol is a two-round protocol with the following steps.

1. The user draws a value a at random and commits to a in two steps, i.e., it computes a commitment $\tilde{a}$ to a, and a commitment $\hat{a}$ to $\tilde{a}$. In the same way

11

it computes a two-step commitment $\hat{b}$ to its public key upk. It also constructs a proof $\pi_{\mathcal{U}}$ of knowledge of a a and coin tosses used in the commitments. It hands $\hat{a}, \hat{b}$ and $\pi_{\mathcal{U}}$ to the bank and stores a together with the coin tosses as the coin secret key csk.

2. The bank verifies that the user is allowed to withdraw a coin and that the proof of knowledge is valid. It then signs the user's public key concatenated with $(\hat{a}, \hat{b})$. The coin consists of the signature, $\hat{a}, \hat{b}$, the user's public key upk, and the proof $\pi_{\mathcal{U}}$.

More precisely, the withdrawal protocols consists of the following two algorithms.

Definition 9 (UWithdraw(msg, state)).

Parse state *as* (upk, usk).
$a \leftarrow_R \{0,1\}^\kappa$
$(\tilde{a}, r_a) \leftarrow \mathsf{Commit}(a)$
$(\hat{a}, r_{\tilde{a}}) \leftarrow \mathsf{Commit}(\tilde{a})$
$(\tilde{b}, r_{\mathsf{upk}}) \leftarrow \mathsf{Commit}(\mathsf{upk})$
$(\hat{b}, r_{\tilde{b}}) \leftarrow \mathsf{Commit}(\tilde{b})$
$\pi_{\mathcal{U}} \leftarrow \mathrm{NIZK}(\alpha, \rho_\alpha, \tilde{\alpha}, \rho_{\tilde{\alpha}}, \tau, \rho_\tau, \rho_{\mathsf{upk}}, \tilde{\beta}, \rho_{\tilde{\beta}} :$
$\quad \mathsf{Reveal}(\tilde{a}, \alpha, r_\alpha) = 1 \wedge \mathsf{Reveal}(\hat{a}, \tilde{\alpha}, \rho_{\tilde{\alpha}}) = 1 \wedge \mathsf{Reveal}(\mathsf{upk}, \tau, \rho_\tau) = 1 \wedge$
$\quad \mathsf{Reveal}(\tilde{\beta}, \mathsf{upk}, \rho_{\mathsf{upk}}) = 1 \wedge \mathsf{Reveal}(\hat{b}, \tilde{\beta}, \rho_{\tilde{\beta}}) = 1)$
return $((a, \tilde{a}, r_a, r_{\tilde{a}}, \tilde{b}, r_{\mathsf{upk}}, r_{\tilde{b}}), (\mathsf{upk}, \hat{a}, \hat{b}, \pi_{\mathcal{U}}))$

Definition 10 (BWithdraw(msg, state)).

Parse state *as* (bsk).
Parse msg *as* $(\mathsf{upk}, \hat{a}, \hat{b}, \pi_{\mathcal{U}})$.
Verify that user with public key upk *is allowed to withdraw a coin, and quit otherwise.*
if $\mathsf{Vf}(\pi_{\mathcal{U}}) = 1$ **then**
$\quad$**return** $(\mathsf{reject}, \emptyset)$
end if
$s \leftarrow_R \mathsf{Sig}_{\mathsf{bsk}}(\mathsf{upk}, \hat{a}, \hat{b})$
$\mathsf{coin} \leftarrow (s, \hat{a}, \hat{b}, \mathsf{upk}, \pi_{\mathcal{U}})$
return $(\mathsf{coin}, \emptyset)$

We also need to be able to verify whether or not a coin has been withdrawn by a certain user by verifying the coin's signature and the user's proof of knowledge.

Definition 11 (VfCoin(coin, upk, bpk)).

Parse coin *as* $(s, \hat{a}, \hat{b}, \mathsf{upk}, \pi_{\mathcal{U}})$.
return $\mathsf{Vf}_{\mathsf{bpk}}((\mathsf{upk}, \hat{a}, \hat{b}), s) \wedge \mathsf{Vf}(\pi_{\mathcal{U}})$

To spend a coin the user first checks that the coin is valid. Then it lets (x, y) be a point on the line $y = ax + \mathsf{upk}$, where a is the coin secret key and upk the public key of the user. The point x is chosen as the concatenation of the

transaction identity and the merchant identity. The user reveals the values $\tilde{a}$ and $\tilde{b}$. The spent coin consists of (p, x, y), and a proof of knowledge of a and upk such that (x, y) is indeed a point on the line and of a bank signature on $(\mathsf{upk}, \hat{a}, \hat{b})$ as well as of coin tosses such that $\hat{a}$ is a commitment of $\tilde{a}$ and $\hat{b}$ of $\tilde{b}$.

Definition 12 ($\mathsf{Spend}(\mathsf{coin}, \mathsf{usk}, \mathsf{csk}, \mathsf{mid}, \mathsf{tid}, \mathsf{bpk})$).

Parse coin *as* $(s, \hat{a}, \hat{u}, \mathsf{upk}, \pi_{\mathcal{U}})$
Parse usk *as* (t, r_t)
Parse csk *as* $(a, \tilde{a}, r_a, r_{\tilde{a}}, \tilde{b}, r_{\mathsf{upk}}, r_{\tilde{b}})$
if $(\mathsf{Vf}_{\mathsf{bpk}}((\mathsf{upk}, \hat{a}, \hat{b}), s) = 0) \vee (\mathsf{Reveal}(\tilde{a}, a, r_a) = 0) \vee (\mathsf{Reveal}(\hat{a}, \tilde{a}, r_{\tilde{a}}) = 0) \vee$
$(\mathsf{Reveal}(\mathsf{upk}, t, r_t) = 0) \vee (\mathsf{Reveal}(\tilde{b}, \mathsf{upk}, r_{\mathsf{upk}}) = 0) \vee (\mathsf{Reveal}(\hat{b}, \tilde{b}, r_{\tilde{b}}) = 0)$ **then**
    return $\perp$
end if
$x \leftarrow \mathsf{mid} || \mathsf{tid}$
$y \leftarrow ax + \mathsf{upk}$
$\pi \leftarrow \mathrm{NIZK}(\iota, \alpha, \rho_\alpha, \hat{\alpha}, \rho_{\tilde{a}}, \rho_{\mathsf{upk}}, \hat{beta}, \rho_{\tilde{b}}, \sigma, \tau, \rho_\tau :$
    $y = \alpha x + \iota \wedge \mathsf{Reveal}(\tilde{a}, \alpha, \rho_\alpha) = 1 \wedge \mathsf{Reveal}(\hat{\alpha}, \tilde{a}, \rho_{\tilde{a}}) = 1 \wedge \mathsf{Reveal}(\tilde{b}, \iota, \rho_{\mathsf{upk}}) \wedge$
    $\mathsf{Reveal}(\hat{\beta}, \tilde{b}, \rho_{\tilde{b}}) = 1 \wedge \mathsf{Vf}_{\mathsf{bpk}}((\iota, \tilde{\alpha}, \tilde{\beta}), \sigma) = 1 \wedge \mathsf{Reveal}(\iota, \rho_\tau, \tau) = 1)$
$\mathsf{spentcoin} \leftarrow (\tilde{a}, \tilde{b}, x, y, \pi)$
return spentcoin

Verification of a spent coin is straight-forward:

Definition 13 ($\mathsf{VfSpentCoin}(\mathsf{spentcoin}, \mathsf{tid}, \mathsf{mid}, \mathsf{bpk})$).

Parse spentcoin *as* $(\tilde{a}, \tilde{b}, x, y, \pi)$.
if $x \neq \mathsf{mid} || \mathsf{tid}$ **then**
    return $\perp$
end if
return $\mathsf{Vf}(\pi)$

Finally we give the algorithm to identify a double-spender. A coin is double-spent if the value b appears twice with different values of x. Finding the double-spender is then simply a task of solving the two equations for upk.

Definition 14 ($\mathsf{VfDoubleSpent}(\mathsf{spentcoin}_1, \mathsf{spentcoin}_2, \mathsf{bpk})$).

Parse $\mathsf{spentcoin}_1$ *as* $(\tilde{a}_1, \tilde{b}_1, x_1, y_1, \pi_1)$.
Parse $\mathsf{spentcoin}_2$ *as* $(\tilde{a}_2, \tilde{b}_2, x_2, y_2, \pi_2)$.
if $((\tilde{a}_1, \tilde{b}_1) \neq (\tilde{a}_2, \tilde{b}_2)) \vee (x_1 = x_2)$ **then**
    return $\perp$
end if
$\mathsf{upk} \leftarrow \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2}$
return upk

5 Proof of Security

In this section we prove the following theorem about the scheme $\mathcal{EC} = (\mathsf{BKg}, \mathsf{UKg}, \mathsf{UWithdraw}, \mathsf{BWithdraw}, \mathsf{VfCoin}, \mathsf{Spend}, \mathsf{VfSpentCoin}, \mathsf{VfDoubleSpent})$ as defined in Section 4.

Theorem 1. *The scheme for electronic cash $\mathcal{EC}$ is correct and secure in the common reference string model if there exists a family of trapdoor permutations.*

We prove the theorem by showing the five properties. Each lemma holds in the CRS-model under the assumption that a family of trapdoor permutations exists, although this is not stated explicitly.

Lemma 1 (Correctness). *The scheme $\mathcal{EC}$ is correct.*

Proof. Follows by the construction of the algorithms.

Lemma 2 (Unforgeability). *The scheme $\mathcal{EC}$ has unforgeability.*

Proof. Let A be an adversary that is successful in Experiment 2 with non-negligible probability. We show how to use A to construct either a machine A_{cma} breaking the CMA-security of the signature scheme $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sig}, \mathsf{Vf})$, a machine A_{binding} breaking the binding property of the commitment scheme $\mathcal{COM}$, or a machine $A_{\mathsf{sim-sound}}$ attacking the simulation soundness of the NIZK-PK.

A_{cma} is given a public key pk for the signature scheme as input. The CRS is created using the simulator $(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$. As in the experiment for CMA security, A_{cma} has access to a signature oracle. It passes pk as parameter bpk to A. The $\mathsf{BWithdraw}$ oracle is run honestly, except that the signature $\mathsf{Sig}_{\mathsf{bsk}}(\cdot)$ is produced by calling the signature oracle.

Let k be the number of $\mathsf{spentcoin}$ produced by A. Recall A has made l withdrawals using its oracle. This implies A_{cma} has made l calls to the CMA oracle. For each $\mathsf{spentcoin}_i = (\tilde{a}_i, \tilde{b}_i, x_i, y_i, \pi_i)$, A_{cma} calls $S(\mathsf{extract}, (\hat{a}_i, \hat{b}_i, x_i, y_i), \pi_i, \xi, \mathsf{simstate})$ to extract (among other parameters) $\sigma_i, \iota, \hat{\alpha}, \hat{\beta}, \rho_\beta$ such that $\mathsf{Vf}_{\mathsf{bpk}}((\iota, \hat{\alpha}, \hat{\beta}), \sigma) = 1$. We now have the following different cases:

1. Signatures on more than l distinct messages are extracted. In this case A_{cma} is successful in breaking the CMA-security of $\mathcal{SS}$.
2. At least one proof π_i cannot be extracted. In this case $A_{\mathsf{sim-sound}}$ uses π_i to break the extractable simulation soundness of the NIZK-PK.
3. All proofs can be extracted, but two proofs yield signatures on the same message $(\iota, \hat{\alpha}, \hat{\beta})$. Since no double-spending is detected, all $(\tilde{\alpha}, \tilde{\beta})$ are distinct. Hence there are two commitments with associated randomness $(\tilde{\alpha}_i, \rho_i)$ and $(\tilde{\alpha}_j, \rho_j)$ such that $\mathsf{Reveal}(\hat{\alpha}, \tilde{\alpha}_i, \rho_i) = \mathsf{Reveal}(\hat{\alpha}, \tilde{\alpha}_j, \rho_j) = 1$. Using these values A_{binding} breaks the binding property of the commitment scheme $\mathcal{COM}$.

Thus we have shown that an adversary which breaks unforgeability can be used to either break the CMA security of $\mathcal{SS}$, the extractable simulation soundness of the NIZK-PK, or break the binding property of $\mathcal{COM}$. Hence $\mathcal{EC}$ has unforgeability.

Lemma 3 (Non-Frameability). *The scheme $\mathcal{EC}$ has non-frameability.*

Proof. Let A be an adversary that succeeds in Experiment 3 with non-negligible probability. We show how to use A to construct either a machine A_{secrecy} breaking the secrecy property of the commitment scheme $\mathcal{COM}$, a machine A_{binding}

breaking the binding property, or a machine $A_{\mathsf{ext-sim-sound}}$, which breaks the extractable simulation soundness of the NIZK-PK.

The machine A_{secrecy} takes part in Experiment 7. It creates a CRS using the simulator $(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$. It randomly draws two message msg_0 and msg_1 which it returns to its experiment, receiving a challenge commitment c. Let the polynomial $p(\kappa)$ be an upper bound on the number of calls to $\mathsf{HonestUKg}$ by A. Since A runs in polynomial time, there exists such a polynomial. A_{secrecy} randomly selects $t \in [p(\kappa)]$. Intuitively A_{secrecy} guesses that A will frame user $\mathcal{U}_t$. All queries to $\mathsf{HonestUKg}$ are executed honestly except for query t, to which A_{secrecy} response c.

When A queries $\mathsf{HonestUWithdraw}$ or $\mathsf{HonestSpend}$ for a user different from $\mathcal{U}_t$, the query is answered honestly. For $\mathcal{U}_t$, the NIZK-PK is constructed by invoking the simulator $S(\mathsf{simulate}, \cdot, \xi, \mathsf{simstate})$.

A outputs two spent coins $\mathsf{spentcoin}_1 = (\tilde{a}_1, \tilde{b}_1, x_1, y_1, \pi_1)$ and $\mathsf{spentcoin}_2 = (\tilde{a}_2, \tilde{b}_2, x_2, y_2, \pi_2)$. By the assumption that they form a double-spending, we have that $(\tilde{a}_1, \tilde{b}_1) = (\tilde{a}_2, \tilde{b}_2)$ and $x_1 \neq x_2$. With probability $1/p(\kappa)$, it holds that $\mathsf{VfDoubleSpent}(\mathsf{spentcoin}_1, \mathsf{spentcoin}_2, \mathsf{bpk}) = \mathsf{upk}_t$. From now on, we will assume that this is the case.

From π_1 and π_2 the machine A tries to extract $\bar{a}, \bar{r}_a, \bar{\mathsf{upk}}, \bar{r}_{\mathsf{upk}}$ such that $\mathsf{Reveal}(\tilde{a}, \bar{a}, \bar{r}_a) = 1$ and $\mathsf{Reveal}(\tilde{b}, \bar{\mathsf{upk}}, \bar{r}_{\mathsf{upk}}) = 1$. We now have the following cases and subcases:

1. None of the proofs were created by the simulator.
   (a) At least one extraction fails. In such case we can construct a machine $A_{\mathsf{ext-sim-sound}}$ using A and breaking the extractable simulation soundness of the NIZK-PK as follows. $A_{\mathsf{ext-sim-sound}}$ takes part in Experiment 13. When A asks for a NIZK-PK, $A_{\mathsf{ext-sim-sound}}$ uses its simulation oracle. The un-extractable NIZK-PK of A is output by $A_{\mathsf{ext-sim-sound}}$, which wins the extractable simulation soundness experiment with non-negligible probability.
   (b) Both extractions succeed but return $(\bar{a}_1, \bar{r}_a^{(1)}, \bar{\mathsf{upk}}_1, \bar{r}_{\mathsf{upk}}^{(1)}) \neq (\bar{a}_2, \bar{r}_a^{(2)}, \bar{\mathsf{upk}}_2, \bar{r}_{\mathsf{upk}}^{(2)})$. In such a case the extracted values can be used by the machine A_{binding} to break the binding property of $\mathcal{COM}$.
   (c) Both extractions succeed and return consistent values. Since the NIZK-PK also proves that $y_i = ax_i + \mathsf{upk}$, it follows that $\bar{\mathsf{upk}} = \mathsf{upk}$. By extracting $\tau, \rho_t au$ such that $\mathsf{Reveal}(\mathsf{upk}, \tau, \rho_\tau$, the machine A_{secrecy} breaks the secrecy of $\mathcal{COM}$ with non-negligible probability.
2. One proof was created by the simulator. Without loss of generality we assume that the simulator created π_2, and let $a_2, r_a^{(2)}, \mathsf{upk}_2, r_{\mathsf{upk}}^{(2)}$ be the values used when responding to the oracle query.
   (a) The extraction of the proof π_1 fails. If this is the case, then $A_{\mathsf{ext-sim-sound}}$ proceeds as in Step 1a to break the simulation soundness of the NIZK-PK.
   (b) The extraction of π_1 succeeds but yields $(\bar{a}_1, \bar{r}_a^{(1)}, \bar{\mathsf{upk}}_1, \bar{r}_{\mathsf{upk}}^{(1)}) \neq (a_2, r_a^{(2)}, \mathsf{upk}_2, r_{\mathsf{upk}}^{(2)})$. Then, as in Step 1b the binding property of $\mathcal{COM}$ is broken.

(c) The extraction π_1 succeeds and gives consistent values. Then, as in Step 1c, the secrecy of $\mathcal{COM}$ is broken.

3. Both proofs were created by the simulator. Since, by the definition of the Spend oracle, each coin can only be spent once, this cannot happen.

We have shown that if A breaks the non-frameability property, then at least one of the machines A_{secrecy}, A_{binding}, and $A_{\mathsf{ext-sim-sound}}$ is successful with non-negligible probability. Since this breaks the assumption, the scheme $\mathcal{EC}$ has non-frameability.

Lemma 4 (Anonymity). *The scheme $\mathcal{EC}$ has anonymity.*

Proof. Assume A wins in the anonymity experiment 4 with non-negligible probability. We show how to construct either A_{secrecy} breaking the secrecy of the commitment scheme $\mathcal{COM}$ or a machine $A_{\mathsf{ad-ind}}$ breaking the adaptive indistinguishability of the NIZK-PK.

We define two variants of the scheme $\mathcal{EC}$. We let $\mathcal{EC}'$ be $\mathcal{EC}$ with the modification that the CRS is created by the simulator for the NIZK-PK in Spend, $(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$ and the NIZK-PK is generated by the simulator, and we let $\mathcal{EC}''$ be $\mathcal{EC}'$ with the difference that the commitment scheme of UWithdraw used to produce $\tilde{a}$ is replaced by a commitment scheme with perfect secrecy.

Since a spentcoin in $\mathcal{EC}''$ contains no information about the spender of a coin, the advantage of A when attacking $\mathcal{EC}''$ is 0. We now have the following two cases.

1. The advantage of A when attacking $\mathcal{EC}'$ is non-negligible. We show how to use A to construct A_{secrecy} which successfully attacks the secrecy of the commitment scheme $\mathcal{COM}$. A_{secrecy} takes part in Experiment 7. All calls to HonestUKg and HonestSpend are answered honestly. When A outputs $(i_0, i_1, \mathsf{mid}, \mathsf{tid})$, A_{secrecy} outputs $(\mathsf{upk}_{i_0}, \mathsf{upk}_{i_1})$ to its experiment, receiving a commitment c in response. Then A_{secrecy} uses c as $\tilde{a}$ when creating the challenge spentcoin and constructs the rest of the coin honestly. (Since $\mathcal{EC}'$ only uses simulated NIZK-PKs, not knowing the message of c is not a problem.) A outputs a bit d, which A_{secrecy} outputs in its experiment.
   From the construction it follows that A_{secrecy} is successful when A is, and hence breaks the secrecy of $\mathcal{COM}$ with non-negligible probability.
2. The advantage of A when attacking $\mathcal{EC}'$ is negligible. In such case we can use A to construct $A_{\mathsf{ad-ind}}$ breaking the adaptive indistinguishability of the NIZK-PK. $A_{\mathsf{ad-ind}}$ takes part in Experiment 10 and 11 while executing Experiment 4 for A. All parts of the experiment are executed honestly, except that NIZK-PKs are created by using the oracle. If A is successful, $A_{\mathsf{ad-ind}}$ responds that it is interacting with Experiment 10, and otherwise that it is interacting with Experiment 11. Since A is successful only when NIZK-PKs are genuine, $A_{\mathsf{ad-ind}}$ has a non-negligible advantage.

We have shown that a machine breaking the anonymity of $\mathcal{EC}$ can be made into a machine either breaking the secrecy of the commitment scheme or a machine breaking the adaptive indistinguishability of the proof system. Since such machines contradicts the assumptions, we conclude that $\mathcal{EC}$ has anonymity.

Lemma 5 (Exculpability). *The scheme $\mathcal{EC}$ has exculpability.*

Proof. Let A be an adversary which wins in Experiment 5 with non-negligible probability. We use A to construct either a machine A_{secrecy} breaking the secrecy property of the commitment scheme $\mathcal{COM}$, A_{binding} breaking the binding property of the $\mathcal{COM}$, a machine $A_{\mathsf{ext-sim-sound}}$ breaking the simulation soundness of the NIZK-PK, or a machine $A_{\mathsf{ad-ind}}$ breaking the adaptive indistinguishability of the NIZK-PK.

We define the scheme $\mathcal{EC}'$ being equal to $\mathcal{EC}$ with the difference that the CRS is setup using the simulator $(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$ and the NIZK-PK of UWithdraw is created using the simulator.

First assume A has negligible probability of breaking the exculpability property of $\mathcal{EC}'$. Then we can use A to construct $A_{\mathsf{ad-ind}}$ in the following way. $A_{\mathsf{ad-ind}}$ takes part in Experiment 10 or 11. It invokes A, answering all queries honestly except that the NIZK-PK is created using its oracle. Hence, if $A_{\mathsf{ad-ind}}$ is run in Experiment 10, it will run $\mathcal{EC}$ for A, but if it is run in 11, it will run $\mathcal{EC}'$. If A is successful, then $A_{\mathsf{ad-ind}}$ returns 0, and otherwise it returns 1. From the construction of $A_{\mathsf{ad-ind}}$ it follows that it breaks the adaptive indistinguishability of the NIZK-PK.

Now assume A has non-negligible probability in winning the exculpability experiment against $\mathcal{EC}'$. Then there are two possiblities.

1. The NIZK-PK $\pi_{\mathcal{U}}$ of coin output by A has been constructed by the simulator. This implies that $\pi_{\mathcal{U}}$ was created by HonestUWithdraw for a certain user and coin secret key $\mathsf{usk}_i, \mathsf{csk}_i$. Hence the coin can be spent using $\mathsf{usk}_i, \mathsf{csk}_i$.
2. The NIZK-PK $\pi_{\mathcal{U}}$ of coin output by A has not been constructed by the simulator. In this case we can construct A_{secrecy} breaking the secrecy of the commitment scheme as follows. The $p(\kappa)$ be an upper bound on the number of calls to HonestUKg. Since A is polynomial, such a bound exists. Let $t \leftarrow_R [p(\kappa)]$. Informally A_{secrecy} guesses that A will frame $\mathcal{U}_t$. A_{secrecy} randomly chooses τ_0, τ_1. It answers queries honestly, except that when asked to generate the public key for $\mathcal{U}_t$, it returns the challenge commitment c created by its experiment.
   With probability $1/p(\kappa)$ A produces a coin that can be verified to belong to $\mathcal{U}_t$. Assume this is the case. Then A_{secrecy} uses the extractor to extact τ, r_τ such that $\mathsf{Reveal}(\mathsf{upk}_t, \tau, r_\tau) = 1$. We now have three cases.
   (a) There exists d such that $\tau_d = \tau$. Then A_{secrecy} returns d hence breaks the secrecy of the commitment scheme $\mathcal{COM}$
   (b) No such d exists. Then two openings of commitment upk_t has been found, violating the binding property of $\mathcal{COM}$.
   (c) The extraction fails, breaking the extractable simulation soundness property of the NIZK-PK.

We can conclude that a machine breaking the exculpability property of $\mathcal{EC}$ implies a machine breaking one of the assumptions. Therefore $\mathcal{EC}$ has exculpability. $\qed$

6 Conclusions

We have given definitions of security that are stronger than what has previously been suggested. We also show that the requirements are realistic by giving a scheme fulfilling them under the assumption of existence of a family of trapdoor permutations.

It remains an open problem to construct a practical scheme which is secure in our sense under some well-established number-theoretical assumptions.

We would like to thank Johan Håstad and Douglas Wikström for pointing out improvements and suggesting relevant references.

References

1. Giuseppe Persiano A. De Santis. Zero-knowledge proofs of knowledge without interaction (extended abstract). In *33rd IEEE Symposium on Foundations of Computer Science – FOCS*, pages 427–436. IEEE Computer Society Press, 1992.
2. M. Bellare, D. Micciancio, and B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 614–629. Springer Verlag, 2003.
3. M. Bellare, H. Shi, and C. Zhang. Foundations of group signatures: The case of dynamic groups. In *RSA Conference 2005, Cryptographers' Track 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 136–153. Springer Verlag, 2005. Full version at `http://eprint.iacr.org/2004/077`.
4. A. Bender, J. Katz, and R. Morselli. Ring signatures: Stronger definitions, and constructions without random oracles. In *Theory of Cryptography Conference – TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 60 – 79. Springer Verlag, 2006. Full version at `http://eprint.iacr.org/2005/304`.
5. M. Blum. How to exchange (secret) keys. *ACM Transactions on Computer Systems – TOCS*, 1(2):175–193, 1983.
6. M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *20th ACM Symposium on the Theory of Computing – STOC*, pages 103–118. ACM Press, 1988.
7. S. Brands. Untraceable off-line cash in wallets with observers. In *Advances in Cryptology – CRYPTO'93*, volume 773 of *Lecture Notes in Computer Science*, pages 302–318. Springer Verlag, 1994.
8. J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact e-cash. In *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 302–321. Springer Verlag, 2005. Full version at `http://eprint.iacr.org/2005/060`.
9. M. Chase and A. Lysyanskaya. On signatures of knowledge. Cryptology ePrint Archive, Report 2006/184, 2006. `http://eprint.iacr.org/2006/184`.
10. D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *Advances in Cryptology – CRYPTO'88*, volume 403 of *Lecture Notes in Computer Science*, pages 319–327. Springer Verlag, 1990.
11. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero-knowledge proofs under general assumptions. *SIAM Journal on Computing*, 29(1):1–28, 1999.

12. N.T. Ferguson. Single term off-line coins. In *Advances in Cryptology – EURO-CRYPT'93*, volume 765 of *Lecture Notes in Computer Science*, pages 318–328. Springer Verlag, 1993.

13. M. Fischlin. Round-optimal composable blind signatures in the common reference string model. In *Advances in Cryptology – CRYPTO2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 60–77. Springer Verlag, 2006.

14. O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *21st ACM Symposium on the Theory of Computing – STOC*, pages 25–32. ACM Press, 1989.

15. S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988.

16. A. Kiayias and M. Yung. Efficient secure group signatures with dynamic joins and keeping anonymity against group managers. In *Mycrypt 2005*, volume 3715 of *Lecture Notes in Computer Science*, pages 151–170. Springer Verlag, 2005.

17. M. Liskov and S. Micali. Amortized e-cash. In *Financial Cryptography 2001*, volume 2339 of *Lecture Notes in Computer Science*, pages 1–20. Springer Verlag, 2001.

18. T. Nakanishi, M. Shiota, and Y. Sugiyama. An efficient online electronic cash with unlinkable exact payments. In *Information Security Conference – ISC 2004*, volume 3225 of *Lecture Notes in Computer Science*, pages 367–378. Springer Verlag, 2004.

19. T. Nakanishi and Y. Sugiyama. Unlinkable divisible electronic cash. In *Information Security Workshop – ISW 2000*, volume 1975 of *Lecture Notes in Computer Science*, pages 121–134. Springer Verlag, 2000.

20. T. Okamoto and K. Ohta. Universal electronic cash. In *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages 324–337. Springer Verlag, 1992.

21. J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM Symposium on the Theory of Computing – STOC*, pages 387–394. ACM Press, 1990.

22. A. Sahai. Non-malleable non-interactive zero-knowledge and adaptive chosen-ciphertext security. In *40th IEEE Symposium on Foundations of Computer Science – FOCS*, pages 543–553. IEEE Computer Society Press, 1999.

23. T. Sander and A. Ta-Shma. Auditable, anonymous electronic cash. In *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 555–572. Springer Verlag, 1999.

24. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 566–598. Springer Verlag, 2001.

25. A. De Santis, G. Di Crescenzo, and G. Persiano. Necessary and sufficient assumptions for non-interactive zero-knowledge proofs of knowledge for all np relations. In *27th International Colloquium on Automata, Languages and Programming – ICALP*, volume 1853 of *Lecture Notes in Computer Science*, pages 451–462. Springer Verlag, 2000.

26. V. Varadharajan, K.Q. Nguyen, and Y. Mu. On the design of efficient RSA-based off-line electronic cash schemes. *Theoretical Computer Science*, 226:173–184, 1999.

27. V. Wei. More compact e-cash with efficient coin tracing. Cryptology ePrint Archive, Report 2005/411, 2005. `http://eprint.iacr.org/2005/411`.

A Definitions

A.1 Trapdoor Permutations

Definition 15 (Trapdoor Permutation Family). *A* trapdoor permutation family *is a tuple of probabilistic polynomial time Turing machines* $\mathcal{F} = ($Gen, Eval, Invert$)$ *such that:*

1. $\mathsf{Gen}(1^\kappa)$ *outputs a pair* (f, f^{-1}) *such that f is a permutation of* $\{0,1\}^\kappa$.
2. $\mathsf{Eval}(1^\kappa, f, x)$ *is a deterministic algorithm which on input f, where* $(f, f^{-1}) \in \mathsf{Gen}(1^\kappa)$, *and* $x \in \{0,1\}^\kappa$ *outputs* $y = f(x)$.
3. $\mathsf{Invert}(1^\kappa, f^{-1}, y)$ *is a deterministic algorithm which on input f^{-1}, where* $(f, f^{-1}) \in \mathsf{Gen}(1^\kappa)$, *and* $y \in \{0,1\}^\kappa$ *outputs some* $x = f^{-1}(y)$.
4. *For all κ,* $(f, f^{-1}) \in \mathsf{Gen}(1^\kappa)$, *and* $x \in \{0,1\}^\kappa$ *we have* $f^{-1}f(x) = x$.
5. *For all adversaries* $A \in \mathrm{PPT}^*$, *the following is negligible*

$$\Pr[(f, f^{-1}) \leftarrow \mathsf{Gen}(1^\kappa), \quad x \leftarrow \{0,1\}^\kappa, \quad A(f, f(x)) = f^{-1}(y)] .$$

A.2 Signature Schemes

A signature scheme $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sig}, \mathsf{Vf})$ is secure against chosen-message attacks, *CMA*-secure [15], if it is infeasible to produce valid message-signature pair for *any* message, even if the adversary has access to a signing oracle $\mathsf{Sig}_{\mathsf{sk}}(\cdot)$. Formally we use the following experiment for the definition

Experiment 6 (CMA, $\mathbf{Exp}^{\mathsf{cma}}_{\mathcal{SS},A}(\kappa)$).

$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Kg}(\kappa)$
$(m, \sigma) \leftarrow A^{\mathsf{Sig}_{\mathsf{sk}}(\cdot)}(\mathsf{pk})$
return $\mathsf{Vf}_{\mathsf{pk}}(m, s) = 1$

The advantage of the adversary is defined as

$$\mathbf{Adv}^{\mathsf{cma}}_{\mathcal{SS},A}(\kappa) = \Pr[\mathbf{Exp}^{\mathsf{cma}}_{\mathcal{SS},A}(\kappa) = 1] .$$

A signature scheme $\mathcal{SS}$ is *CMA*-secure if $\mathbf{Adv}^{\mathsf{cma}}_{\mathcal{SS},A}(\kappa)$ is negligible for all polynomial-time adversaries A.

A.3 Commitment Schemes

A (non-interactive) commitment scheme $\mathcal{COM} = (\mathsf{Commit}, \mathsf{Reveal})$ consists of two algorithms, the commitment algorithm and the reveal algorithm. The commitment algorithm takes as input a message $\mathsf{msg} \in \{0,1\}^\kappa$ and outputs a pair (c, r). The reveal algorithm takes a commitment c, a message msg, and the secret r and determines whether or not c is a commitment to msg under commitment secret r.

The following two experiments defines secrecy and binding of a commitment scheme.

Experiment 7 (Secrecy, $\mathbf{Exp}^{\text{secrecy}-b}_{\mathcal{COM},A}(\kappa)$).

  $(\text{msg}_0, \text{msg}_1, \text{state}) \leftarrow A(\text{choose}, 1^\kappa)$
  $(c, r) \leftarrow \text{Commit}(\text{msg}_b)$
  $d \leftarrow A(\text{guess}, c, \text{state})$
  return d

The advantage of an adversary A is

$$\mathbf{Adv}^{\text{secrecy}}_{\mathcal{COM},A}(\kappa) = |\Pr[\mathbf{Exp}^{\text{secrecy}-0}_{\mathcal{COM},A}(\kappa) = 1] - \Pr[\mathbf{Exp}^{\text{secrecy}-1}_{\mathcal{COM},A}(\kappa) = 1]| \ .$$

The commitment scheme $\mathcal{COM}$ has secrecy if $\mathbf{Adv}^{\text{secrecy}}_{\mathcal{COM},A}(\kappa)$ is negligible for any polynomial-time adversary A.

Experiment 8 (Binding, $\mathbf{Exp}^{\text{binding}}_{\mathcal{COM},A}(\kappa)$).

  $(c, r_0, \text{msg}_0, r_1, \text{msg}_1) \leftarrow A(\text{guess})$
  if $(\text{Reveal}(c, \text{msg}_0, r_0) = \text{Reveal}(c, \text{msg}_1, r_1) = 1) \wedge (\text{msg}_0 \neq \text{msg}_1)$ **then**
    return 1
  else
    return 0
  end if

The advantage of an adversary A is

$$\mathbf{Adv}^{\text{binding}}_{\mathcal{COM},A}(\kappa) = \Pr[\mathbf{Exp}^{\text{binding}}_{\mathcal{COM},A}(\kappa) = 1] \ .$$

The commitment scheme $\mathcal{COM}$ is binding if $\mathbf{Adv}^{\text{binding}}_{\mathcal{COM},A}(\kappa)$ is negligible for any polynomial-time adversary A.

One could give stronger definitions, but in our case the above experiments suffice. As an example, they do not rule out the existence of an adversary which, after seeing one commitment, creates another commitment to the same value, which can be opened after the first commitment has been opened.

It is known that secret and binding commitment schemes exist if there exists a family of one-way permutations [14]. The construction even gives a perfectly binding scheme, i.e., even an unbounded adversary cannot decommit to more than one value.

A.4 Indistinguishable Encryption Schemes

Informally an encryption scheme $\mathcal{CS} = (\text{Kg}, E, D)$ is called indistinguishable if it is infeasible to distinguish between the encryptions of two plaintexts of the same length. The experiment below formalizes this assumption.

Experiment 9 (Indistinguishability, $\mathbf{Exp}^{\text{ind}-b}_{\mathcal{CS},A}(\kappa)$).

  $(\text{msg}_0, \text{msg}_1, \text{state}) \leftarrow A(\text{choose}, 1^\kappa)$
  $c \leftarrow E(\text{msg}_b)$
  $d \leftarrow A(\text{guess}, c, \text{state})$
  return d

The advantage of an adversary A is

$$\mathbf{Adv}^{\mathsf{ind}}_{\mathcal{CS},A}(\kappa) = |\Pr[\mathbf{Exp}^{\mathsf{ind}-0}_{\mathcal{CS},A}(\kappa) = 1] - \Pr[\mathbf{Exp}^{\mathsf{ind}-1}_{\mathcal{CS},A}(\kappa) = 1]| \ .$$

The encryption scheme $\mathcal{CS}$ is indistinguishable if $\mathbf{Adv}^{\mathsf{ind}}_{\mathcal{COM},A}(\kappa)$ is negligible for any polynomial-time adversary A.

The notion of indistinguishability is equivalent to the well-known definition of semantic security, which informally says that no information about the plaintext can be efficiently computed from the cipher-text. We use the terms "indistinguishable encryption scheme" and "semantically secure encryption scheme" interchangeably in this text.

A.5 Proofs of Knowledge

Non-interactive zero-knowledge proofs (NIZK) were introduced by Blum, Feldman, and Micali [6]. Several works have since refined and extended the notion in various ways. Following [2] we employ the definition of adaptive zero-knowledge for NIZK introduced by Feige, Lapidot, and Shamir [11] and we use the notion of simulation soundness introduced by Sahai [22]. The notion of simulation soundness is strengthened by De Santis et al. [24]. In contrast to [2], the NIZK we use must be adaptive zero-knowledge for polynomially many statements, and not only for a single statement. The requirement on simulation soundness is in fact unchanged compared with [2], i.e. single statement simulation soundness suffices.

Definition 16 (NIPS). *A triple $(p(\kappa), P, V)$ is an efficient adaptive non-interactive proof system (NIPS) for a language $L \in \mathbf{NP}$ with witness relation R if $p(\kappa)$ is a polynomial and P and V are probabilistic polynomial time machines such that*

1. *Completeness. $(x, w) \in R$ and $\xi \in \{0, 1\}^{p(\kappa)}$ implies $V(x, P(x, w, \xi), \xi) = 1$.*
2. *Soundness. For all computable functions A, $\Pr_{\xi \in \{0,1\}^{p(\kappa)}}[A(\xi) = (x, \pi) \wedge x \notin L \wedge V(x, \pi, \xi) = 1]$ is negligible in κ.*

We suppress p in our notation of a NIPS and simply write (P, V).

Loosely speaking a non-interactive zero-knowledge proof system is a NIPS, which is also zero-knowledge, but there are several flavors of zero-knowledge. We need a NIZK which is adaptive zero-knowledge (for a single statement) in the sense of Feige, Lapidot, and Shamir [11].

Experiment 10 (Adaptive Indistinguishability, $\mathbf{Exp}^{\mathsf{ad-ind}-0}_{(P,V,S),A}(\kappa)$).

$\xi \leftarrow_R \{0, 1\}^{f(\kappa)}$
$(\mathsf{state}, x, w) \leftarrow A(\mathsf{setup}, \xi)$
while $(x, w) \in R$ **do**
   $(\mathsf{state}, x, w) \leftarrow A(\mathsf{choose}, P(x, w, \xi))$
end while
return $A(\mathsf{guess}, \mathsf{state})$

Experiment 11 (Adaptive Indistinguishability, $\mathbf{Exp}^{\mathsf{ad-ind-1}}_{(P,V,S),A}(\kappa)$).

$(\xi, \mathsf{simstate}) \leftarrow S(1^\kappa)$
$(\mathsf{state}, x, w) \leftarrow A(\mathsf{setup}, \xi)$
while $(x, w) \in R$ **do**
   $(\mathsf{state}, x, w) \leftarrow A(\mathsf{choose}, S(x, \xi))$
end while
return $A(\mathsf{guess}, \mathsf{state})$

The advantage in the experiment is defined

$$\mathbf{Adv}^{\mathsf{ad-ind}}_{(P,V,S),A}(\kappa) = |\Pr[\mathbf{Exp}^{\mathsf{ad-ind-0}}_{(P,V,S),A}(\kappa) = 1] - \Pr[\mathbf{Exp}^{\mathsf{ad-ind-1}}_{(P,V,S),A}(\kappa) = 1]|$$

and the notion of adaptive zero-knowledge is given below.

Definition 17 (Adaptive Zero-Knowledge (cf. [11])). *A NIPS (P, V) is adaptive zero-knowledge (NIZK) if there exists a polynomial time Turing machine S such that $\mathbf{Adv}^{\mathsf{ad-ind}}_{(P,V,S),A}(\kappa)$ is negligible for all $A \in \mathcal{A}$.*

In cryptographic proofs one often performs hypothetic experiments where the adversary is run with simulated NIZKs. If the experiment simulates NIZKs to the adversary, the adversary could potentially gain the power to compute valid proofs of false statements. For a simulation sound NIZK this is not possible.

Experiment 12 (Simulation Soundness, $\mathbf{Exp}^{\mathsf{sim-sound}}_{(P,V,S),A}(\kappa)$ (cf. [24])).

$(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$
$(x, \pi) \leftarrow A^{S(\mathsf{simulate}, \cdot, \xi, \mathsf{simstate})}(\mathsf{guess}, \xi)$
if $(\pi \notin \mathfrak{R}_S) \wedge (x \notin L) \wedge (V(x, \pi, \xi) = 1)$ **then**
   return 1
else
   return 0
end if

Definition 18 (Simulation Soundness (cf. [22,24])). *A NIZK (P, V) with polynomial time simulator S for a language L is unbounded simulation sound if*

$$\mathbf{Adv}^{\mathsf{sim-sound}}_{(P,V,S),A}(\kappa) = \mathbf{Exp}^{\mathsf{sim-sound}}_{(P,V,S),A}(\kappa)$$

is negligible for all $A \in \mathcal{A}$.

De Santis et al. [24] extend the results in [11] and [22] and prove the following result.

Theorem 2. *If there exists a family of trapdoor permutations, then there exists a simulation sound NIZK for any language in* **NP** *in the CRS-model.*

In the this paper we abbreviate "efficient non-interactive adaptive zero-knowledge unbounded simulation sound proof" by NIZK.

It is important to note that the above definitions do not require that it is possible to extract the witness, i.e., they are not proofs of knowledge. To our

knowledge, there are no results on the existence of simulation-sound proofs of knowledge, although signatures of knowledge [9] are similar.

One must be careful when defining the experiment for extractability. As for simulation-soundness, we want to give the adversary the ability to request simulated proofs for theorems of its choice, and if it outputs a valid proof, the extractor should be able to extract a witness. In the original definitions of NIZK proofs of knowledge [1,24], *soundness* and *validity*, i.e., the requirement on extractability, are two separate properties. Such a definition would be hard to use when designing protocols. In a protocol, we need to produce a single CRS which is used both for simulation and for extraction. Therefore it makes sense to combine the two properties in a single experiment.

Experiment 13 (Extractable Simulation Soundness, $\mathbf{Exp}^{\mathsf{ext-sim-sound}}_{(P,V,S),A}(\kappa)$ (cf. [24])).

$(\xi, \mathsf{simstate}) \leftarrow S(\mathsf{setup}, 1^\kappa)$
$(x, \pi) \leftarrow A^{S(\mathsf{simulate}, \cdot, \xi, \mathsf{simstate})}(\mathsf{guess}, \xi)$
$w \leftarrow S(\mathsf{extract}, x, \pi, \xi, \mathsf{simstate})$
if $(\pi \notin \mathfrak{R}_S) \wedge ((x, w) \notin R) \wedge (V(x, \pi, \xi) = 1)$ **then**
    return 1
else
    return 0
end if

Definition 19 (Extractable Simulation Soundness). *A NIZK (P, V) with polynomial time simulator S for a language L is unbounded extractable simulation sound if*
$$\mathbf{Adv}^{\mathsf{ext-sim-sound}}_{(P,V,S),A}(\kappa) = \mathbf{Exp}^{\mathsf{ext-sim-sound}}_{(P,V,S),A}(\kappa)$$
is negligible for all $A \in \mathcal{A}$.

We now give a construction of an extractable simulation sound proof system based on an unbounded simulation sound proof system. The idea behind the construction is the same as for [1] which is also used in [9], namely to encrypt the witness using a semantically secure encryption scheme where the public key is derived from the common reference string. Extraction is performed by letting the extractor choose the CRS in such a way that it knows the private key.

Let L be a language with witness relation R, i.e., $x \in L$ exactly when there exists w such that $(x, w) \in R$. We define a proof system (P, V) with simulator S and prove that it is an unbounded simulation sound zero-knowledge proof of knowledge. Note that S plays the role both of the simulator and the extractor.

In all the below experiments, we let the common reference string ξ consist of two parts, ξ_0 and ξ_1. We let pk be the public key defined by ξ_0 for the encryption scheme $\mathcal{CS} = (\mathsf{Kg}, E, Dec)$, and let sk be the corresponding secret key . We also let $L' = \{(x, c) \mid x \in L \wedge (x, D_{\mathsf{sk}}(c)) \in R$ with witness relation $R' = \{((x, c), (w, r)) \mid (x, w) \in R \wedge E_{\mathsf{pk},r}(w) = c\}$. Let (P', V') be a unbounded simulation sound proof system for L', and let S' be its simulator. .

Definition 20 (Prover $P(x, w)$).

$$(c, r) \leftarrow E_{\mathsf{pk}}(w)$$
$$\pi' \leftarrow P'((x, c), (w, r), \xi_1)$$
$$\pi \leftarrow (c, \pi')$$
return π

Definition 21 (Verifier $P(x, \pi)$).

Parse π as (c, π')
return $V'((x, c), \pi', \xi_1)$

Definition 22 (Simulator $S(\mathsf{tag}, \mathsf{params})$).

if $\mathsf{tag} = \mathsf{setup}$ **then**
 Parse params *as* 1^κ
 $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Kg}(1^\kappa)$
 $\xi_0 \leftarrow \mathsf{pk}$
 $(\xi_1, \mathsf{simstate}') \leftarrow S'(\mathsf{setup}, 1^\kappa)$
 $\mathsf{simstate} \leftarrow (\mathsf{sk}, \mathsf{simstate})$
 $\xi \leftarrow (\xi_0, \xi_1)$
 return $(\xi, \mathsf{simstate})$
else if $\mathsf{tag} = \mathsf{simulate}$ **then**
 Parse params *as* $(x, \xi, \mathsf{simstate})$
 Parse $\mathsf{simstate}$ *as* $(\mathsf{sk}, \mathsf{simstate}')$
 $c \leftarrow E_{\mathsf{pk}}(0)$
 $\pi' \leftarrow S'(\mathsf{simulate}, (x, c), \xi_1, \mathsf{simstate}')$
 $\pi \leftarrow (c, \pi')$
 return π
else if $\mathsf{tag} = \mathsf{extract}$ **then**
 Parse params *as* $(\pi, x, \xi, \mathsf{simstate})$
 Parse $\mathsf{simstate}$ *as* $(\mathsf{sk}, \mathsf{simstate}')$
 Parse π *as* (c, π')
 $w \leftarrow D_{\mathsf{sk}}(c)$
 return w
else
 return $\bot$
end if

We now prove the following theorem.

Theorem 3. *The proof system (P, V) is an unbounded simulation-sound non-interactive zero-knowledge proof of knowledge for the language L with witness relation R.*

Proof. We prove, in order, the properties adaptive indistinguishability and extractable simulation-soundness.

ADAPTIVE INDISTINGUISHABILITY Assume (P, V) is not adaptively indistinguishable. Let A be an adversary such that $\mathbf{Adv}^{\mathsf{ad-ind}}_{(P,V,S),A}(\kappa)$ is non-negligible. We will use A to construct $A_{\mathsf{ad-ind}}$, breaking either the adaptive indistinguishability of (P', V'), or $A_{\mathsf{sem-sec}}$, breaking the semantic security of $\mathcal{CS}$. Consider the

proof system $(\tilde{P}, \tilde{V})$, which is identical to (P, V) except that instead of outputting (c, π'), the prover outputs $(E_p k(0), \pi')$. If A wins the indistinguishability experiment in this case with non-negligible probability, then $A_{\mathsf{ad-ind}}$ can use A by prepending $E_{\mathsf{pk}}(0)$ to each query and distinguish between interaction with P' and the simulator.

If A does not win with non-negligible probability against $(\tilde{P}, \tilde{V})$, then it can be used by $A_{\mathsf{sem-sec}}$ in the following way. $A_{\mathsf{sem-sec}}$ chooses a theorem and a witness $(x, w) \in R$ and requests and encryption c of either w or 0. It runs $\mathbf{Exp}^{\mathsf{ad-ind-0}}_{(P,V,S),A}(\kappa)$ and $\mathbf{Exp}^{\mathsf{ad-ind-1}}_{(P,V,S),A}(\kappa)$ which the modification that P returns (c, π') instead of $(E_{\mathsf{pk}}(w), \pi')$. If A is able to distinguish between the experiments, then $A_{\mathsf{sem-sec}}$ concludes that it received an encryption of w, and otherwise that in received an encryption of 0.

Since, by assumption, (P', V') has adaptive indistinguishability and $\mathcal{CS}$ is semantically secure, the existence of either $A_{\mathsf{ad-ind}}$ or $A_{\mathsf{sem-sec}}$ with the above properties is a contradiction. Hence (P, V) has adaptive indistinguishability.

EXTRACTABLE SIMULATION-SOUNDNESS Assume (P, V) does not have extractable simulation-soundness, and let A be an adversary which wins in Experiment 13 with non-negligible probability. We describe how to construct an adversary $A_{\mathsf{sim-sound}}$ which breaks the simulation soundness of (P', V').

$A_{\mathsf{sim-sound}}$ runs A in Experiment 13, while taking part in Experiment 12 itself. It answers queries to S by algorithm in Definition 22, using its oracle S' where necessary. When A outputs (x, π) on the call $A(\mathsf{guess}, \xi)$, $A_{\mathsf{sim-sound}}$ parses π as (c, π') and outputs $((x, c), \pi')$ on its call $A_{\mathsf{sim-sound}}(\mathsf{guess}, \xi)$. If $w \leftarrow D_{\mathsf{sk}}(c)$ is not a witness of x, then $(x, c) \notin L'$. Thus $A_{\mathsf{sim-sound}}$ wins in its experiment exactly when A wins. Thus $A_{\mathsf{sim-sound}}$ breaks the simulation-soundness of (P', V'), which is a contradiction. We conclude that (P, V) is extractable simulation sound.

In this paper we write NIZK-PK for unbounded simulation sound non-interactive zero-knowledge proof of knowledge.