

The Ring of Gyges: Investigating the Future of Criminal Smart Contracts

Ari Juels
Cornell Tech, Jacobs Institute
IC3[†]
juels@cornell.edu

Ahmed Kosba
University of Maryland
akosba@cs.umd.edu

Elaine Shi
Cornell University
IC3[†]
rs2358@cornell.edu

[†]Initiative for CryptoCurrencies and Contracts

ABSTRACT

Thanks to their anonymity (pseudonymity) and elimination of trusted intermediaries, cryptocurrencies such as Bitcoin have created or stimulated growth in many businesses and communities. Unfortunately, some of these are criminal, e.g., money laundering, illicit marketplaces, and ransomware.

Next-generation cryptocurrencies such as Ethereum will include rich scripting languages in support of *smart contracts*, programs that autonomously intermediate transactions. In this paper, we explore the risk of smart contracts fueling new criminal ecosystems. Specifically, we show how what we call *criminal smart contracts* (CSCs) can facilitate leakage of confidential information, theft of cryptographic keys, and various real-world crimes (murder, arson, terrorism).

We show that CSCs for leakage of secrets (à la Wikileaks) are efficiently realizable in existing scripting languages such as that in Ethereum. We show that CSCs for theft of cryptographic keys can be achieved using primitives, such as Succinct Non-interactive ARGuments of Knowledge (SNARKs), that are already expressible in these languages and for which efficient supporting language extensions are anticipated. We show similarly that authenticated data feeds, an emerging feature of smart contract systems, can facilitate CSCs for real-world crimes (e.g., property crimes).

Our results highlight the urgency of creating policy and technical safeguards against CSCs in order to realize the promise of smart contracts for beneficial goals.

Keywords

Criminal Smart Contracts; Ethereum

The Ring of Gyges is a mythical magical artifact mentioned by the philosopher Plato in Book 2 of his Republic. It granted its owner the power to become invisible at will.—Wikipedia, “Ring of Gyges”

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS'16, October 24-28, 2016, Vienna, Austria

© 2016 ACM. ISBN 978-1-4503-4139-4/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2976749.2978362>

“[On wearing the ring,] no man would keep his hands off what was not his own when he could safely take what he liked out of the market, or go into houses and lie with anyone at his pleasure, or kill or release from prison whom he would...”—Plato, *The Republic*, Book 2 (2.360b) (trans. Benjamin Jowett)

1. INTRODUCTION

Cryptocurrencies such as Bitcoin remove the need for trusted third parties from basic monetary transactions and offer anonymous (more accurately, pseudonymous) transactions between individuals. While attractive for many applications, these features have a dark side. Bitcoin has stimulated the growth of ransomware [6], money laundering [38], and illicit commerce, as exemplified by the notorious Silk Road [31].

New cryptocurrencies such as Ethereum (as well as systems such as Counterparty [45] and SmartContract [1]) offer even richer functionality than Bitcoin. They support *smart contracts*, a generic term denoting programs written in Turing-complete cryptocurrency scripting languages. In a fully distributed system such as Ethereum, smart contracts enable general fair exchange (atomic swaps) without a trusted third party, and thus can effectively guarantee payment for successfully delivered data or services. Given the flexibility of such smart contract systems, it is to be expected that they will stimulate not just new beneficial services, but new forms of crime.

We refer to smart contracts that facilitate crimes in distributed smart contract systems as *criminal smart contracts* (CSCs). An example of a CSC is a smart contract for (private-)key theft. Such a CSC might pay a reward for (confidential) delivery of a target key sk , such as a certificate authority’s private digital signature key.

We explore the following key questions in this paper. *Could CSCs enable a wider range of significant new crimes than earlier cryptocurrencies (Bitcoin)? How practical will such new crimes be? And What key advantages do CSCs provide to criminals compared with conventional online systems?* Exploring these questions is essential to identifying threats and devising countermeasures.

1.1 CSC challenges

Would-be criminals face two basic challenges in the construction of CSCs. First, it is not immediately obvious whether a CSC is at all feasible for a given crime, such as key theft. This is because it is challenging to ensure

that a CSC achieves a key property in this paper that we call *commission-fair*, meaning informally that its execution guarantees *both* commission of a crime and commensurate payment for the perpetrator of the crime or *neither*. (We formally define commission-fairness for individual CSCs in the paper.) Fair exchange is necessary to ensure commission-fairness, but not sufficient: We show how CSC constructions implementing fair exchange still allow a party to a CSC to cheat. Correct construction of CSCs can thus be delicate.

Second, even if a CSC can in principle be constructed, given the limited opcodes in existing smart contract systems (such as Ethereum), it is not immediately clear that the CSC can be made *practical*. By this we mean that the CSC can be executed without unduly burdensome computational effort, which in some smart contract systems (e.g., Ethereum) would also mean unacceptably high execution fees levied against the CSC.

The following example illustrates these challenges.

EXAMPLE 1A (KEY COMPROMISE CONTRACT). Contractor \mathcal{C} posts a request for theft and delivery of the signing key sk_V of a victim certificate authority (CA) CertoMart. \mathcal{C} offers a reward $\$reward$ to a perpetrator \mathcal{P} for (confidentially) delivering the CertoMart private key sk_V to \mathcal{C} .

To ensure fair exchange of the key and reward in Bitcoin, \mathcal{C} and \mathcal{P} would need to use a trusted third party or communicate directly, raising the risks of being cheated or discovered by law enforcement. They could vet one another using a reputation system, but such systems are often infiltrated by law enforcement authorities [55]. In contrast, a decentralized smart contract can achieve self-enforcing fair exchange. For key theft, this is possible using the CSC Key-Theft in the following example:

EXAMPLE 1B (KEY COMPROMISE CSC). \mathcal{C} generates a private / public key pair (sk_C, pk_C) and initializes Key-Theft with public keys pk_C and pk_V (the CertoMart public key). Key-Theft awaits input from a claimed perpetrator \mathcal{P} of a pair (ct, π) , where π is a zero-knowledge proof that $ct = enc_{pk_C}[sk_V]$ is well-formed. Key-Theft then verifies π and upon success sends a reward of $\$reward$ to \mathcal{P} . The contractor \mathcal{C} can then download and decrypt ct to obtain the compromised key sk_V .

Key-Theft implements a fair exchange between \mathcal{C} and \mathcal{P} , paying a reward to \mathcal{P} if and only if \mathcal{P} delivers a valid key (as proven by π), eliminating the need for a trusted third party. But it is *not commission-fair*, as it does not ensure that sk_{vict} actually has value. The CertoMart can neutralize the contract by preemptively revoking its own certificate and then itself claiming \mathcal{C} 's reward $\$reward$!

As noted, a major thrust of this paper is showing how, for CSCs such as Key-Theft, criminals will be able to bypass such problems and still construct commission-fair CSCs. (For key compromise, it is necessary to enable contract cancellation should a key be revoked.) Additionally, we show that these CSCs can be efficiently realized using existing cryptocurrency tools or features currently envisioned for cryptocurrencies (e.g., zk-SNARKS [22]).

1.2 This paper

We show that it is or will be possible in smart contract systems to construct commission-fair CSCs for three types of crime:

1. Leakage / sale of secret documents;
2. Theft of private keys; and
3. “Calling-card” crimes, a broad class of physical-world crimes (murder, arson, etc.)

The fact that CSCs are possible in principle is not surprising. Previously, however, it was not clear how practical or extensively applicable CSCs might be. As our constructions for commission-fair CSCs show, constructing CSCs is not as straightforward as it might seem, but new cryptographic techniques and new approaches to smart contract design can render them feasible and even practical. Furthermore, criminals will undoubtedly devise CSCs beyond what this paper and the community in general are able to anticipate.

Our work therefore shows how imperative it is for the community to consider the construction of defenses against CSCs. Criminal activity committed under the guise of anonymity has posed a major impediment to adoption for Bitcoin. Yet there has been little discussion of criminal contracts in public forums on cryptocurrency [16] and the launch of Ethereum took place in July 2015. *It is only by recognizing CSCs early in their lifecycle that the community can develop timely countermeasures to them*, and see the promise of distributed smart contract systems fully realized.

While our focus is on preventing evil, happily the techniques we propose can also be used to create beneficial contracts. We explore both techniques for structuring CSCs and the use of cutting-edge cryptographic tools, e.g., Succinct Non-interactive Arguments of Knowledge (SNARKs), in CSCs. Like the design of beneficial smart contracts, CSC construction requires a careful combination of cryptography with commission-fair design [34].

In summary, our contributions are:

- *Criminal smart contracts:* We initiate the study of CSCs as enabled by Turing-complete scripting languages in next-generation cryptocurrencies. We explore CSCs for three different types of crimes: leakage of secrets in Section 4 (e.g., pre-release Hollywood films), key compromise / theft (of, e.g., a CA signing key) in Section 5, and “calling-card” crimes, such as assassination, that use data sources called “authenticated data feeds” (described below) in Section 6. We explore the challenges involved in crafting such criminal contracts and demonstrate (anticipate) new techniques to resist neutralization and achieve commission-fairness. We emphasize that because commission-fairness means informally that contracting parties obtain their “expected” utility, an application-specific metric, commission-fairness must be defined in a way specific to a given CSC. We thus formally specify commission-fairness for each of our CSC constructions in the online full version [42].
- *Proof of concept:* To demonstrate that even sophisticated CSC are realistic, we report (in their respective sections) on implementation of the CSCs we explore. Our CSC for leakage of secrets is *efficiently realizable today* in existing smart contract languages (e.g., that of Ethereum). Those for key theft and “calling-card” crimes rely respectively for efficiency and realizability on features currently envisioned by the cryptocurrency community.
- *Countermeasures:* We briefly discuss in Section 7 how our work in this paper can help prevent a proliferation of CSCs. Briefly, to be most effective, CSCs must be advertised, making them detectible given community vigilance. Miners have an economic incentive not to include

CSC transactions in blocks, as CSCs degrade the market value of a cryptocurrency. Consequently, awareness and robust detection strategies may offer an effective general defense. A key contribution of our work is to show the need for such countermeasures and stimulate exploration of their implementation in smart contract systems such as Ethereum.

We also briefly discuss in the online full version [42] how maturing technologies, such as hardware roots of trust (e.g., Intel SGX [40]) and program obfuscation can enrich the space of possible CSCs—as they can, of course, beneficial smart contracts.

2. BACKGROUND AND RELATED WORK

Emerging decentralized cryptocurrencies [52, 59] rely on a novel blockchain technology where miners reach consensus not only about *data*, but also about *computation*. Loosely speaking, the Bitcoin blockchain (i.e., miners) verifies transactions and stores a global *ledger*, which may be modeled as a piece of public memory whose integrity relies on correct execution of the underlying distributed consensus protocol. Bitcoin supports a limited range of programmable logic to be executed by the blockchain. Its scripting language is restrictive, however, and difficult to use, as demonstrated by previous efforts at building smart contract-like applications atop Bitcoin [23, 17, 7, 53, 46].

When the computation performed by the blockchain (i.e., miners) is generalized to arbitrary Turing-complete logic, we obtain a more powerful, general-purpose *smart contract* system. The first embodiment of such a decentralized smart contract system is the recently launched Ethereum [59]. Informally, a smart contract in such a system may be thought of as an autonomously executing piece of code whose inputs and outputs can include money. (We give more formalism below.) Hobbyists and companies are already building atop or forking off Ethereum to develop various smart contract applications such as security and derivatives trading [45], prediction markets [5], supply chain provenance [11], and crowd fund raising [2, 47].

Figure 1 shows the high-level architecture of a smart contract system instantiated over a decentralized cryptocurrency such as Bitcoin or Ethereum. When the underlying consensus protocol employed the cryptocurrency is secure, a majority of the miners (as measured by computational resources) are assumed to correctly execute the contract’s programmable logic.

Gas. Realistic instantiations of decentralized smart contract systems rely on *gas* to protect miners against denial-of-service attacks (e.g., running an unbounded contract). Gas is a form of transaction fee that is, roughly speaking, proportional to the runtime of a contract.

In this paper, although we do not explicitly express gas in our smart contract notation, we attempt to factor program logic away from the contract as an optimization when possible, to keep gas and thus transactional fees low. For example, some of the contracts we propose involve program logic executed on the user side, with no loss in security.

2.1 Smart contracts: the good and bad

Decentralized smart contracts have many beneficial uses, including the realization of a rich variety of new financial instruments. As Bitcoin does for transactions, in a decen-




Figure 1: **Schematic of a decentralized cryptocurrency system with smart contracts**, as illustrated by Delmolino et al. [34]. A smart contract’s state is stored on the public blockchain. A smart contract program is executed by a network of miners who reach consensus on the outcome of the execution, and update the contract’s state on the blockchain accordingly. Users can send money or data to a contract; or receive money or data from a contract.

tralized smart contract system, *the consensus system enforces autonomous execution of contracts*; no one entity or small set of entities can interfere with the execution of a contract. As contracts are self-enforcing, they eliminate the need for trusted intermediaries or reputation systems to reduce transactional risk. Decentralized smart contracts offer these advantages over traditional cryptocurrencies such as Bitcoin:

- *Fair exchange* between mutually distrustful parties with *rich contract rules expressible in a programmable logic*. This feature prevents parties from cheating by aborting an exchange protocol, yet removes the need for physical rendezvous and (potentially cheating) third-party intermediaries.
- *Minimized interaction* between parties, reducing opportunities for unwanted monitoring and tracking.
- *Enriched transactions with external state* by allowing as input *authenticated data feeds* (attestations) provided by brokers on physical and other events outside the smart-contract system, e.g., stock tickers, weather reports, etc. These are in their infancy in Ethereum, but their availability is growing and use of trusted hardware promises to stimulate their deployment [61].

Unfortunately, for all of their benefit, these properties have a dark side, potentially facilitating crime because:

- *Fair exchange* enables transactions between mutually distrustful criminal parties, eliminating the need for today’s fragile reputation systems and/or potentially cheating or law-enforcement-infiltrated third-party intermediaries [55, 39].
- *Minimized interaction* renders illegal activities harder for law enforcement to monitor. In some cases, as for the key-theft and calling-card CSCs we present, a criminal can set up a contract and walk away, allowing it to execute autonomously with no further interaction.
- *Enriched transactions with external state* broaden the scope

of possible CSCs to, e.g., physical crimes (terrorism, arson, murder, etc.).

As decentralized smart contract systems typically inherit the anonymity (pseudonymity) of Bitcoin, they offer similar secrecy for criminal activities. Broadly speaking, therefore, there is a risk that the capabilities enabled by decentralized smart contract systems will enable new underground ecosystems and communities.

2.2 Digital cash and crime

Bitcoin and smart contracts do not represent the earliest emergence of cryptocurrency. Anonymous e-cash was introduced in 1982 in a seminal paper by David Chaum [29]. Nacache and von Solms noted that anonymous currency would render “perfect crimes” such as kidnapping untraceable by law enforcement [57]. This observation prompted the design of fair blind signatures or “escrow” for e-cash [26, 58], which enables a trusted third party to link identities and payments. Such linkage is possible in classical e-cash schemes where a user identifies herself upon withdrawal of anonymous cash, but not pseudonymous cryptocurrencies such as Bitcoin.

Ransomware has appeared in the wild since 1989 [18]. A major cryptovirological [60] “improvement” to ransomware has been use of Bitcoin [44], thanks to which CryptoLocker ransomware has purportedly netted hundreds of millions of dollars in ransom [25]. Assassination markets using anonymous digital cash were first proposed in a 1995-6 essay entitled “Assassination Politics” [19].

There has been extensive study of Bitcoin-enabled crime, such as money laundering [51], Bitcoin theft [49], and illegal marketplaces such as the Silk Road [31]. Meiklejohn et al. [49] note that Bitcoin is pseudonymous and that mixes, mechanisms designed to confer anonymity on Bitcoins, do not operate on large volumes of currency and in general today it is hard for criminals to cash out anonymously in volume.

On the other hand, Ron and Shamir provide evidence that the FBI failed to locate most of the Bitcoin holdings of Dread Pirate Roberts (Ross Ulbricht), the operator of the Silk Road, even after seizing his laptop [56]. Möser, Böhome, and Breuker [51] find that they cannot successfully deanonymize transactions in two of three mixes under study, suggesting that the “Know-Your-Customer” principle, regulators’ main tool in combatting money laundering, may prove difficult to enforce in cryptocurrencies. Increasingly practical proposals to use NIZK proofs for anonymity in cryptocurrencies [20, 33, 50], and at least one currently in the early stages of commercial deployment [13], promise to make stronger anonymity available to criminals.

3. NOTATION AND THREAT MODEL

We adopt the formal blockchain model of Kosba et al. [43]. As background, we give a high-level description of that model in this section. We use the model to specify cryptographic protocols in our paper; these protocols encompass criminal smart contracts and corresponding user-side protocols.

Protocols in the smart contract model. Our model treats a *contract* as a special party that is *entrusted to enforce correctness but not privacy*, as noted above. (In reality, of course, a contract is enforced by the network.) All messages sent to the contract and its internal state are publicly visible. A contract interacts with users and other con-

tracts by exchanging messages (also referred to as transactions). Money, expressed in the form of account balances, is recorded in the global ledger (on the blockchain). Contracts can access and update the ledger to implement money transfers between users, who are represented by pseudonymous public keys.

3.1 Threat Model

We adopt the following threat model in this paper.

- *Blockchain: Trusted for correctness but not privacy.* We assume that the blockchain always correctly stores data and performs computations and is always available. The blockchain exposes all of its internal states to the public, however, and retains no private data.
- *Arbitrarily malicious contractual parties.* We assume that contractual parties are mutually distrustful, and they act solely to maximize their own benefit. Not only can they deviate arbitrarily from the prescribed protocol, they can also abort from the protocol prematurely.
- *Network influence of the adversary.* We assume that messages between the blockchain and players are delivered within a bounded delay, i.e., not permanently dropped. (A player can always resend a transaction dropped by a malicious miner.) In our model, an adversary immediately receives and can arbitrarily reorder messages, however. In real-life decentralized cryptocurrencies, the winning miner sets the order of message processing. An adversary may collude with certain miners or influence message-propagation among nodes. As we show in Section 5, for key-theft contracts, message-reordering enables a rushing attack that a commission-fair CSC must prevent.

The formal model we adopt (reviewed later in this section and described in full by Kosba et al. [43]) captures all of the above aspects of our threat model.

3.2 Security definitions

For a CSC to be commission-fair requires two things:

- *Correct definition of commission-fairness.* There is no universal formal definition of commission fairness: It is application-specific, as it depends on the goals of the criminal (and perpetrator). Thus, for each CSC, we specify in the online full version [42] a corresponding definition of commission-fairness by means of a UC-style ideal functionality that achieves it. Just specifying a correct ideal functionality is itself often challenging! We illustrate the challenge in Section 5 and the online full version [42] with a naive-key functionality that represents seemingly correct but in fact flawed key-theft contract.
- *Correct protocol implementation.* To prove that a CSC is commission-fair, we must show that its (real-world) protocol emulates the corresponding ideal functionality. We prove this for our described CSCs in the standard Universally Composable (UC) simulation paradigm [28] adopted in the cryptography literature, against arbitrarily malicious contractual counterparties as well as possible network adversaries. Our protocols are also secure against aborting adversaries, e.g., attempts to abort without paying the other party. Fairness in the presence of aborts is well known in general to be impossible in standard models of distributed computation [32]. Several recent works show that a blockchain that is correct, available, and aware of the progression of time can enforce financial fairness against aborting parties [23, 43, 17]. Specifi-

cally, when a contract lapses, the blockchain can cause the aborting party to lose a deposit to the honest parties.

3.3 Notational Conventions

We now explain some notational conventions for writing contracts. The online full version [42] gives a warm-up example.

- **Currency and ledger.** We use $\text{ledger}[\mathcal{P}]$ to denote party \mathcal{P} 's balance in the global ledger. For clarity, variables that begin with a \$ sign denote money, but otherwise behave like ordinary variables.
- **Pseudonymity.** Parties can use pseudonyms to obtain better anonymity. In particular, a party can generate arbitrarily many public keys. In our notational system, when we refer to a party \mathcal{P} , \mathcal{P} denotes the party's pseudonym. The formal blockchain model [43] we adopt provides a contract wrapper that manages the pseudonym generation and the message signing necessary for establishing an authenticated channel to the contract. These details are abstracted away from the main contract program.
- **Timer.** Time progresses in rounds. At the beginning of each round, the contract's **Timer** function will be invoked. The variable T encodes the current time.
- **Entry points and variable scope.** A contract can have various entry points, each of which is invoked when receiving a corresponding message type. Thus entry points behave like function calls invoked upon receipt of messages.

All variables are assumed to be globally scoped, with the following exception: When an entry point says "Upon receiving a message from *some* party \mathcal{P} ," this allows the registration of a new party \mathcal{P} . In general, contracts are open to any party who interacts with them. When a message is received from \mathcal{P} (without the keyword "some"), party \mathcal{P} denotes a fixed party – and a well-formed contract has already defined \mathcal{P} .

This notational system [43] is not only designed for convenience, but is also endowed with precise, formal meanings compatible with the Universal Composability framework [28]. We refer the reader to [43] for formal modeling details. While our proofs in the online full version [42] rely on this supporting formalism, the main body can be understood without it.

4. CSCS FOR LEAKAGE OF SECRETS

As a first example of the power of smart contracts, we show how an existing type of criminal contract deployed over Bitcoin can be made more robust and functionally enhanced as a smart contract and can be practically implemented in Ethereum.

Among the illicit practices stimulated by Bitcoin is payment-incentivized *leakage*, i.e., public disclosure, of secrets. The recently created web site Darkleaks [3] (a kind of subsidized Wikileaks) serves as a decentralized market for crowd-funded public leakage of a wide variety of secrets, including, "Hollywood movies, trade secrets, government secrets, propri-

etary source code, industrial designs like medicine or defence, [etc.]."

Intuitively, we define commission-fairness in this setting to mean that a contractor \mathcal{C} receives payment iff it leaks a secret in its entirety within a specified time limit. (See the online full version [42] for a formal definition.) As we show, Darkleaks highlights the inability of Bitcoin to support commission-fairness. We show how a CSC can in fact achieve commission-fairness with high probability.

4.1 Darkleaks

In the Darkleaks system, a contractor \mathcal{C} who wishes to sell a piece of content M partitions it into a sequence of n segments $\{m_i\}_{i=1}^n$. At a time (block height) T_{open} pre-specified by \mathcal{C} , a randomly selected subset $\Omega \subset [n]$ of k segments is publicly disclosed as a sample to entice donors / purchasers—those who will contribute to the purchase of M for public leakage. When \mathcal{C} determines that donors have collectively paid a sufficient price, \mathcal{C} decrypts the remaining segments for public release. The parameter triple (n, k, T_{open}) is set by \mathcal{C} (where $n = 100$ and $k = 20$ are recommended defaults).

To ensure a fair exchange of M for payment without direct interaction between parties, Darkleaks implements a (clever) protocol on top of the Bitcoin scripting language. The main idea is that for a given segment m_i of M that is not revealed as a sample in Ω , donors make payment to a Bitcoin account a_i with public key pk_i . The segment m_i is encrypted under a key $\kappa = H(\text{pk}_i)$ (where $H = \text{SHA-256}$). To spend its reward from account a_i , \mathcal{C} is forced by the Bitcoin transaction protocol to disclose pk_i ; thus the act of *spending the reward automatically enables the community to decrypt m_i* .

We give further details in the online full version [42].

Shortcomings and vulnerabilities. The Darkleaks protocol has three major shortcomings / vulnerabilities that appear to stem from fundamental functional limitations of Bitcoin's scripting language when constructing contracts without direct communication between parties. The first two undermine commission-fairness, while the third limits functionality.¹

1. *Delayed release:* \mathcal{C} can refrain from spending purchasers' / donors' payments and releasing unopened segments of M until after M loses value. E.g., \mathcal{C} could withhold segments of a film until after its release in theaters, of an industrial design until after it is produced, etc.

2. *Selective withholding:* \mathcal{C} can choose to forego payment for selected segments and not disclose them. For example, \mathcal{C} could leak and collect payment for all of a leaked film but the last few minutes (which, with high probability, will not appear in the sample Ω), significantly diminishing the value of leaked segments.

3. *Public leakage only:* Darkleaks can only serve to leak secrets *publicly*. It does not enable fair exchange for *private leakage*, i.e., for payment in exchange for a secret M encrypted under the public key of a purchaser \mathcal{P} .

¹That these limitations are fundamental is evidenced by calls for new, time-dependent opcodes. One example is CHECKLOCKTIMEVERIFY; apart from its many legitimate applications, proponents note that it can facilitate secret leakage as in Darkleaks [35].

Additionally, Darkleaks has a basic protocol flaw:

4. *Reward theft*: In the Darkleaks protocol, the Bitcoin private key sk_i corresponding to pk_i is derived from m_i ; specifically $\text{sk}_i = \text{SHA-256}(m_i)$. Thus, the source of M (e.g., the victimized owner of a leaked film) can derive sk_i and steal rewards received by \mathcal{C} . (Also, when \mathcal{C} claims a reward, a malicious node that receives the transaction can decrypt m_i , compute $\text{sk}_i = \text{SHA-256}(m_i)$, and potentially steal the reward by flooding the network with a competing transaction [36].)

This last problem is easily remedied by generating the set $\{\kappa_i\}_{i=1}^n$ of segment encryption keys pseudorandomly or randomly, which we do in our CSC designs.

Remark: In *any* protocol in which goods are represented by a random sample, not just Darkleaks, \mathcal{C} can insert a small number of valueless or duplicate segments into M . With non-negligible probability, these will not result in an invalid-looking sample Ω , so Ω necessarily provides only a *weak* guarantee of the global validity of M . The larger k and n , the smaller the risk of such attack. Formal analysis of human-verified proofs of this kind and/or ways of automating them is an interesting problem beyond the scope of this paper, but important in assessing end-to-end security in a CSC of this kind.

4.2 A generic public-leakage CSC

We now present a smart contract that realizes public leakage of secrets using blackbox cryptographic primitives. (We later present efficient realizations.) This contract overcomes limitation 1. of the Darkleaks protocol (delayed release) by enforcing disclosure of M at a pre-specified time T_{end} —or else immediately refunding buyers’ money. It addresses limitation 2. (selective withholding) by ensuring that M is revealed in an all-or-nothing manner. (We later explain how to achieve private leakage and overcome limitation 3.)

Again, we consider settings where \mathcal{C} aims to sell M for public release after revealing sample segments M^* .

Informal protocol description. Informally, the protocol involves the following steps:

- *Create contract*. A seller \mathcal{C} initializes a smart contract with the encryption of a randomly generated *master secret key msk*. The master secret key is used to generate (symmetric) encryption keys for the segments $\{m_i\}_{i=1}^n$. \mathcal{C} provides a cryptographic commitment $c_0 := \text{Enc}(\text{pk}, \text{msk}, r_0)$ of msk to the contract. (To meet the narrow technical requirements of our security proofs, the commitment is an encryption with randomness r_0 under a public key pk created during a trusted setup step.) The master secret key msk can be used to decrypt all leaked segments of M .
- *Upload encrypted data*. For each $i \in [n]$, \mathcal{C} generates encryption key $\kappa_i := \text{PRF}(\text{msk}, i)$, and encrypts the i -th segment as $\text{ct}_i = \text{enc}_{\kappa_i}[m_i]$. \mathcal{C} sends all encrypted segments $\{\text{ct}_i\}_{i \in [n]}$ to the contract (or, for efficiency, provides hashes of copies stored with a storage provider, e.g., a peer-to-peer network). Interested purchasers / donors can download the segments of M , but cannot decrypt them yet.
- *Challenge*. The contract generates a random challenge set $\Omega \subset [n]$, in practice today in Ethereum based on the hash of a recent block. Another future possibility is some well known randomness source, e.g., the NIST randomness

beacon [9], perhaps relayed through an authenticated data feed.

- *Response*. \mathcal{C} reveals the set $\{\kappa_i\}_{i \in \Omega}$ to the contract, and gives ZK proofs that the revealed secret keys $\{\kappa_i\}_{i \in \Omega}$ are generated correctly from the msk encrypted as c_0 .
- *Collect donations*. During a donation period, potential purchasers / donors can use the revealed secret keys $\{\kappa_i\}_{i \in \Omega}$ to decrypt the corresponding segments. If they like the decrypted segments, they can donate money to the contract as contribution for the leakage.
- *Accept*. If enough money has been collected, \mathcal{C} decommits msk for the contract (sends the randomness for the ciphertext along with msk). If the contract verifies the decommitment successfully, all donated money is paid to \mathcal{C} . The contract thus enforces a fair exchange of msk for money. (If the contract expires at time T_{end} without release of msk , all donations are refunded.)

The contract. Our proposed CSC PublicLeaks for implementing this public leakage protocol is given in Figure 2. The corresponding user side is as explained informally above (and inferable from the contract).

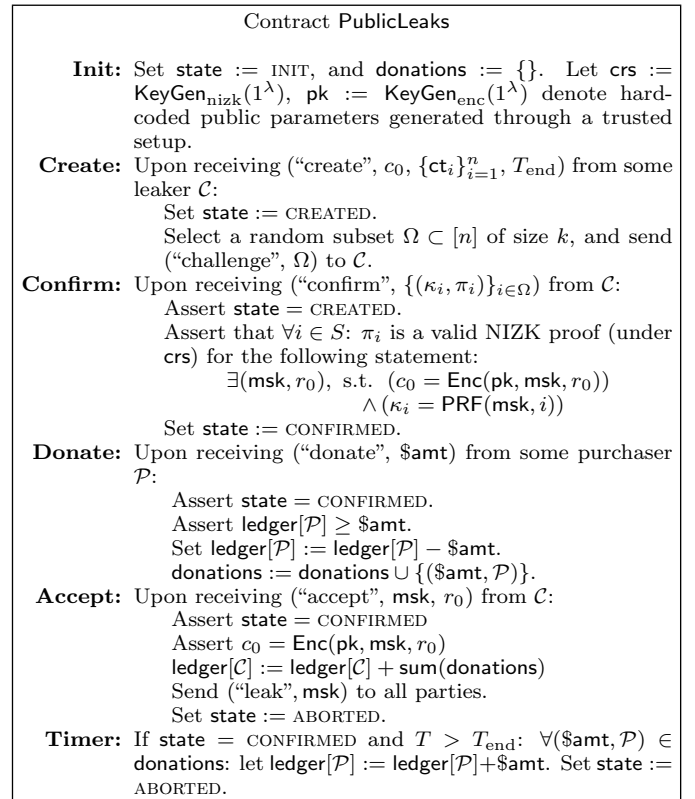


Figure 2: A contract PublicLeaks that leaks a secret M to the public in exchange for donations.

4.3 Commission-fairness: Formal definition and proof

In the online full version [42], we give a formal definition of commission-fairness for public leakage (explained informally above) as an ideal functionality. We also prove that PublicLeaks realizes this functionality assuming all revealed segments are valid—a property enforced with high (but not

overwhelming) probability by random sampling of M in `PublicLeaks`.

4.4 Optimizations and Ethereum implementation

The formally specified contract `PublicLeaks` uses generic cryptographic primitives in a black-box manner. We now give a practical, optimized version, relying on the random oracle model (ROM), that eliminates trusted setup, and also achieves better efficiency and easy integration with Ethereum [59].

A practical optimization. During contract creation, \mathcal{C} chooses random $\kappa_i \xleftarrow{\$} \{0, 1\}^\lambda$ for $i \in [n]$, and computes

$$c_0 := \{H(\kappa_1, 1), \dots, H(\kappa_n, n)\}.$$

The master secret key is simply $\text{msk} := \{\kappa_1, \dots, \kappa_n\}$, i.e., the set of hash pre-images. As in `PublicLeaks`, each segment m_i will still be encrypted as $\text{ct}_i := \text{enc}_{\kappa_i}[m_i]$. (For technical reasons—to achieve simulatability in the security proof—here $\text{enc}_{\kappa_i}[m_i] = m_i \oplus [H(\kappa_i, 1, \text{“enc”}) \parallel H(\kappa_i, 2, \text{“enc”}) \dots \parallel H(\kappa_i, z, \text{“enc”})]$ for suitably large z .)

\mathcal{C} submits c_0 to the smart contract. When challenged with the set Ω , \mathcal{C} reveals $\{\kappa_i\}_{i \in \Omega}$ to the contract, which then verifies its correctness by hashing and comparing with c_0 . To accept donations, \mathcal{C} reveals the entire msk .

This optimized scheme is asymptotically less efficient than our generic, black-box construction `PublicLeaks`—as the master secret key scales linearly in the number of segments n . But for typical, realistic document set sizes in practice (e.g., $n = 100$, as recommended for `Darkleaks`), it is more efficient.

Ethereum-based implementation. To demonstrate the feasibility of implementing leakage contracts using currently available technology, we implemented a version of the contract `PublicLeaks` atop Ethereum [59], using the `Serpent` contract language [10]. We specify the full implementation in detail in the online full version [42].

The version we implemented relies on the practical optimizations described above. As a technical matter, Ethereum does not appear at present to support timer-activated functions, so we implemented **Timer** in such a way that purchasers / donors make explicit withdrawals, rather than receiving automatic refunds.

This public leakage Ethereum contract is highly efficient, as it does not require expensive cryptographic operations. It mainly relies on hashing (SHA3-256) for random number generation and for verifying hash commitments. The total number of storage entries (needed for encryption keys) and hashing operations is $O(n)$, where, again, `Darkleaks` recommends $n = 100$. (A hash function call in practice takes a few micro-seconds, e.g., 3.92 μsecs measured on a core i7 processor.)

4.5 Extension: private leakage

As noted above, shortcoming 3. of `Darkleaks` is its inability to support *private* leakage, in which \mathcal{C} sells a secret exclusively to a purchaser \mathcal{P} . In the online full version [42], we show how `PublicLeaks` can be modified for this purpose. The basic idea is for \mathcal{C} not to reveal msk directly, but to provide a ciphertext $\text{ct} = \text{enc}_{\text{pk}_{\mathcal{P}}}[\text{msk}]$ on msk to the contract for a purchaser \mathcal{P} , along with a proof that ct is correctly formed. We describe a black-box variant whose security can be proven in essentially the same way as `PublicLeaks`. We also describe

Contract KeyTheft-Naive

Init: Set $\text{state} := \text{INIT}$. Let $\text{crs} := \text{KeyGen}_{\text{nizk}}(1^\lambda)$ denote a hard-coded NIZK common reference string generated during a trusted setup process.

Create: Upon receiving (“create”, $\$reward, \text{pk}_{\mathcal{V}}, T_{\text{end}}$) from some contractor $\mathcal{C} := (\text{pk}_{\mathcal{C}}, \dots)$:
 Assert $\text{state} = \text{INIT}$.
 Assert $\text{ledger}[\mathcal{C}] \geq \$reward$.
 $\text{ledger}[\mathcal{C}] := \text{ledger}[\mathcal{C}] - \$reward$.
 Set $\text{state} := \text{CREATED}$.

Claim: Upon receiving (“claim”, ct, π) from some purported perpetrator \mathcal{P} :
 Assert $\text{state} = \text{CREATED}$.
 Assert that π is a valid NIZK proof (under crs) for the following statement:
 $\exists r, \text{sk}_{\mathcal{V}}$ s.t. $\text{ct} = \text{Enc}(\text{pk}_{\mathcal{C}}, (\text{sk}_{\mathcal{V}}, \mathcal{P}), r)$
 and $\text{match}(\text{pk}_{\mathcal{V}}, \text{sk}_{\mathcal{V}}) = \text{true}$
 $\text{ledger}[\mathcal{P}] := \text{ledger}[\mathcal{P}] + \$reward$.
 Set $\text{state} := \text{CLAIMED}$.

Timer: If $\text{state} = \text{CREATED}$ and current time $T > T_{\text{end}}$:
 $\text{ledger}[\mathcal{C}] := \text{ledger}[\mathcal{C}] + \$reward$
 $\text{state} := \text{ABORTED}$

Figure 3: A naïve, flawed key theft contract (lacking commission-fairness)

a practical variant that combines a *verifiable random function* (VRF) of Chaum and Pedersen [30] (for generation of $\{\kappa_i\}_{i=1}^n$) with a *verifiable encryption* (VE) scheme of Camenisch and Shoup [27] (to prove correctness of ct). This variant can be deployed today using beta support for big number arithmetic in Ethereum.

5. A KEY-COMPROMISE CSC

Example 1b in the paper introduction described a CSC that rewards a perpetrator \mathcal{P} for delivering to \mathcal{C} the stolen key $\text{sk}_{\mathcal{V}}$ of a victim \mathcal{V} —in this case a certificate authority (CA) with public key $\text{pk}_{\mathcal{V}}$. Recall that \mathcal{C} generates a private / public key encryption pair $(\text{sk}_{\mathcal{C}}, \text{pk}_{\mathcal{C}})$. The contract accepts as a claim by \mathcal{P} a pair (ct, π) . It sends reward $\$reward$ to \mathcal{P} if π is a valid proof that $\text{ct} = \text{enc}_{\text{pk}_{\mathcal{C}}}[\text{sk}_{\mathcal{V}}]$ and $\text{sk}_{\mathcal{V}}$ is the private key corresponding to $\text{pk}_{\mathcal{V}}$.

Intuitively, a key-theft contract is commission-fair if it rewards a perpetrator \mathcal{P} for delivery of a private key that: (1) \mathcal{P} was responsible for stealing and (2) Is valid for a substantial period of time. (See the online full version [42] for a formal definition.)

This form of contract can be used to solicit theft of any type of private key, e.g., the signing key of a CA, the private key for a SSL/TLS certificate, a PGP private key, etc. (Similar contracts could solicit abuse, but not full compromise of a private key, e.g., forged certificates.)

Figure 3 shows the contract of Example 1b in our notation for smart contracts. We let crs here denote a common reference string for a NIZK scheme and $\text{match}(\text{pk}_{\mathcal{V}}, \text{sk}_{\mathcal{V}})$ denote an algorithm that verifies whether $\text{sk}_{\mathcal{V}}$ is the corresponding private key for some public key $\text{pk}_{\mathcal{V}}$ in a target public-key cryptosystem.

As noted above, this CSC is *not* commission-fair. Thus we refer to it as `KeyTheft-Naive`. We use `KeyTheft-Naive` as a helpful starting point for motivating and understanding the construction of a commission-fair contract proposed later, called `KeyTheft`.

5.1 Flaws in KeyTheft-Naive

The contract KeyTheft-Naive fails to achieve commission-fairness due to two shortcomings.

Revoke-and-claim attack. The CA \mathcal{V} can revoke the key $\text{sk}_{\mathcal{V}}$ and then itself submit the key for payment. The CA then not only negates the value of the contract but actually profits from it! This *revoke-and-claim* attack demonstrates that KeyTheft-Naive is not commission-fair in the sense of ensuring the delivery of a *usable* private key $\text{sk}_{\mathcal{V}}$.

Rushing attack. Another attack is a rushing attack. As noted in Section 3, an adversary can arbitrarily reorder messages—a reflection of possible attacks against the network layer in a cryptocurrency. (See also the formal blockchain model [43].) Thus, given a valid claim from perpetrator \mathcal{P} , a corrupt \mathcal{C} can decrypt and learn $\text{sk}_{\mathcal{V}}$, construct another valid-looking claim of its own, and make its own claim arrive before the valid one.

5.2 Fixing flaws in KeyTheft-Naive

We now show how to modify KeyTheft-Naive to prevent the above two attacks and achieve commission-fairness.

Thwarting revoke-and-claim attacks. In a revoke-and-claim attack against KeyTheft-Naive, \mathcal{V} preemptively revokes its public key $\text{pk}_{\mathcal{V}}$ and replaces it with a fresh one $\text{pk}'_{\mathcal{V}}$. As noted above, the victim can then play the role of perpetrator \mathcal{P} , submit $\text{sk}_{\mathcal{V}}$ to the contract and claim the reward. The result is that \mathcal{C} pays $\text{\$reward}$ to \mathcal{V} and obtains a stale key.

We address this problem by adding to the contract a feature called *reward truncation*, whereby the contract accepts evidence of revocation Π_{revoke} .

This evidence Π_{revoke} can be an Online Certificate Status Protocol (OCSP) response indicating that $\text{pk}_{\mathcal{V}}$ is no longer valid, a new certificate for \mathcal{V} that was unknown at the time of contract creation (and thus not stored in **Contract**), or a certificate revocation list (CRL) containing the certificate with $\text{pk}_{\mathcal{V}}$.

\mathcal{C} could submit Π_{revoke} , but to minimize interaction by \mathcal{C} , KeyTheft could provide a reward $\text{\$smallreward}$ to a third-party submitter. The reward could be small, as Π_{revoke} would be easy for ordinary users to obtain.

The contract then provides a reward based on the interval of time over which the key $\text{sk}_{\mathcal{V}}$ remains valid. Let T_{claim} denote the time at which the key $\text{sk}_{\mathcal{V}}$ is provided and T_{end} be an expiration time for the contract (which must not exceed the expiration of the certificate containing the targeted key). Let T_{revoke} be the time at which Π_{revoke} is presented ($T_{\text{revoke}} = \infty$ if no revocation happens prior to T_{end}). Then the contract assigns to \mathcal{P} a reward of $f(\text{reward}, t)$, where $t = \min(T_{\text{end}}, T_{\text{revoke}}) - T_{\text{claim}}$.

We do not explore choices of f here. We note, however, that given that a CA key $\text{sk}_{\mathcal{V}}$ can be used to forge certificates for rapid use in, e.g., malware or falsified software updates, much of its value can be realized in a short interval of time which we denote by δ . (A slant toward up-front realization of the value of exploits is common in general [24].) A suitable choice of reward function should be front-loaded and rapidly decaying. A natural, simple choice with this property is

$$f(\text{\$reward}, t) = \begin{cases} 0 & : t < \delta \\ \text{\$reward}(1 - ae^{-b(t-\delta)}) & : t \geq \delta \end{cases}$$

for $a < 1/2$ and some positive real value b . Note that a majority of the reward is paid provided that $t \geq \delta$.

Thwarting rushing attacks. To thwart rushing attacks, we separate the claim into two phases. In the first phase, \mathcal{P} expresses an intent to claim by submitting a commitment of the real claim message. \mathcal{P} then waits for the next round to open the commitment and reveal the claim message. (Due to technical subtleties in the proof, the commitment must be *adaptively secure*; in the proof, the simulator must be able to simulate a commitment without knowing the string s being committed to, and later, be able to claim the commitment to any string s .) In real-life decentralized cryptocurrencies, \mathcal{P} can potentially wait multiple block intervals before opening the commitment, to have higher confidence that the blockchain will not fork. In our formalism, one round can correspond to one or more block intervals.

Figure 4 gives a key theft contract KeyTheft that thwarts revoke-and-claim and the rushing attacks.

5.3 Target and state exposure

An undesirable property of KeyTheft-Naive is that its target / victim and state are publicly visible. \mathcal{V} can thus learn whether it is the target of KeyTheft-Naive. \mathcal{V} also observes successful claims—i.e., whether $\text{sk}_{\mathcal{V}}$ has been stolen—and can thus take informed defensive action. For example, as key revocation is expensive and time-consuming, \mathcal{V} might wait until a successful claim occurs and only then perform a revoke-and-claim attack.

To limit target and state exposure, we note two possible enhancements to KeyTheft. The first is a *multi-target* contract, in which key theft is requested for any one of a set of multiple victims. The second is what we call *cover claims*, false claims that conceal any true claim. Our implementation of KeyTheft, as specified in Figure 4, is a multi-target contract, as this technique provides both partial target and partial state concealment.

Multi-target contract. A multi-target contract solicits the private key of any of m potential victims $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_m$. There are many settings in which the private keys of different victims are of similar value. For example, a multi-target contract KeyTheft could offer a reward for the private key $\text{sk}_{\mathcal{V}}$ of *any* CA able to issue SSL/TLS certificates trusted by, e.g., Internet Explorer (of which there are more than 650 [37]).

A challenge here is that the contract state is public, thus the contract must be able to verify the proof for a valid claim (private key) $\text{sk}_{\mathcal{V}_i}$ *without knowing which key was furnished, i.e., without learning i* . Our implementation shows that constructing such proofs as zk-SNARKs is practical. (The contractor \mathcal{C} itself can easily learn i by decrypting $\text{sk}_{\mathcal{V}_i}$, generating $\text{pk}_{\mathcal{V}_i}$, and identifying the corresponding victim.)

Cover claims. As the state of a contract is publicly visible, a victim \mathcal{V} learns whether or not a successful claim has been submitted to KeyTheft-Naive. This is particularly problematic in the case of single-target contracts.

Rather than sending the NIZK proof π with **ct**, it is possible instead to delay submission of π (and payment of the reward) until T_{end} . (That is, **Claim** takes as input (“claim”, **ct**.) This approach conceals the validity of **ct**. Note that even without π , \mathcal{C} can still make use of **ct**.

A contract that supports such concealment can also sup-

Contract KeyTheft	
Init:	Set $\text{state} := \text{INIT}$. Let $\text{crs} := \text{KeyGen}_{\text{nizk}}(1^\lambda)$ denote a hard-coded NIZK common reference string generated during a trusted setup process.
Create:	Same as in Contract KeyTheft-Naive (Figure 3), except that an additional parameter ΔT is additionally submitted by \mathcal{C} .
Intent:	Upon receiving (“intent”, cm) from some purported perpetrator \mathcal{P} : Assert $\text{state} = \text{CREATED}$ Assert that \mathcal{P} has not sent “intent” earlier Store cm, \mathcal{P}
Claim:	Upon receiving (“claim”, ct, π, r) from \mathcal{P} : Assert $\text{state} = \text{CREATED}$ Assert \mathcal{P} submitted (“intent”, cm) earlier such that $\text{cm} = \text{comm}(\text{ct} \pi, r)$. Continue in the same manner as in contract KeyTheft-Naive, except that the ledger update $\text{ledger}[\mathcal{P}] := \text{ledger}[\mathcal{P}] + \reward does not take place immediately.
Revoke:	On receive (“revoke”, Π_{revoke}) from some \mathcal{R} : Assert Π_{revoke} is valid, and $\text{state} \neq \text{ABORTED}$. $\text{ledger}[\mathcal{R}] := \text{ledger}[\mathcal{R}] + \smallreward . If $\text{state} = \text{CLAIMED}$: Let $t :=$ (time elapsed since successful Claim). Let $\mathcal{P} :=$ (successful claimer). $\text{reward}_{\mathcal{P}} := f(\$ \text{reward}, t)$. $\text{ledger}[\mathcal{P}] := \text{ledger}[\mathcal{P}] + \text{reward}_{\mathcal{P}}$. Else, $\text{reward}_{\mathcal{P}} := 0$ $\text{ledger}[\mathcal{C}] := \text{ledger}[\mathcal{C}] + \$\text{reward} - \$\text{smallreward} - \text{reward}_{\mathcal{P}}$ Set $\text{state} := \text{ABORTED}$.
Timer:	If $\text{state} = \text{CLAIMED}$ and at least ΔT time elapsed since Claim : $\text{ledger}[\mathcal{P}] := \text{ledger}[\mathcal{P}] + \reward ; Set $\text{state} := \text{ABORTED}$. Else if current time $T > T_{\text{end}}$ and $\text{state} \neq \text{ABORTED}$: $\text{ledger}[\mathcal{C}] := \text{ledger}[\mathcal{C}] + \reward . Set $\text{state} := \text{ABORTED}$. <i>// \mathcal{P} should not submit claims after $T_{\text{end}} - \Delta T$.</i>

Figure 4: Key compromise CSC that thwarts the revoke-and-claim attack and the rushing attack.

port an idea that we refer to as *cover claims*. A cover claim is an *invalid* claim of the form (“claim”, ct), i.e., one in which ct is not a valid encryption of sk_V . Cover claims may be submitted by \mathcal{C} to conceal the true state of the contract. So that \mathcal{C} need not interact with the contract after creation, the contract could parcel out small rewards at time T_{end} to third parties that submit cover claims. We do not implement cover claims in our version of KeyTheft nor include them in Figure 4.

5.4 Commision-fairness: Formal definition and proof

We define commission-fairness for key theft in terms of an ideal functionality in the online full version [42] and also provide a formal proof of security there for KeyTheft.

5.5 Implementation

We rely on zk-SNARKs for efficient realization of the protocols above. zk-SNARKs are zero-knowledge proofs of knowledge that are succinct and very efficient to verify. zk-SNARKs have weaker security than what is needed in UC-style simulation proofs. We therefore use a generic transformation described in the Hawk work [43] to lift security such that the zero-knowledge proof ensures *simulation-extractable soundness*. (In brief, a one-time key generation phase is

1-Target		#threads	RSA-2048	ECDSA_P256
Key Gen. [\mathcal{C}]		1	124.88 sec	242.30 sec
	Eval. Key	4	33.53 sec	73.38 sec
	Ver. Key		215.93 MB	448.24 MB
Prove [\mathcal{P}]		1	41.02 sec	83.63 sec
	Proof	4	15.7 sec	32.19 sec
			711 B	711 B
Verification [Contract]			0.0089 sec	0.0087 sec
500-Target		#threads	RSA-2048	ECDSA_P256
Key Gen. [\mathcal{C}]		1	161.56 sec	263.07 sec
	Eval. Key	4	43.35 sec	78.31 sec
	Ver. Key		279.41 MB	490.85 MB
Prove [\mathcal{P}]		1	54.15 sec	84.69 sec
	Proof	4	23.54 sec	33.49 sec
			711 B	711 B
Verification [Contract]			0.0087 sec	0.0087 sec

Table 1: Performance of the key-compromise zk-SNARK circuit for **Claim** in the case of a 1-target and 500-target contracts. [\cdot] refers to the entity performing the computational work.

needed to generate two keys: a public evaluation key, and a public verification key. To prove a certain NP statement, an untrusted prover uses the evaluation key to compute a succinct proof; any verifier can use the public verification key to verify the proof. The verifier in our case is the contract.) In our implementation, we assume the key generation is executed confidentially by a trusted party; otherwise a prover can produce a valid proof for a false statement. To minimize trust in the key generation phase, secure multi-party computation techniques can be used as in [21].

zk-SNARK circuits for Claim. To estimate the proof computation and verification costs required for **Claim**, we implemented the above protocol for theft of RSA-2048 and ECDSA_P256 keys, which are widely used in SSL/TLS certificates currently. The circuit has two main sub-circuits: a key-check circuit, and an encryption circuit.² The encryption circuit was realized using RSAES-OAEP [41] with a 2048-bit key. Relying on compilers for high-level implementation of these algorithms may produce expensive circuits for the zk-SNARK proof computation. Instead, we built customized circuit generators that produce more efficient circuits. We then used the state-of-the-art zk-SNARK library [22] to obtain the evaluation results. Table 1 shows the results of the evaluation of the circuits for both single-target and multi-target contracts. The experiments were conducted on an Amazon EC2 r3.2xlarge instance with 61GB of memory and 2.5 GHz processors.

The results yield two interesting observations: i) Once a perpetrator obtains the secret key of a TLS public key, computing the zk-SNARK proof would require less than two minutes, costing less than 1 USD [4] for either single or multi-target contracts; ii) The overhead introduced by using a multi-target contract with 500 keys on the prover’s side is only 13 seconds in the worst case. In the same time, the verification overhead by the contract is still the same as in the single-target case. This is achieved by the use of an efficient Merkle tree circuit that proves the membership of

²The circuit also has other commitment and encryption sub-circuits needed for simulation extractability – see the online full version [42].

the compromised public key in the target key set, while using the same components of the single-target circuit as is.

Validation of revoked certificates. The reward function in the contract above relies on certificate revocation time, and therefore the contract needs modules that can process certificate revocation proofs, such as CRLs and OCSP responses, and verify the CA digital signatures on them. As an example, we measured the running time of `openssl verify -crl_check` command, testing the revoked certificate at [12] and the CRL last updated at [8] on Feb 15th, 2016, that had a size of 143KB. On average, the verification executed in about 0.016 seconds on a 2.3 GHz i7 processor. The signature algorithm was SHA-256 with RSA encryption, with a 2048-bit key. Since OCSP responses can be smaller than CRLs, the verification time could be even less for OCSP.

The case of multi-target contracts. Verifying the revocation proof for single-target contracts is straightforward: The contract can determine whether a revocation proof corresponds to the targeted key. In multi-target contracts, though, the contract does not know which target key corresponds to the proof of key theft \mathcal{P} submitted. Thus, a proof is needed that the revocation corresponds to the stolen key, and it must be submitted by \mathcal{C} .

We built a zk-SNARK circuit through which \mathcal{C} can prove the connection between the ciphertext submitted by the perpetrator and the compromised target key. For efficiency, we eliminated the need for the key-check sub-circuit in **Revoke** by forcing \mathcal{P} to append the secret index of the compromised public key to the secret key before applying encryption in **Claim**. The evaluation in Table 2 illustrates the efficiency of the verification done by the contract receiving the proof, and the practicality for \mathcal{C} of constructing the proof. In contrast to the case for **Claim**, the one-time key generation for this circuit must be done independently from \mathcal{C} , so that \mathcal{C} cannot cheat the contract. We note that the **Revoke** circuit we built is invariant to the cryptosystem of the target keys.

	<i>#threads</i>	RSA-2048	ECDsa_P256
Key Gen.	1	124.64 sec	124.35 sec
	4	33.52 sec	33.38 sec
Eval. Key		215.41 MB	214.81 MB
Ver. Key		5.51 KB	4.88 KB
Prove [\mathcal{C}]	1	41.08 sec	40.96 sec
	4	15.94 sec	15.59 sec
Proof		711 B	711 B
Verification [Contract]		0.0087 sec	0.0086 sec

Table 2: Performance of the key-compromise zk-SNARK circuit for **Revoke** needed in the case of multi-target contract. [\cdot] refers to the entity performing the computational work.

6. CALLING-CARD CRIMES

As noted above, decentralized smart contract systems (e.g., Ethereum) have supporting services that provide authenticated data feeds, digitally signed attestations to news, facts about the physical world, etc. While still in its infancy, this powerful capability is fundamental to many applications of smart contracts and will expand the range of CSCs very broadly to encompass events in the physical world, as in the following example:

EXAMPLE 2 (ASSASSINATION CSC). *Contractor \mathcal{C} posts a contract **Assassinate** for the assassination of Senator X . The contract rewards the perpetrator \mathcal{P} of this crime.*

*The contract **Assassinate** takes as input from a perpetrator \mathcal{P} a commitment **vcc** specifying in advance the details (day, time, and place) of the assassination. To claim the reward, \mathcal{P} decommits **vcc** after the assassination. To verify \mathcal{P} 's claim, **Assassinate** searches an authenticated data feed on current events to confirm the assassination of Senator X with details matching **vcc**.*

This example also illustrates the use of what we refer to as a *calling card*, denoted **cc**. A calling card is an unpredictable feature of a to-be-executed crime (e.g., in Example 2, a day, time, and place). Calling cards, alongside authenticated data feeds, can support a *general framework for a wide variety of CSCs*.

A generic construction for a CSC based on a calling card is as follows. \mathcal{P} provides a commitment **vcc** to a calling card **cc** to a contract in advance. After the commission of the crime, \mathcal{P} proves that **cc** corresponds to **vcc** (e.g., decommits **vcc**). The contract refers to some trustworthy and authenticated data feed to verify that: (1) The crime was committed and (2) The calling card **cc** matches the crime. If both conditions are met, the contract pays a reward to \mathcal{P} .

Intuitively, we define commission fairness to mean that \mathcal{P} receives a reward iff it was responsible for carrying out a commissioned crime. (A formal definition is given in the online full version [42].)

In more detail, let **CC** be a set of possible calling cards and $\text{cc} \in \text{CC}$ denote a calling card. As noted above, it is anticipated that an ecosystem of authenticated data feeds will arise around smart contract systems such as Ethereum. We model a data feed as a sequence of pairs from a source \mathcal{S} , where $(s(t), \sigma(t))$ is the emission for time t . The value $s(t) \in \{0, 1\}^*$ here is a piece of data released at time t , while $\sigma(t)$ is a corresponding digital signature; \mathcal{S} has an associated private / public key pair $(\text{sk}_{\mathcal{S}}, \text{pk}_{\mathcal{S}})$ used to sign / verify signatures.

Note that once created, a calling-card contract requires *no further interaction from \mathcal{C}* , making it hard for law enforcement to trace \mathcal{C} using subsequent network traffic.

6.1 Example: website defacement contract

As an example, we specify a simple CSC **SiteDeface** for website defacement. The contractor \mathcal{C} specifies a website url to be hacked and a statement **stmt** to be displayed. (For example, **stmt** = "Anonymous. We are Legion. We do not Forgive..." and **url** = whitehouse.gov.)

We assume a data feed that authenticates website content, i.e., $s(t) = (w, \text{url}, t)$, where w is a representation of the webpage content and t is a timestamp, denoted for simplicity in contract time. (For efficiency, w might be a hash of and pointer to the page content.) Such a feed might take the form of, e.g., a digitally signed version of an archive of hacked websites (e.g., zone-h.com). The function **SigVer** denotes the signature verification operation.

As example parameterization, we might let $\text{CC} = \{0, 1\}^{256}$, i.e., **cc** is a 256-bit string. A perpetrator \mathcal{P} simply selects a calling card $\text{cc} \xleftarrow{\$} \{0, 1\}^{256}$ and commitment $\text{vcc} := \text{commit}(\text{cc}, \mathcal{P}; \rho)$, where **commit** denotes a commitment scheme, and $\rho \in \{0, 1\}^{256}$ a random string. (In practice, HMAC-SHA256 is a suitable choice for easy implementation in Ethereum, given its support for SHA-256.) \mathcal{P} decommits by revealing all arguments to **commit**.

The CSC **SiteDeface** is shown in Figure 5. The example we use is simplified for clarity. We assume in this example

that the published webpage will only contain the calling card and the statement, but it is possible to support arbitrarily rich content in the published webpage.

<p>Contract SiteDeface</p> <p>Init: On receiving $(\\$reward, pk_S, url, stmt)$ from some \mathcal{C}:</p> <p style="padding-left: 2em;">Store $(\\$reward, pk_S, url, stmt)$</p> <p style="padding-left: 2em;">Set $i := 0, T_{start} := T$</p> <p>Commit: Upon receiving commitment vcc from some \mathcal{P}:</p> <p style="padding-left: 2em;">Store $vcc_i := vcc$ and $P_i := \mathcal{P}; i := i + 1$.</p> <p>Claim: Upon receiving as input a tuple (cc, ρ, σ, w, t) from some \mathcal{P}:</p> <p style="padding-left: 2em;">Find smallest i such that $vcc_i = \text{commit}(cc, \mathcal{P}; \rho)$, abort if not found.</p> <p style="padding-left: 2em;">Assert $w = cc \parallel stmt$</p> <p style="padding-left: 2em;">Assert $t \geq T_{start}$</p> <p style="padding-left: 2em;">Assert $\text{SigVer}(pk_S, (w, url, t), \sigma) = \text{true}$</p> <p style="padding-left: 2em;">Send $\\$reward$ to P_i and abort.</p>
--

Figure 5: CSC for website defacement

Remarks. **SiteDeface** could be implemented alternatively by having \mathcal{P} generate cc as a digital signature. Our implementation, however, also accommodates short, low-entropy calling cards cc , which is important for general calling-card CSCs. See the online full version [42].

Implementation. Given an authenticated data feed, implementing **SiteDeface** would be straightforward and efficient. The main overhead lies in the **Claim** module, where the contract computes a couple of hashes and validates the feed signature on retrieved website data. As noted in Section 4, a hash function call can be computed in very short time ($4\mu\text{sec}$), while checking the signature would be more costly. For example, if the retrieved content is 100KB, the contract needs only about 10msec to verify an RSA-2048 signature.

6.2 Commission-fairness: Formal definition

We give a formal definition of commission-fairness for a general calling-card CSC in the online full version [42].

We do not provide a security proof, as this would require modeling of physical-world systems, which is outside the scope of this paper.

6.3 Other calling-card crimes / data feed corruption

Using a CSC much like **SiteDeface**, a contractor \mathcal{C} can solicit many other crimes, e.g., assassination, assault, sabotage, hijacking, kidnapping, denial-of-service attacks, and terrorist attacks. A perpetrator \mathcal{P} must be able to designate a calling card that is reliably reported by an authenticated data feed.

A natural countermeasure to a calling-card-based CSC, however, is for an adversary, i.e., party trying to neutralize the CSC, to corrupt an authenticated data feed or data source such that it furnishes invalid data. In cases involve highly dangerous CSCs, such approaches might be workable. For example, if there is a well-funded CSC calling for the assassination of a public official, a data feed provider or news source might be persuaded to issue a false report of the target’s death.

Conversely, if \mathcal{C} is concerned about suppression of information in one source, it can of course create a CSC that references multiple sources, e.g., multiple news feeds. The

CSC might treat an event as authentic only if it is reported by a certain fraction, e.g., a majority, of these sources. Such diversification of sources would render corruption far more challenging (and is a good idea to provide resilience to data corruption for benign contracts too).

We discuss a number of the general issues surrounding the construction of calling-card CSCs in the online full version [42].

7. COUNTERMEASURES

The main aim of our work is to emphasize the importance of developing countermeasures against CSCs for emerging smart contract systems such as Ethereum. We briefly discuss this challenge here.

Ideas such as blacklisting “tainted” coins / transactions—those with known criminal use—have been brought forward for cryptocurrencies such as Bitcoin. A proactive alternative noted in Section 2 is an identity-escrow idea in early (centralized) e-cash systems sometimes referred as “trustee-based tracing” [26, 58]. Trustee-tracing schemes permitted a trusted party (“trustee”) or a quorum of such parties to trace monetary transactions that would otherwise remain anonymous. In decentralized cryptocurrencies, however, users do not register identities with authorities—and many would object to doing so. It would be possible for users to register voluntarily and to choose only to accept only currency they deem suitably registered. The idea of tainting coins, though, has been poorly received by the cryptocurrency community because it undermines the basic cash-like property of fungibility [14, 48], and trustee-based tracing would have a similar drawback. It is also unclear what entities should be granted the authority to perform blacklisting or register users.

We observe, however, that it is economically advantageous for most users of a cryptocurrency to monitor and/or restrain criminal activity, which can degrade acceptance and therefore market value. This observation has stimulated the creation, for instance, of the Blockchain Alliance [15], whose mission is to combat criminal activity on blockchains. Similarly, core developers of cryptocurrencies such as Ethereum have indicated the desirability of filtering blockchain content, by analogy with censorship of hate speech [54].

Identifying Bitcoin transactions as criminal is challenging, as transactions themselves carry no information about payment context. In contrast, CSCs, if identified as such (e.g., an assassination contract), are self-incriminating objects. Reversing CSC binaries could be challenging, but we note that for CSCs to be effective—as in our examples above—their deployers must advertise, drawing attention to the nature of their contracts. (For example, a contractor looking to have an assassination performed must find an assassin.) Our hypothesis, therefore, is that a sufficiently vigilant cryptocurrency community can detect the presence of many CSCs and will be incentivized to filter or purge associated transactions. One simple potential mechanism is for miners to omit transactions from blocks when they are flagged by reputable communities as CSCs.

A more aggressive approach is possible as well, a notion that we call *trustee-neutralizable smart contracts*. A smart contract system might be designed such that an authority, quorum of authorities, or suitable set of general system participants is empowered to remove a contract from the blockchain. Such an approach would have a big advantage over traditional trustee-based protections, in that

it would not require users to register identities. Whether the idea would be palatable to cryptocurrency communities and whether a broadly acceptable set of authorities could be identified are, of course, open questions, as are the right supporting technical mechanisms.

In general, our work here is therefore important in *sensitizing the cryptocurrency community to the threat of CSCs*, enabling structures for monitoring and appropriate countermeasures to be set in place. We believe it is important to create awareness in the early stages of development of decentralized smart contract ecosystems such as Ethereum.

8. CONCLUSION

We have demonstrated that a range of commission-fair *criminal smart contracts* (CSCs) are practical in decentralized currencies with smart contracts. We presented three—leakage of secrets, key theft, and calling-card crimes—and showed that they are efficiently implementable with existing cryptographic techniques, given suitable support in smart contract systems such as Ethereum. The contract **PublicLeaks** and its private variant can today be efficiently implemented in **Serpent**, an Ethereum scripting language. **KeyTheft** would require only modest, already envisioned opcode support for zk-SNARKs for efficient deployment. Calling-card CSCs will be possible given a sufficiently rich data-feed ecosystem. Many more CSCs are no doubt possible.

We emphasize that smart contracts in distributed cryptocurrencies have numerous promising, legitimate applications and that banning smart contracts would be neither sensible nor, in all likelihood, possible. The urgent open question raised by our work is thus how to create safeguards against the most dangerous abuses of such smart contracts while supporting their many beneficial applications.

Acknowledgements

This work is funded in part by NSF grants CNS-1314857, CNS-1330599, CNS-1453634, CNS-1514261, CNS-1514163, CNS-1518765, a Packard Fellowship, a Sloan Fellowship, Google Faculty Research Awards, and a VMWare Research Award. We gratefully acknowledge Rafael Pass, abhi she-lat, Jonathan Katz, and Andrew Miller for helpful technical discussions regarding the cryptographic constructions and formalism. We especially thank Rafael Pass for pointing out the rushing attack.

9. REFERENCES

- [1] <http://www.smartcontract.com>.
- [2] <http://koinify.com>.
- [3] <https://github.com/darkwallet/darkleaks>.
- [4] Amazon EC2 pricing. <http://aws.amazon.com/ec2/pricing/>.
- [5] Augur. <http://www.augur.net/>.
- [6] Bitcoin ransomware now spreading via spam campaigns. <http://www.coindesk.com/bitcoin-ransomware-now-spreading-via-spam-campaigns/>.
- [7] bitoinj. <https://bitoinj.github.io/>.
- [8] CRL issued bby Symantec Class 3 EV SSL CA - G3. <http://ss.symcb.com/sr.crl>.
- [9] NIST randomness beacon. <https://beacon.nist.gov/home>.
- [10] Serpent. <https://github.com/ethereum/wiki/wiki/Serpent>.
- [11] Skuchain. <http://www.skuchain.com/>.
- [12] Verisign revoked certificate test page. <https://test-sspev.verisign.com>.
- [13] Zcash. Referenced Aug. 2016 at z.cash.
- [14] Mt. Gox thinks it's the Fed. freezes acc based on "tainted" coins. (unlocked now). <https://bitcointalk.org/index.php?topic=73385.0>, 2012.
- [15] Blockchain Alliance. www.blockchainalliance.org, 2016.
- [16] Ethereum and evil. Forum post at Reddit; url: <http://tinyurl.com/k8awj2j>, Accessed May 2015.
- [17] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek. Secure Multiparty Computations on Bitcoin. In *S & P*, 2013.
- [18] J. Bates. Trojan horse: AIDS information introductory diskette version 2.0. In E. Wilding and F. Skulason, editors, *Virus Bulletin*, pages 3–6. 1990.
- [19] J. Bell. Assassination politics. <http://www.outpost-of-freedom.com/jimbella.htm>, 1995-6.
- [20] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *S & P*. IEEE, 2014.
- [21] E. Ben-Sasson, A. Chiesa, M. Green, E. Tromer, and M. Virza. Secure sampling of public parameters for succinct zero knowledge proofs. In *S & P*, 2015.
- [22] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. Succinct non-interactive zero knowledge for a von Neumann architecture. In *USENIX Security*, 2014.
- [23] I. Bentov and R. Kumaresan. How to Use Bitcoin to Design Fair Protocols. In *CRYPTO*, 2014.
- [24] L. Bilge and T. Dumitras. Before we knew it: an empirical study of zero-day attacks in the real world. In *CCS*, 2012.
- [25] V. Blue. Cryptolocker's crimewave: A trail of millions in laundered Bitcoin. *ZDNet*, 22 December 2013.
- [26] E. F. Brickell, P. Gemmel, and D. W. Kravitz. Trustee-based tracing extensions to anonymous cash and the making of anonymous change. In *SODA*, volume 95, pages 457–466, 1995.
- [27] J. Camenisch and V. Shoup. Practical verifiable encryption and decryption of discrete logarithms. In *CRYPTO '03*. 2003.
- [28] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
- [29] D. Chaum. Blind signatures for untraceable payments. In *CRYPTO*, pages 199–203, 1983.
- [30] D. Chaum and T. P. Pedersen. Wallet databases with observers. In *CRYPTO '92*, pages 89–105, 1993.
- [31] N. Christin. Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In *WWW*, 2013.
- [32] R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *STOC*, 1986.
- [33] G. Danezis, C. Fournet, M. Kohlweiss, and B. Parno. Pinocchio Coin: building Zerocoin from a succinct pairing-based proof system. In *PETShop*, 2013.
- [34] K. Delmolino, M. Arnett, A. Kosba, A. Miller, and E. Shi. Step by step towards creating a safe smart contract: Lessons and insights from a cryptocurrency lab. <https://eprint.iacr.org/2015/460>.
- [35] P. T. et al. Darkwallet on twitter: "DARK LEAKS coming soon. <http://t.co/k4ubs16scr>". Reddit: <http://bit.ly/1A9USHy>.
- [36] I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *FC*, 2014.
- [37] E. F. Foundation. EFF SSL observatory. URL: <https://www.eff.org/observatory>, August 2010.
- [38] A. Greenberg. 'Dark Wallet' is about to make Bitcoin money laundering easier than ever. <http://www.wired.com/2014/04/dark-wallet/>.
- [39] A. Greenberg. Alleged silk road boss Ross Ulbricht now accused of six murders-for-hire, denied bail. *Forbes*, 21 November 2013.

- [40] Intel. Intel software guard extensions programming reference. Whitepaper ref. 329298-002US, October 2014.
- [41] J. Jonsson and B. Kaliski. Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specifications Version 2.1, 2003. RFC 3447.
- [42] A. Juels, A. Kosba, and E. Shi. The ring of gyges: Investigating the future of criminal smart contracts. Cryptology ePrint Archive, Report 2016/358, 2016. <http://eprint.iacr.org/2016/358>.
- [43] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *S & P*. IEEE, 2016.
- [44] V. Kotov and M. Rajpal. Understanding crypto-ransomware. Bromium whitepaper, 2014.
- [45] A. Krellenstein, R. Dermody, and O. Slama. Counterparty announcement. <https://bitcointalk.org/index.php?topic=395761.0>, January 2014.
- [46] R. Kumaresan and I. Bentov. How to Use Bitcoin to Incentivize Correct Computations. In *CCS*, 2014.
- [47] D. Mark, V. Zamfir, and E. G. Sirer. A call for a temporary moratorium on “The DAO” (v0.3.2). Referenced Aug. 2016 at <http://bit.ly/2aWDhyY>, 30 May 2016.
- [48] J. Matonis. Why Bitcoin fungibility is essential. *CoinDesk*, 1 Dec. 2013.
- [49] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage. A fistful of bitcoins: characterizing payments among men with no names. In *IMC*, 2013.
- [50] I. Miers, C. Garman, M. Green, and A. D. Rubin. Zerocoin: Anonymous Distributed E-Cash from Bitcoin. In *S & P*, 2013.
- [51] M. Moser, R. Bohme, and D. Breuker. An inquiry into money laundering tools in the bitcoin ecosystem. In *eCRS*, 2013.
- [52] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <http://bitcoin.org/bitcoin.pdf>, 2009.
- [53] R. Pass and a. shelat. Micropayments for peer-to-peer currencies. Manuscript.
- [54] M. Peck. Ethereum developer explores the dark side of Bitcoin-inspired technology. *IEEE Spectrum*, 19 May 2016.
- [55] K. Poulsen. Cybercrime supersite ‘DarkMarket’ was FBI sting, documents confirm. *Wired*, 13 Oct. 2008.
- [56] D. Ron and A. Shamir. How did Dread Pirate Roberts acquire and protect his bitcoin wealth? In *FC*. 2014.
- [57] S. V. Solms and D. Naccache. On blind signatures and perfect crimes. *Computers Security*, 11(6):581–583, 1992.
- [58] M. Stadler, J.-M. Piveteau, and J. Camenisch. Fair blind signatures. In *Eurocrypt*, pages 209–219, 1995.
- [59] G. Wood. Ethereum: A secure decentralized transaction ledger. <http://gavwood.com/paper.pdf>, 2014.
- [60] A. Young and M. Yung. Cryptovirology: Extortion-based security threats and countermeasures. In *S & P*, 1996.
- [61] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi. Town Crier: An authenticated data feed for smart contracts. In *ACM CCS*, 2016. (To appear.).