

Filling the Tax Gap via Programmable Money

Dimitris Karakostas
University of Edinburgh
dimitris.karakostas@ed.ac.uk

Aggelos Kiayias
University of Edinburgh and IOHK
akiayias@inf.ed.ac.uk

July 27, 2021

Abstract

In this work, we discuss the problem of facilitating tax auditing assuming “programmable money”, i.e., digital monetary instruments that are managed by an underlying distributed ledger. We explore how a taxation authority can verify the declared returns of its citizens and create a counter-incentive to tax evasion by two distinct mechanisms. First, we describe a design which enables auditing it as a built-in feature with minimal changes on the underlying ledger’s consensus protocol. Second, we offer an application-layer extension, which requires no modification in the underlying ledger’s design. Both solutions provide a high level of privacy, ensuring that, apart from specific limited data given to the taxation authority, no additional information — beyond the information already published on the underlying ledger — is leaked.

1 Introduction

A tax gap [11] is a difference between the reported and the real tax revenue, for a given jurisdiction and period of time. Research estimated that the tax gap in the USA was 16.4% of revenue owed [17] between 2008-2010, the total loss throughout the EU due to the tax gap to €151.5 in 2015 [12], while $\frac{1}{3}$ of taxpayers in the UK under-report their earnings [1] (albeit half of UK’s lost taxes are product of a small, wealthy fraction of misbehaving taxpayers). Therefore, reducing the tax gaps can significantly enhance the efforts of tax-collecting authorities.

Central bank digital currencies (CBDC) have also come to prominence in recent years. In the past decade, distributed ledger-based financial systems, which were kick-started with the creation of Bitcoin [13], were accompanied by the increasing digitalization of payments [5]. CBDCs are the culmination of these trends, enabling fast, cheap, and safe transactions in fiat assets. Crucially though, although still mostly on a research stage,¹ CBDCs have caused great concerns on citizens regarding transaction privacy [4].

¹<https://cbdctracker.org> [July 2021]

Our work offers two mechanisms that facilitate tax auditing and the identification of tax gaps in distributed ledger-based currency systems. The first is a wrapper around a generic distributed ledger, which enables taxpayers to declare their assets directly to the authorities, while undeclared assets are frozen. The second is a proof mechanism that enables the sender of some assets to prove, in a privacy-preserving manner, whether the transferred assets have been taxed. Both mechanisms are examples of programmable money (also referred to as smart money [3]), where currency is programmed to be transferable under a suitable set of circumstances or its transfer has specific implications.

1.1 Desiderata

In distributed ledger-based currency systems, a user \mathcal{U} manages their assets via addresses. Each address α is associated with a key pair (sk, vk) , such that the private key sk is used to claim ownership of the assets, e.g., by signing special messages; typically $\alpha = H(vk)$ for some hash function H . Each address α is associated with a (public) balance $\text{bal}(\alpha)$ so, given a list $[\alpha_1, \dots, \alpha_n]$ of all addresses that \mathcal{U} controls, \mathcal{U} 's total assets are $\Theta = \sum_{i=1}^n \text{bal}(\alpha_i)$. Our goal will be to retain as much privacy as possible, so Θ should be the only information that is leaked to \mathcal{T} , without de-anonymization of individual transaction data.

To showcase the limitations of current systems, consider the following example. Assume that Alice tax evades, i.e., creates a secret, undeclared address α and controls some assets θ in it. Given the pseudonymous nature of the ledger, α cannot be correlated with Alice, until she uses it. Following, Alice issues a transaction τ which sends θ assets from α to Bob. If Bob suspects that Alice evaded taxation for these θ assets, they might want to report her to the authorities for inspection. However, the complaint should be accompanied by a proof that α is controlled by Alice, i.e., a proof that Alice knows the private key associated with α . This is necessary as \mathcal{T} needs to distinguish between two scenarios: i) Alice controls α and tax evades; ii) Bob is lying about Alice owning α . In the first scenario, Bob *does* know that α is controlled by Alice, but τ is not sufficient to prove it. Instead, Bob needs a proof which can only be supplied by Alice, e.g., a signature from Alice which acknowledges τ or α . However, if Alice tax evades, naturally she would not create such incriminating proof.

It is important that we retain as many good features of existing ledger systems as possible. The most notable such feature is transaction privacy, thus our work considers pseudonymous, Bitcoin-like levels of privacy, and minimizes the information leaked to the authorities during a tax auditing. Another important aspect is the mechanism's performance. A fundamental ingredient of payment systems is the seamless transaction experience, so it is important to allow users to transact at all times, while also avoiding significant strain during taxation periods. Finally, our mechanisms aim to minimize the amount of (additional) published data, since storage in distributed ledgers is particularly costly.

In summary, the desiderata of our mechanisms are as follows:

- *Tax gap identification and counterincentive:* Tax evasion, i.e., failure of a user \mathcal{U} to declare the amount of assets they own, should be either detectable by a tax authority \mathcal{T} , with access to the ledger, or render the assets unusable.
- *High level of privacy:* \mathcal{T} should — at most — learn the total amount of

assets owned by each taxpayer at the end of a fiscal year; this information should be leaked only to \mathcal{T} and no additional information should be leaked to any other party, apart from the information already published on the ledger.

- *Unobstructed operation*: The introduction of a taxation mechanism should not result in any period during which the — tax compliant — users are prohibited from transacting.
- *Low performance overhead*: The taxation mechanism should not introduce a major performance overhead, in terms of computation and storage requirements from the users and the taxation authority.
- *Balanced load*: The computation and storage overhead of taxation should be spread over a period of time, rather than introduce performance spikes.

1.2 Related work

Literature offers various works on auditing of distributed ledger-based assets. A holistic approach is taken in zkLedger [14], which combines a permissioned ledger with zero-knowledge proofs to create a tamper-resistant, verifiable ledger of transactions. PRCash [19] also employs a permissioned ledger and offers a regulation mechanism that restricts the total amount of assets a user can receive anonymously for a period of time. Also Garman *et al.* [10] propose an anonymous ledger, which can enforce specific transaction policies. In our paper, Section 2 aims at offering a simpler design, which can be more easily integrated in existing pseudonymous distributed ledgers, compared to the aforementioned works. Another interesting research thread considers proofs of solvency. The first such scheme for Bitcoin exchanges, proposed by Maxwell [18], leaks the total amount of both assets and liabilities of the exchange; more importantly, it enables an attack that allows the exchange to hide assets, as detailed by in Zeroledge [7], which also proposed a privacy-preserving system that allows exchanges to prove properties about their holdings. Provisions [6] is a zero-knowledge proof of solvency mechanism for Bitcoin exchanges, based on Sigma protocols i.e., without the need to reveal the addresses or the amount of assets that an exchange controls. Similarly, Agrawal *et al.* [2] describe a proof of solvency which achieves better performance compared to Provisions, although assuming a trusted setup. The mechanism of Section 3 extends Provisions and is also applicable to [2].

2 Tax Auditable Distributed Ledger

In this section we describe a ledger with a built-in tax auditing mechanism. Our design is generic, such that existing ledgers can incorporate it with minimal changes in the underlying consensus protocol. An *auditable ledger* enforces a user \mathcal{U} to declare the amount of assets they own to a taxation authority \mathcal{T} , with failure to do so rendering the assets unusable. We achieve this while leaking to \mathcal{T} only the total amount of assets that \mathcal{U} owns at a specific point in time, e.g., the end of a fiscal year. We note that we consider only pseudonymous ledgers, so potentially de-anonymizable data may be published on the ledger, e.g., addresses which may be linked to the user who controls them.

We assume that \mathcal{T} holds a list of all taxpayers and is identified by a key $(sk_{\mathcal{T}}, vk_{\mathcal{T}})$. Also there exist taxation periods, which last for a pre-specified amount of time d . For example, a taxation period may last 1 calendar year, at the end of which taxpayers need to declare their assets to the authorities.

The core idea is that assets unaccounted for, at the end of the taxation period, are frozen, until their owners declare them to the authority. Specifically, at the end of a taxation period, all assets are frozen. To unfreeze an asset, a taxpayer \mathcal{U} declares it to \mathcal{T} as follows.

First, \mathcal{U} creates a new key pair $(sk_{\mathcal{U}}, vk_{\mathcal{U}})$ and the corresponding address $\alpha_{\mathcal{U}}$ and sends $\alpha_{\mathcal{U}}$ to \mathcal{T} as part of a KYC process. Next, \mathcal{T} certifies $\alpha_{\mathcal{U}}$ by issuing the signature $\sigma = \text{Sign}(\alpha_{\mathcal{U}}, sk_{\mathcal{T}})$, which it gives to \mathcal{U} . The tuple $\alpha_{\mathcal{U}}^t = \langle \alpha_{\mathcal{U}}, \sigma \rangle$ is the *certified address*, which is used by the user to transact with frozen assets. \mathcal{T} maintains a mapping of taxpayers and certified addresses, i.e., for every taxpayer \mathcal{U} it holds a list $A_{\mathcal{U}}$ of all certified taxation addresses that \mathcal{U} requested.

A transaction $\tau = \langle \alpha_s, \alpha_d, \Theta \rangle$, which moves Θ frozen assets from an address α_s , is valid only if $\alpha_d = \langle \alpha, \sigma \rangle \wedge \text{Verify}(\alpha, \sigma, vk_{\mathcal{T}}) = 1$. Consequently, miners accept transactions that unfreeze assets only as long as said assets are transferred to a certified address. Therefore, \mathcal{T} can compute the amount of \mathcal{U} 's assets as $\Theta_{\mathcal{U}} := \sum_{i=1}^n \text{bal}(\alpha_{\mathcal{U}}[i])$, n being the total number of \mathcal{U} 's certified addresses.

We note that the system can accommodate multiple taxation authorities from different countries. In that case, \mathcal{T} is a federation of authorities, each identified by a single key. Each authority's key is published on the ledger and a taxpayer can certify their addresses and declare their assets to the respective authorities.

Naturally, this mechanism introduces some challenges. Although standard pay-to-public-key-hash addresses are 25 bytes, certified addresses may be significantly larger, due to the certification signature of \mathcal{T} . For instance, ECDSA signatures in the DER format result in 72 additional bytes, thus making certified addresses 99 bytes long. Nevertheless, certified addresses are expected to be used only once, to declare the assets, thus the overall storage cost should not be significant. Another important consideration regards to the private state of the taxation authority; given the statute of limitations, \mathcal{T} might need to maintain its taxation private key and the mapping of certified addresses for a significant period, possibly resulting in significant maintenance costs.

We showcase our design via an auditable variation of Bitcoin ledger, denoted as \mathcal{L}^t . \mathcal{L}^t is initially parameterized by the public key of the authority $(sk_{\mathcal{T}}, vk_{\mathcal{T}})$, which is part of the ledger's genesis block. During the execution, \mathcal{T} can update its key by simply signing a new key $vk'_{\mathcal{T}}$ with $sk_{\mathcal{T}}$ and publishing it on the ledger. A taxation period lasts 52560 blocks, i.e., roughly 1 calendar year, so block 52560 and its multiples are "tax-auditing" blocks. When a tax-auditing block is issued, all assets on \mathcal{L}^t which are controlled by non-certified addresses are frozen. To transact with assets from a frozen address, a user sends them to a certified address, as described above.

Freezing complicates the system in a number of ways. First, the liveness of a transaction [9] may be affected. For instance, a transaction which spends from a non-certified address will be rejected, if it is created before but published after a tax-auditing block. We sidestep this issue by enabling users to use certified addresses before the freezing period, hence the liveness guarantees of the ledger apply unconditionally on certified addresses. Second, it is possible that multiple competing tax-auditing blocks are created, e.g., multiple blocks

which extend the tax-auditing block. Therefore, \mathcal{T} needs to pick one and certify it. Afterwards, this certified block cannot be reverted and acts as a “checkpoint”.

We note that \mathcal{L}^t covers the desiderata proposed in Section 1.1. Regarding privacy, although \mathcal{T} can de-anonymize the set of \mathcal{L}^t users at a specific point in time, i.e., when the assets freeze, the users can employ standard Bitcoin addresses and transactions outwith this period. Additionally, as with standard Bitcoin addresses, third parties cannot obtain information regarding the identity of a certified address’s owner (as long as the signature itself does not leak it). In terms of performance, a user can transact with their assets effortlessly, as long as they use certified addresses to receive or unfreeze assets around the taxation period. Importantly, users can certify their addresses ahead of the freezing time, thus the additional load can be spread over a period of a few days or weeks.

3 A Tax-Auditing Extension for Provisions

We now build a tax auditing mechanism for existing ledgers based on Provisions [6]. The goal of this mechanism is to enable all payment recipients to verify whether the assets used by a sender \mathcal{E} in a transaction have been properly declared to the authority \mathcal{T} . This is achieved in two stages, first with an asset declaration stage that involves \mathcal{T} and second with a payer address auditing protocol, which is created in tandem with the transaction that pays a recipient, and after \mathcal{E} commits to owning the assets. If \mathcal{E} fails to provide such proof, the implication is that \mathcal{E} performs tax evasion. To build this proof mechanism we rely on Provisions [6], particularly its *proof of assets*. Our scheme comprises of two simple protocols, which \mathcal{E} runs with the taxation authority and their counter-party respectively. As we show, our protocols retain the privacy guarantees of Provisions.

Provisions is a privacy-preserving auditing mechanism for Bitcoin exchanges. Using Provisions a party can verify that a (cooperating) Bitcoin exchange is solvent, i.e., possesses enough assets to cover the liabilities towards its users. In order to achieve this, Provisions defines three protocols: i) proof of assets, ii) proof of liabilities, and iii) proof of solvency. Our work is only concerned in the assets owned by the exchange, thus we focus on the proof of assets. All proofs are considered under a group G of prime order q with fixed public generators g, h . The proof of assets considers the following:

- $\text{PK} = \{y_1, \dots, y_n\}$: the total (anonymity) set of public keys;
- s_i : a bit such that, if the exchange controls y_i , i.e., if it possesses the private key of y_i , then $s_i = 1$, otherwise $s_i = 0$;
- $\text{bal}(y_i)$: the amount of assets that the address corresponding to y_i controls;
- $\Theta = \sum_{i=1}^n s_i \cdot \text{bal}(y_i)$: the amount of assets that the exchange controls;
- $b_i = g^{\text{bal}(y_i)}$: a bidding (but not hiding) commitment to the balance of y_i .

The exchange publishes the Pedersen commitments [15] for each $s_i \cdot \text{bal}(y_i), s_i$:

$$p_i = b_i^{s_i} \cdot h^{v_i} = g^{\text{bal}(y_i) \cdot s_i} \cdot h^{v_i} \quad (1)$$

$$l_i = y_i^{s_i} h^{t_i} = g^{\hat{x}_i} h^{t_i} \quad (2)$$

where $v_i, t_i \in \mathbb{Z}_q$ are chosen at random, x_i is the private key for y_i , and $\hat{x}_i = x_i \cdot s_i$.

Asset Declaration. \mathcal{E} declares the total amount of assets it controls, i.e., the value Θ , to \mathcal{T} who verifies that \mathcal{E} 's commitments correspond to Θ . We obtain the condition $Z_\Theta = \prod_{i=1}^n p_i = g^\Theta \cdot h^v$, where $v = \sum_{i=1}^n v_i$, is a (publicly computable) Pedersen commitment to \mathcal{E} 's assets. Given that \mathcal{T} knows Θ , \mathcal{E} needs only to prove knowledge of a value v , such that this condition is satisfied. This is done via the Schnorr protocol [16] of Figure 1, which guarantees privacy as described in Lemma 1.

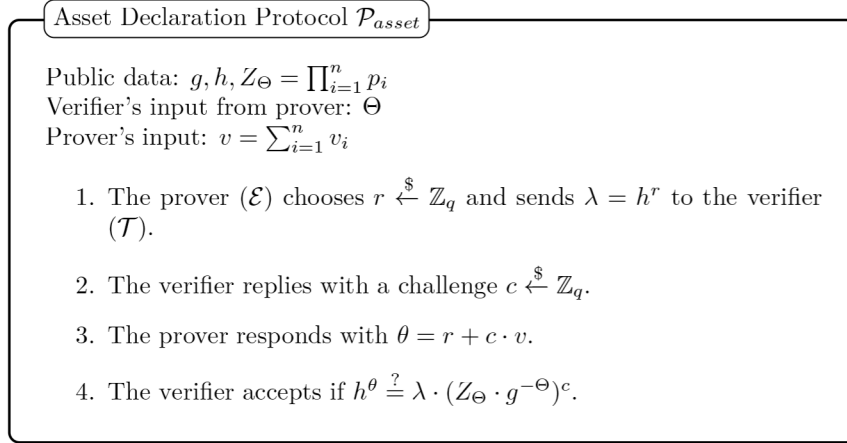


Figure 1: Tax-auditing between \mathcal{E} (prover) and \mathcal{T} (verifier).

Lemma 1. *For public values g, h and Z_Θ , the protocol \mathcal{P}_{asset} is an honest-verifier zero-knowledge argument of knowledge of quantity v satisfying $Z_\Theta = \prod_{i=1}^n p_i = g^\Theta \cdot h^v$ for $i \in [1, n]$.*

Payer Address Auditing. The second part of our taxation proof enables the tax auditing of a specific address used by a payer \mathcal{E} whenever a payment is made to an arbitrary user \mathcal{U} . \mathcal{E} will prove two conditions to \mathcal{U} : i) for some $i \in [1, n]$, the public key y_i (which is published as part of the Provisions scheme) corresponds to the address from which \mathcal{U} receives their assets; ii) the corresponding bit s_i for y_i in the commitment condition (2) is $s_i = 1$. The first condition can be easily proven by providing \mathcal{U} with an index i , such that \mathcal{U} confirms that the address in question is equal to the hash of y_i . To prove the second condition, we observe that, for $s_i = 1$, $p_i = g^{\text{bal}(y_i)} h^{v_i}$ and $l_i = y_i h^{t_i}$. Therefore, \mathcal{E} needs only to prove knowledge of t_i and v_i , such that this statement is satisfied, which can be achieved via the Schnorr protocol of Figure 2, its privacy properties formalized in Lemma 2.

Lemma 2. *For public values g, h and $y_i, l_i, p_i, \text{bal}(y_i)$, the protocol $\mathcal{P}_{address}$ is an honest-verifier zero-knowledge argument of knowledge of quantities t_i, v_i satisfying $l_i = y_i h^{t_i}$ and $p_i = g^{\text{bal}(y_i)} h^{v_i}$ respectively.*

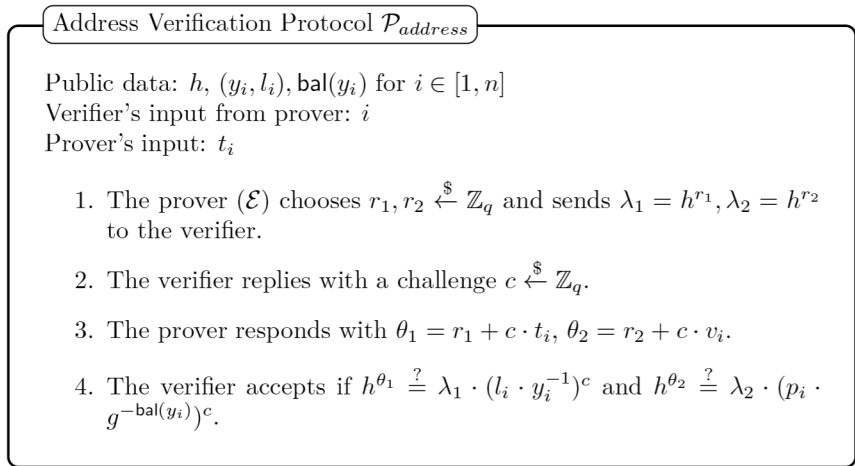


Figure 2: Address verification between \mathcal{E} (prover) and a user \mathcal{U} (verifier).

Finally, both protocols can be turned into non-interactive zero-knowledge (NIZK) proofs of knowledge in the random oracle model by using the Fiat-Shamir transformation [8].

4 Conclusion

Our work offers a programmable money approach for authorities to audit the citizens' tax returns and create a tax-gap counter-incentive: undeclared fund transfers are programmed to be frozen in the ledger. We identify a number of limitations and desiderata and present two basic designs, which can act as a stepping stone for more concrete solutions. Our mechanisms can be employed by different tax authorities and be applied on different ledger designs. Naturally, to efficiently utilize it on a global scale for decentralized systems, like Bitcoin, tax authorities of all countries would need to collaborate, an assumption which seems infeasible in our current fragmented landscape. Nevertheless, a single country's sovereign could deploy it as a feature of, for example, a central bank digital currency. Particular points of interest for future work are the effect of freezing on user experience, as well as the storage overhead. Additionally, our scheme considers pseudonymous systems; future work could explore fully anonymous applications, which utilize zero-knowledge schemes to achieve cryptographic-grade transaction anonymity. Finally, an interesting direction is the design of incentive schemes that motivate the system's adoption and reduce the dependence on enforcement by the authorities.

References

[1] Arun Advani. Who does and doesn't pay taxes? *Fiscal Studies*, 2020.

[2] Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In Hovav Shacham and

- Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 643–673, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-96878-0_22.
- [3] Michel Avital, Jonas Hedman, and Lars Albinsson. Smart money: Blockchain-based customizable payments system. *Dagstuhl Reports*, 7(3):104–106, 2017.
- [4] European Central Bank. Eurosystem report on the public consultation on a digital euro, 2021. URL: https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_c
- [5] Codruta Boar and Róbert Szemere. Payments go (even more) digital*, 2011. URL: https://www.bis.org/statistics/payment_stats/commentary2011.htm.
- [6] Gaby G. Dagher, Benedikt Bünz, Joseph Bonneau, Jeremy Clark, and Dan Boneh. Provisions: Privacy-preserving proofs of solvency for bitcoin exchanges. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015: 22nd Conference on Computer and Communications Security*, pages 720–731, Denver, CO, USA, October 12–16, 2015. ACM Press. doi:10.1145/2810103.2813674.
- [7] Jack Doerner, Abhi Shelat, and David Evans. Zeroledge: Proving solvency with privacy.
- [8] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194, Santa Barbara, CA, USA, August 1987. Springer, Heidelberg, Germany. doi:10.1007/3-540-47721-7_12.
- [9] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-46803-6_10.
- [10] Christina Garman, Matthew Green, and Ian Miers. Accountable privacy for decentralized anonymous payments. In Jens Grossklags and Bart Preneel, editors, *FC 2016: 20th International Conference on Financial Cryptography and Data Security*, volume 9603 of *Lecture Notes in Computer Science*, pages 81–98, Christ Church, Barbados, February 22–26, 2016. Springer, Heidelberg, Germany.
- [11] FISCALIS Tax Gap Project Group. the concept of tax gaps. report ii: Corporate income tax gap estimation methodologies, 2018. URL: <https://op.europa.eu/en/publication-detail/-/publication/a5da4716-e7c1-11e8-b690-01aa>
- [12] R Murphy and A Guter-Sandu. Resources allocated to tackling the tax gap: a comparative eu study. *Working paper for Combating Financial Fraud and*

Empowering Regulators (COFFERS) Horizon 2020 project, November(A), 2018.

- [13] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [14] Neha Narula, Willy Vasquez, and Madars Virza. zkledger: Privacy-preserving auditing for distributed ledgers. In *15th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 18)*, pages 65–80, 2018.
- [15] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO’91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140, Santa Barbara, CA, USA, August 11–15, 1992. Springer, Heidelberg, Germany. doi:10.1007/3-540-46766-1_9.
- [16] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 239–252, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany. doi:10.1007/0-387-34805-0_22.
- [17] Internal Revenue Service. Federal tax compliance research: Tax gap estimates for tax years 2008–2010, 2016.
- [18] Zak Wilcox. Proving your bitcoin reserves, 2014.
- [19] Karl Wüst, Kari Kostianen, Vedran Capkun, and Srdjan Capkun. PRCash: Centrally-issued digital currency with privacy and regulation. Cryptology ePrint Archive, Report 2018/412, 2018. <https://eprint.iacr.org/2018/412>.