Certified Hardware Requirements Undermine Digital Currency

Geoffrey Goodell
University College London
g.goodell@ucl.ac.uk

This Version: 2022–09–29

Abstract

Design approaches based on certified hardware have featured prominently in the recent debate concerning the deployment of large-scale digital currency systems in general and systems to support central bank digital currency (CBDC) in particular. Certified hardware is a cornerstone of the so-called "trusted computing" paradigm, wherein a user carries a device designed to operate in a manner that respects the will of a third party in preference to the will of the user. The justification for such approaches is usually given in terms of the perceived need to facilitate offline payments or to facilitate the recovery of lost assets. In this article, we critically assess the essential problems that digital currency solutions are being proposed to solve, particularly with respect to the future of payments and the future of cash. Next, we examine the characteristics of trusted computing within the context of its application to digital currency systems and its potential impact on the power relationships between the users of devices, the manufacturers of devices, and other powerful actors such as corporations and the state. Finally, we assess the validity of common justifications for certified hardware in the context of alternative designs, limitations, and trade-offs. We conclude that the interests of consumers would be better served by design approaches to digital currency that do not require the use of trusted hardware.

1 Introduction

The proliferation of cashless payments has given rise to a growing preponderance of merchants who refuse to accept cash, as well as the suggestion that private-sector payment systems could undermine the authority of central banks to implement monetary policy. This development has, in turn, motivated the attention of central banks and other financial institutions around the world on digital currency, particularly the idea that central banks could issue *central bank digital currency* (CBDC) for retail use. In principle, retail CBDC would be a direct obligation of the central bank, like cash or central bank reserves, but existing in digital form and held by ordinary consumers and non-financial businesses.

For ordinary consumers and non-financial businesses to hold CBDC directly, there must be some way for CBDC to be held outside accounts. Accounts generally involve a relationship between a beneficial owner (the *account holder* and a fiduciary, typically a regulated financial institution. A salient property of accounts is that access to assets held within accounts is determined by the identity of the account holder [1]. In contrast, *tokens* are assets that can be held outside of accounts. Tokens can be *bearer instruments*, meaning that they can be used in a way that does not index the identity of their owners.

Compared to assets held within accounts, assets that are held directly by their owners have a different set of security requirements for their owners. For example, owners of assets held within accounts can rely upon various identification mechanisms, such as PIN numbers and two-factor authentication, to protect their assets if their cards or mobile devices are stolen, and asset custodians can also provide credit to account holders, such that if assets are spent unlawfully, the account holders can ask the asset custodians cover the value of their losses. In contrast, the bearer is responsible for bearer instruments. If cash is lost or stolen, owners generally have little recourse beyond law enforcement mechanisms to recover their money, and such mechanisms often fail to recover the value. The potential for loss or theft exists not only physically but also electronically. Devices can malfunction, and devices

can be compromised. The prevailing model for handling devices that no longer work or have been compromised is to wipe their contents completely or to replace them wholesale. When such devices are used to access accounts, such replacement is not much of a problem, although bearer instruments inexorably expose their owners to new risks that require different solutions.

Many proponents of retail CBDC have argued that it must be "cash-like" to deliver value to end-users. This argument is at least partially right. For example, it might be expected that a cash-like system would afford users privacy by design, resistance to discrimination and profiling, and the ability to possess and control the assets that they use to engage with the economy. These are certainly desirable properties. However, the argument that a digital payment system must be cash-like has also been interpreted to mean that a digital payment system must enable fully "offline" transactions, without network connections of any kind. We suggest that this argument constitutes a leap of logic that must be examined carefully, starting with its implicit assumptions that any new payment system must be appropriate for all scenarios, that existing solutions such as cash are insufficient to handle payment scenarios in which neither party has access to a network connection, and that there is a specific need for future CBDC infrastructure to enable a kind of digital transaction that is not commonly available to users today.

Wholly "offline" transactions impose additional security requirements to prevent double-spending attacks, in addition to the new security requirements applicable to digital bearer instruments in general. Some prominent experts have proposed "trusted computing" [2, 3] as a solution to the security problem, reasoning that if the security of bearer instruments is a cause for concern, then perhaps there is benefit to be realised in making an effort to ensure that the devices that can hold them must meet a certain standard. After all, many users are accustomed to carrying certified hardware, and the prevalence of certified hardware among common user devices might be seen as a justification for their use in the context of digital currency. For example, smart cards commonly rely upon tamper-resistant chips containing secrets issued by their manufacturers [4]. Mobile phones often contain secure enclaves to restrict what their users can do with their devices, for example to prevent users from copying multimedia files.

Of course, the fact that certain methods are commonly used in some contexts does not justify their use in others, and certainly does not justify mandating their use in the context of future infrastructure that intrinsically must be available for general use by the public. Mandating the use of certified hardware carries many costs, and the benefits of certified hardware are not as clear as its proponents claim. In this article, we debunk the argument that trusted computing is appropriate as an enabler of digital currency solutions. In the next section, we consider the arguments in favour of offline payments, toward an understanding of precisely what is meant by the use of the term "offline", as well as the specific kinds of offline payments for which trusted computing has been proposed as a necessary solution. In Section 4, we critically examine the requirement for trusted computing in the context of central bank digital currency, identifying several key problems that undermine its perceived value. In Section 3, we assess the various motivations and expectations for CBDC in general, and whether the assumptions about offline payments make sense. In Section 5, we consider an alternative design that does not rely upon trusted computing and argue that this and similar designs are better-suited for retail CBDC infrastructure. In Section 6, we summarise our various arguments and conclude.

2 Offline payments

Recently, several prominent organisations, including the International Monetary Fund [5] as well as central banks such as the Bank of England [6] and the Bank of Canada [7], have released reports suggesting that "offline" payments are desirable or even essential to the functionality of a future retail CBDC. Such bold statements are striking, considering the prevailing lack of consensus about what the requirements for a CBDC should be in general. But what exactly constitutes an "offline" payment? We consider various kinds of payment mechanisms and assess them in terms of whether they are online or offline:

- *E-commerce transactions.* Transactions that involve consumers using personal computers or mobile devices to interact with websites require the internet by definition. They represent a significant and growing share of the total value consumers spend in the economy, and naturally, they require consumers to be online.

- *Bank transfers.* Bank transfers are mechanisms for transferring money from a sender's account with a sending bank to a recipient's account with a receiving bank. As account-holders, both the sender and the recipient must be identified. Bank transfers can be arranged online, or an account-holder can arrange a bank transfer via a conversation with a bank teller or a computer controlled by the bank. The sending bank debits the sender's account, and the receiving bank credits the recipient's account. Settlement between the banks is managed via an interbank clearing and settlement network. In most cases, interbank settlement networks cannot guarantee "immediate" settlement within a few seconds, and bank transfers are usually not considered appropriate at the point of sale or for transactions that require acceptance in real-time.

- *Electronic card payments at the point of sale (POS).* Electronic card payments provide a mechanism by which merchants can accept in-person payments from consumers in real-time. The consumer's bank (the *issuer*) and the merchant's bank (the *acquirer*) both participate in a *payment network*. The issuer provides a card to the consumer, and the acquirer offers to connect a merchant's POS device to the payment network. The card serves as a form of identification: Its purpose is to specify a particular bank account and to identify its user as the beneficiary of that account. To protect against unauthorised copying of the identification credential, cards are usually designed with features that provide tamper-resistance. The client is not assumed to have an Internet connection, and the POS device provides a way for the card to authenticate itself to the payment network. If the authentication is successful, then the payment network effectively underwrites the transaction, implicitly providing credit to the issuer and prompting the acquirer to accept the transaction immediately. The payment network informs the merchant that the transaction is approved, and the merchant can then accept the payment in real-time. Merchants pay interchange fees to support this mechanism.

- *Closed-loop payment systems at the point of sale.* With closed-loop payment networks, such as Alipay and WeChat Pay [8, 9], both consumers and merchants are primary customers of a *platform operator*, which maintains accounts with balances for all of its customers. When a consumer pays a merchant in-person, the platform operator debits the consumer's account and credits the merchant's account. For a transaction to succeed, either the consumer or the merchant must have an active connection to the platform operator. The consumer must authenticate successfully to the platform operator, although this can be done via the merchant's device, as it is with card payment networks. Crucially, banks are not generally involved in transactions; the platform operator manages the process.

- *Cash.* Cash comprises a set of physical tokens, each of which is a direct obligation of a central bank or monetary authority. Cash is held directly by individuals and non-financial businesses as a store of value, and it is also held by banks to service *withdrawals*, wherein banks debit accounts of customers in exchange for providing them with cash. Cash is *fungible*, which is to say that individual cash tokens are mutually substitutable and undifferentiated in practice. Since the consummation of a cash transaction is the exchange of a physical token, cash transactions between parties are offline by definition. Regulators generally do not have a direct view into cash transactions, and assessment of regulatory and tax compliance depends upon accounting and reporting.

E-commerce transactions are certainly online, and cash transactions are certainly offline. However, whether point of sale transactions and bank transfers are online or offline is a matter of perspective. For example:

- Is it enough for one party but not both to be connected to the Internet?
- Is it necessary to be able to reach a particular service run by a specific operator, or is it sufficient to be able to reach any one of a set of operators?
- Can a transaction to be provisionally allocated using the network but consummated when both parties are disconnected from the network?

Notwithstanding some measure of understandable confusion and inconsistency in the use of the term "offline" accompanying calls for considering offline payments to be an essential feature of CBDC, it can be said that some measure of online technology supports virtually all retail payments today, excepting only payments in cash, which has seen a dramatic decline in use throughout the world over the past decade. With this in mind, it is reasonable to ask: What problem do proponents of retail CBDC intend to solve, for which offline payments might constitute a solution?

3 The motivation for offline retail CBDC

One might argue that conducting transactions using digital currency held by custodians is, from the perspective of users and counterparties, not much different than conducting transactions practice of using ordinary bank accounts, with a possible exception of a different regulatory environment for the digital currency custodian. The custodian is still a gatekeeper that could fail to meet its obligations; the user is still bound to the custodian by identification requirements and still does not hold the assets directly. One might also argue that digital currency held by custodians is even less useful than bank deposits, if it is assumed that digital currency tokens cannot be rehypothecated and therefore cannot contribute to money creation.

We might reasonably conclude that *bearer wallets* (or "non-custodial wallets"), which allow users to hold digital assets directly, are a foundational benefit of CBDC. However, even if we accept that bearer wallets are essential to CBDC, the idea that offline payments are necessary for CBDC is certainly false. Users have an interest to possess and control digital assets directly, regardless of the context in which they spend them. Holding digital assets directly takes custodians out of outbound transactions, allowing users to avoid the risk that the custodians might block certain (or all) transactions from taking place, or build profiles of their behaviour by monitoring their transactions. Bearer wallets can be used in a variety of contexts, including but not limited to in-person POS transactions. Importantly, the benefits of bearer wallets can also be realised in a fully online context such as e-commerce.

Proponents of CBDC are right to acknowledge the secular decline in the use of cash in retail consumer transactions as a primary motivator for CBDC. As fewer consumers use cash for transactions, the variable revenues associated with operating a cash infrastructure are falling below the fixed costs. This pattern is evident in much of the developed world. For example, in the UK, banks are closing ATMs and branches [10], making cash more difficult for consumers to withdraw [11]. Some merchants have decided that the marginal cost of the various expenses related to accepting cash payments, such as tills, employee time, cash accounts with banks, and the means of transporting cash, outweighs the benefit of providing goods and services to cash-only customers. As a result, the value of cash to consumers is reduced, since customers no longer have assurance that their cash will be accepted at the point of sale, creating a feedback loop.

And so, when proponents of CBDC say that CBDC should allow "cash-like" transactions for retail users of money, they are referencing the idea that CBDC can provide individuals with a payment option that provides the key affordances of cash, including direct ownership, privacy, and equal access without discrimination. Users of CBDC would hold direct obligations of the central bank, and, given the right regulations, they might have reason to believe that their CBDC would be accepted by retailers that require electronic transactions.

CBDC can certainly address problems associated with the systematic decline in the set of available options for consumers to access cash and the increasing tendency for merchants to refuse payments in cash at the point of sale. Users can load bearer wallets with the help of a personal device with network access, as well as via facilities provided by banks or ATM networks. Consumers can use bearer wallets in-person, to interface with point of sale devices in the same manner that they use debit cards or other custodial payment mechanisms today. Such payments can be considered "offline" both in the sense that they are in-person and in the sense that that the bearer wallets themselves do not require Internet connections.

There is a clear case for using CBDC with bearer wallets as a substitute for electronic card payments or network platform transactions at the point of sale. What remains is the question of whether there exists a need to support payments between two parties such that neither of which is connected to a network. There is no evidence that consumers who cannot withdraw cash or merchants who refuse to accept cash are using alternatives to cash that do not make use of a network connection instead, and there is no evidence of a significant demand for one, either. As a means of facilitating a transaction between two parties in situations in which neither party has a networked device, cash is remarkably effective. Although it is possible to imagine narrowly construed use cases for new capabilities, such use cases do not justify system requirements.

It is worth asking whether governments or central banks are harbouring an intention to do away with cash entirely, or whether the presumption of such an intention is shaping the view of what CBDC should be. For example, some authorities have argued that a larger share of domestic transactions should be brought inside the scope of the formal economy []. Legitimate or not, that goal is not inconsistent with prevailing mechanisms to roll

out digital transactions in general and does not require CBDC. Moreover, few monetary authorities or government regulators have called for the abolition of cash, whereas many, including those of Switzerland, the United Kingdom, and the United States, have pledged to preserve access to cash for the foreseeable future. Notably, Sweden has supported cash infrastructure in a near-cashless economy, at great cost. So the idea that CBDC is mostly a justification for phasing out cash, as has been speculated by some, is not corroborated by the facts.

4 The perils of trusted computing

If we define a *fair exchange* as a transaction mechanism that ensures that transacting parties either each receive or each do not receive what it expects (and do not receive "any additional information about the other's item") [13], then it is not possible to have a fair exchange between two parties without the involvement of a third party; this is a theoretical truth [14]. With the exchange of physical objects, possession of the object itself is sufficient to demonstrate the veracity of the exchange that has taken place. But with digital objects the situation is more complicated. It is possible to have fair exchange without connection to the Internet, provided that there are third parties accessible locally, perhaps via a local network, as might be possible with ad hoc networks.

Proponents of approaches requiring certified hardware have proposed embedding into end-user devices a tamper-resistant "secure element" that can be used to systematically restrict what the device can do, disclose the immutable identity of the device to a third party, or furnish information to third parties about what the device had observed or done. The secure element would be impenetrable to the user, and the user would not be able to modify its behaviour without rendering the device inoperable. In effect, such an element effectively acts as the third party by proxy, by embedding the will of the third party into the operation of a device carried by a user.

In the case of digital currency, the secure element would presumably be used as an on-board authority that would prevent double-spending by maintaining a record of transactions made by the device and disallowing users from transacting tokens that have already been signed over to others. This design model in particular, and trusted hardware in general, introduces a variety of problems. We consider some of those problems here:

1. *Trusted computing introduces significant security risks particular to digital currency.* The risks of relying upon the security of a hardware manufacturer are too great for a token-based CBDC context. In particular, a system design that relies upon trusted hardware components is fragile by design. If the trusted hardware fails to function as advertised or is compromised by an attacker, then users could potentially transact the same tokens an arbitrary number of times without bound. The system-level costs of such an attack would be huge, and the incentives would be tremendously enticing for attackers. We note that modern devices, such as debit cards, that rely upon tamper-proof hardware for the exchange of value generally interact with accounts. If their security is compromised, then the value available to attackers is limited by rules and operational procedures implemented by custodians. Without custodial accounts, such backup mechanisms are not possible.

2. *Secure elements are unaccountable to their users.* Because trusted hardware is impenetrable to users, users cannot know with certainty what their devices are doing. The inability to analyse the behaviour of a device undermines the trust that users can and should place in the device. In essence, the device serves a second master that is not the user. Users cannot trust what they cannot verify, so they have every reason to be suspicious of unchecked promises on the part of hardware manufacturers or other privileged guarantors [15]. For example, secure elements can furnish information that can be used to identify the user or the user's previous transactions. By contrast, with open hardware, skilled members of the community can conduct audits and provide reports to the public. Of course, users with the means to do so can personally conduct the audits themselves or hire contractors to conduct the audits for them. However, it is not necessary for every user to actually run an audit; evidence suggests that similar benefits are achieved so long as all users *can* perform audits if they choose to do so [16].

3. *Certified hardware discourages innovation.* Users who cannot create their own devices cannot innovate. Instituting a certification requirement for hardware serves to ensure that only those with the time and resources to pursue certification will be able to produce working devices. Regulations that increase the cost of entry are known to decrease entrepreneurship as well as innovation by existing firms [17]. Ordinary users

and hobbyists will not be able to develop their own devices, creating a barrier to public understanding the of the core technology, encumbering public debate and consideration of alternative designs [18, 19].

4. *Certified hardware distorts the market by privileging its manufacturers.* The requirement to use trusted hardware effectively locks the hardware manufacturer in to a privileged position in the market, wherein prospective manufacturers require certification by some authority (either directly by the government or by a delegated authority such as a trade organisation), and concomitant lobbying and regulation can restrict free entry, increasing market concentration [20]. Furthermore, even if all manufacturers could have certification at zero marginal cost, the fact that fabrication facilities carry high fixed costs means that the market is likely to be highly concentrated, with manufacturers capturing the surplus [21].

5. *A certified device is tantamount to a relationship with a custodian.* Hardware manufacturers become gatekeepers at best, and globally trusted third parties in the general case. But unlike trusted third parties on the network, users bind themselves to long-term relationships with personal devices that contain trusted hardware; the relationships are not much different between account-holders and custodians, wherein account-holders are subject to satisfying the requirements of custodians. The cost of switching bank accounts is generally only a matter of time and effort on the part of account-holders, whereas the cost of switching devices generally includes per-device manufacturing and retailing cost plus taxes and costs associated with physical acquisition and disposal of the devices. In addition, while it is possible to implement both technical protocols and regulations that rely upon locally-trusted third parties, economies of scale mean that hardware manufacturers are likely to have global reach. It is reasonable to conclude that relationships between users and their devices featuring specific trusted hardware will be sticky indeed.

For these reasons, we are unconvinced that trusted hardware solutions are appropriate as requirements for the design of tools that allow users to possess and control their own assets. At the same time, we would not suggest that trusted hardware should be forbidden. In some fraction of cases in which physical cash is not an option and network connections are not available, trusted hardware might be an alternative kind of third party for transacting parties to consider, although this possibility is not sufficient justification to institute a requirement. In all cases, both sender and recipient must willingly choose to trust the third party to consummate the transaction, and given the final three points in the last above, it is difficult to imagine that a user preference for a particular device manufacturer would be real. Consider, for example, that existing users of secure hardware in the form of electronic payment cards generally have little say in which payment network their issuing banks choose.

It is plain to see why a hardware manufacturer would lobby in favour of regulations and standards that declare "offline transactions" to be an essential requirement for CBDC. But this argument is unsound: Cash is an excellent offline payment method, and to our knowledge, few if any central banks have suggested that the purpose of CBDC is to replace cash outright. The argument implicitly equivocates between a concept of "offline" that denotes transactions wherein a user is not required to carry a device with network connectivity and a concept of "offline" that involves two counterparties physically meeting in a location without third parties or network infrastructure of any kind nearby. The former scenario is the only one in which users commonly use certified hardware today. The idea that cash could suddenly become unavailable for use in the latter scenario is an overly simplistic justification for hardware manufacturers to tout their products.

5 Alternatives to trusted computing

For a set of specific technical requirements for CBDC in general to be credible, it must address the requirements of its various stakeholders. One of the unfortunate characteristics of the current debate is the diversity of conflicting assumptions about the nature of payments and the relative importance of different kinds of payment scenarios. In reality, there are many different ways that people make payments today: with and without custodians, locally and at a distance, with and without electronic devices, with and without money issued by a central bank, in exchange for goods or services, as a means to make purchases, as a means to give gifts, as a means to pay taxes. These different scenarios imply a plethora of different costs and risk profiles: Buying a cup of coffee is different from buying a house. We should accept that there will not be a single payment system or key mechanism to rule them all, although it might be possible to develop a common framework that can accommodate a variety of different payment systems.

The case for trusted hardware, as for other core mechanisms of proposed payment systems, involves the requirement for assessing the *integrity* of money. In principle, trusted hardware can act as a trusted third party to enforce system-level rules on user devices directly, by implementing rules that prevent assets from being spent more than once.

These rules can be enforced at the system level as well, for example by invoking a trusted third party to validate all transactions, as proposed by Chaum and others [23, 24]. However, there is reason to believe the costs, risks, and public accountability concerns associated with having a central operator dictate that history unchecked are too great for the CBDC use case [25], and the most important costs and risks would be borne not only by the users but also by the central operator. A modern alternative involves using a distributed ledger system to put the money itself on a blockchain, the way that UTXO designs (e.g. Bitcoin [22], Monero [26]) and state-transition designs (e.g. Ethereum [27]) do. Such approaches are unwieldy because they require the ledger to track all of the tokens individually, and transaction finality requires global consensus. In practice, most cryptocurrency users rely upon custodial accounts, negating a foundational benefit of digital currency.

However, distributed ledgers and blockchains can be used to support and enable payment systems in a variety of ways [28]. The digital currency design that we have proposed addresses the scalability problems associated with modern cryptocurrency systems while strongly preventing double spending without trusted hardware [29]. In particular, it avoids requiring the issuer to participate in every transaction of its tokens, it avoids requiring system operators to manage a list of tokens that have been created or spent, and it avoids requiring users to have access to a global network at the time of the transaction. Any third party chosen by the spender and accessible to the recipient at the time of the transaction can suffice.

Proposed solutions, including ours, can work technically without distributed ledgers, but in the end, digital currency systems require a means to support an immutable, shared history mutually accepted among independent actors, which is the salient characteristic of a distributed ledger [30]. The immutability of a distributed ledger is emergent property of the independence of the participants. So, although using a distributed ledger to record and process tokens directly on the ledger might bring unacceptable costs in terms of scale, efficiency, and operational burden, we should not rule out using distributed ledgers for what they do best: externalising commitments so that parties responsible for maintaining the history do not equivocate [29].

Trusted computing is not necessary to protect users from loss or theft, either. To protect their assets, users may, at their option, choose devices that offer security features such as encryption, passcodes, or local biometrics to protect against theft, and they may make copies of their tokens to mitigate the risk of accidental loss of their devices.

A better way to address the limited set of cases in which "fully offline" payments (that is, those where neither the payer nor the payee have access to any suitable third parties or network infrastructure) are necessary would be to support the use of cash where cash continues to work well, such as in remote locations and environments with low connectivity, while avoiding all of the aforementioned problems associated with trusted hardware. To address the reduction of cash in the "built" environments by providing CBDC, whilst encouraging the provision of facilities for converting between cash and CBDC at the boundary of "unbuilt" environments, would be an entirely legitimate approach. In particular, if one potential benefit of CBDC in an age of deteriorating cash infrastructure is the ability to electronically transport cash to remote population centres during crises [31], then there is no reason to require that CBDC distributed in this manner must remain digital. For example, it can be exchanged for physical tokens specific to that population centre, much as cash and bank money are exchanged for chips at a casino.

None of this is to say that we should forbid limited groups of users from trusting certified devices if that is what they freely choose. However, at a system level, placing trust in certified devices introduces costs and risks that are entirely avoidable.

6 Conclusion

In the end, we must ask ourselves: What problem are we trying to solve? Manifestly, the problem is the decline of cash, and the concomitant challenges of maintaining cash infrastructure when the variable revenues of cash infrastructure fall below the fixed costs. But the decline of cash really only exists in environments with high network connectivity, wherein "fully offline" payments are not necessary. For environments with low network connectivity, the decline of cash has not occurred, and cash continues to be the best solution for such environments.

Many central banks, including the Bank of England, have affirmed a commitment to cash [32], while underscoring the risks associated with the decline of cash use in favour of online payment schemes such as card payment networks [33]. Really, the use cases that CBDC payments should be targeting are cases for which at least one counterparty is already online. To argue for digital currencies with the premise that support for "fully offline" digital payments are its principal source of added value mainly supports the view that digital currencies are solutions looking for a problem, as some observers have prominently suggested [34].

The purpose of bearer wallets (and indeed, retail CBDC in general) is not primarily about enabling transactions in which both parties are offline. The purpose of bearer wallets is to allow users to hold assets directly, without a custodian. Put differently, this is a way for users to have *private property* that they can exchange. By facilitating such payments, CBDC can protect the interests of consumers in all cases, including fully networked ones. If the recipient of digital assets has network connectivity (for example, an electronic point of sale device), then transactions involving bearer wallets can be consummated without special hardware. Even if we accept that fair exchange is impossible without the participation of a trusted third party, and that the trusted third party can participate via the network or via trusted hardware, it is important to underscore that a non-networked bearer wallet can conduct a fair exchange with a networked counterparty, without the aid of special security mechanisms built into the wallet.

It is a mistake to assume the primary use for bearer wallets is to allow a government to phase out cash, or that custodial wallets are sufficient for all scenarios except the narrow case in which neither the sender nor the receiver have a communication line to a third party. Custodial wallets require network access as a rule, because consumers must be able to use their devices to access their custodial accounts. However, bearer wallets do not require either network access or special hardware in the general case. To imply otherwise is to impede the development of a technology that may become crucial for the digital economy.

Acknowledgements

We thank Professor Tomaso Aste for his continued support for our work on digital currencies and digital payment systems. We also acknowledge the support of the Centre for Blockchain Technologies at University College London and the Systemic Risk Centre at the London School of Economics, and we acknowledge EPSRC and the PETRAS Research Centre EP/S035362/1 for the FIRE Project.

References

[1] R. Auer and R. Böhme. "The technology of retail central bank digital currency." *BIS Quarterly Review*, March 2020, pp. 85–100. [online] `https://www.bis.org/publ/qtrpdf/r_qt2003j.pdf` [retrieved 2020-05-19]

[2] Trusted Computing Group. "TCG Specification Architecture Overview." Revision 1.2, 2004-04-28. [online] `https://web.archive.org/web/20081117214013/https://www.trustedcomputinggroup.org/downloads/TCG_1_0_Architecture_Overview.pdf` [retrieved 2008-11-17]

[3] E Roberts, "What is Trusted Computing?" Stanford University. [online] `https://cs.stanford.edu/people/eroberts/cs201/projects/trusted-computing/what.html` [retrieved 2021-07-29]

[4] European Telecommunications Standards Institute (ETSI). "Secure Elements (Smart Cards)." [online] `https://web.archive.org/web/20220201021409/https://www.etsi.org/technologies/smart-cards` [retrieved 2022-09-24]

[5] J Kiff. "Taking Digital Currencies Offline." International Monetary Fund, August 2022. https://www.imf.org/en/Publications/fandd/issues/2022/09/kiff-taking-digital-currencies-offline [retrieved 2022-08-03]

[6] Bank of England. Minutes of the CBDC Technology Forum, June 2022. [online] https://www.bankofengland.co.uk/minutes/2022/june/minutes-of-cbdc-forum-june-2022 [retrieved 2022-07-06]

[7] J Miedema, C Minwalla, M Warren, and D Shah. "Designing a CBDC for universal access." Staff Analytical Note 2020-10, Bank of Canada, June 2020. [online] https://www.bankofcanada.ca/2020/06/staff-analytical-note-2020-10/ [retrieved 2020-10-14]

[8] A Klein. "China's Digital Payments Revolution." Brookings Institution, April 2020. [online] https://www.brookings.edu/wp-content/uploads/2020/04/FP_20200427_china_digital_payments_klein.pdf [retrieved 2022-09-28]

[9] J Ehrentraud, J Prenio, C Boar, M Janfils, and A Lawson. "Fintech and payments: regulating digital payment services and e-money." Financial Stability Institute, Bank for International Settlements, FSI Insights on policy implementation No 33, July 2021. [online] https://www.bis.org/fsi/publ/insights33.pdf [retrieved 2022-09-28]

[10] Access to Cash Review (UK), Final Report, March 2019. https://www.accesstocash.org.uk/media/1087/final-report-final-web.pdf

[11] D Tisher, J Evans, K Cross, R Scott, and I Oxley. "Where to Withdraw? Mapping access to cash across the UK." University of Bristol, November 2020. http://www.bristol.ac.uk/media-library/sites/geography/pfrc/Where%20to%20withdraw%20-%20mapping%20access%20to%20cash%20across%20the%20UK.pdf

[12] L Klapper, M Miller, and J Hess. "Leveraging digital financial solutions to promote formal business participation." International Bank for Reconstruction and Development, The World Bank, 2019. [online] https://documents1.worldbank.org/curated/en/486541556177550649/pdf/Leveraging-Digital-Financial-Solutions-to-Promote-Formal-Business-Participation.pdf [retrieved 2022-07-11]

[13] N Asokan. "Fairness in Electronic Commerce." Ph.D. thesis, University of Waterloo, 1998. [online] https://www.researchgate.net/profile/N-Asokan/publication/2281851_Fairness_in_Electronic_Commerce/links/00463529834ec77716000000/Fairness-in-Electronic-Commerce.pdf [retrieved 2022-09-24]

[14] H Pagnia and F Gärtner, "On the impossibility of fair exchange without a trusted third party." Darmstadt University of Technology Department of Computer Science Technical Report TUD-BS-1999-02, 1999. [online] http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=9D84F245D05FE33F5B67E3BB4D1C9489?doi=10.1.1.44.7863&rep=rep1&type=pdf [retrieved 2021-02-13]

[15] R Stallman. "Can You Trust Your Computer?" Free Software Foundation, 2002. [online] [retrieved 2022-01-13]

[16] T Heydt-Benjamin et al, H Chae, B Defend, and K Fu. "Privacy for Public Transportation." In *Lecture Notes in Computer Science* **4258**, ISBN: 3540687904, 2006, pp. 1–19. [online] https://doi.org/10.1007/11957454_1

[17] N Kozeniauskas. "What's Driving the Decline in Entrepreneurship?" 2022-08-09. [online] https://nicjkoz.com/files/220809_WDDE.pdf [retrieved 2022-09-20]

[18] E Felten. "Freedom to Tinker: The Struggle to Access Devices You Own." Princeton University, 2002. [online] https://www.naefrontiers.org/21886/Freedom-to-Tinker-Paper [retrieved 2022-09-20]

[19] P Samuelson. "Freedom to Tinker." *Theoretical Inquiries in Law* **563**, 2016. [online] https://lawcat.berkeley.edu/record/1127533/files/fulltext.pdf [retrieved 2022-09-20]

[20] G Gutiérrez and T Philippon. "The Failure of Free Entry." National Bureau of Economic Research Working Paper 26001, June 2019. [online] https://germangutierrezg.com/GutierrezPhilippon_FreeEntry_2019.pdf [retrieved 2022-09-20]

[21] W Gao and M Kehrig. "Returns to Scale, Productivity and Competition: Empirical Evidence from U.S. Manufacturing and Construction Establishments." Working paper, 2017-05-01. [online] https://events.barcelonagse.eu/live/files/1775-matthias-kehrig-54183pdf [retrieved 2022-09-20]

[22] S Nakamoto. "Bitcoin: A Peer-to-Peer Electronic Cash System." 2009-01-03. [online] https://bitcoin.org/bitcoin.pdf [retrieved 2018-10-04]

[23] D. Chaum. "Blind Signatures for Untraceable Payments." *Advances in Cryptology: Proceedings of Crypto* **82**(3), pp. 199–203, 1983. [online] http://www.hit.bme.hu/~buttyan/courses/BMEVIHIM219/2009/Chaum.BlindSigForPayment.1982.PDF [retrieved 2018-09-28]

[24] D Chaum, C Grothoff, and T Moser. "How to issue a central bank digital currency." Swiss National Bank Working Paper 3/2021, January 2021. [online] https://www.snb.ch/n/mmr/reference/working_paper_2021_03/source/working_paper_2021_03.n.pdf [retrieved 2021-02-25]

[25] G Goodell, H Nakib, and P Tasca. "A Digital Currency Architecture for Privacy and Owner-Custodianship." *Future Internet* 2021, **13**(5), May 2021. [online] https://doi.org/10.3390/fi13050130

[26] N Van Saberhagen. "CryptoNote v 2.0." 2013-10-17. [online]
https://web.archive.org/web/20201028121818/https://cryptonote.org/whitepaper.pdf [retreived 2022-07-14]

[27] V Buterin. "Ethereum Whitepaper." 2013. [online] https://ethereum.org/en/whitepaper/ [retrieved 20201-01-08]

[28] G Goodell. "Tokens and Distributed Ledgers in Digital Payment Systems." Discussion paper, July 2022. [online]
https://arxiv.org/abs/2207.07530

[29] G Goodell, D Toliver, and H Nakib. "A Scalable Architecture for Electronic Payments." Presented at *WTSC 2022*, Grenada, May 2022. To appear, *Lecture Notes in Computer Science* **13412**. [online] http://dx.doi.org/10.2139/ssrn.3951988

[30] International Organization for Standardization (ISO). "Blockchain and distributed ledger technologies – Vocabulary." *ISO/22739:2020*, First edition, 2020. [online] https://www.iso.org/obp/ui/#iso:std:iso:22739:ed-1:v1:en [retrieved 2020-12-01]

[31] G Goodell, H Nakib, and P Tasca. "Digital Currency and Economic Crises: Helping States Respond." LSE Systemic Risk Centre Special Papers SP 20, September 2020; Presented at 6th Annual Peer-to-Peer Financial Systems Workshop (P2PFISY 2020). [online] https://systemicrisk.ac.uk/sites/default/files/2020-09/SP-20_0.pdf

[32] Bank of England. "Responses to the Bank of England's Discussion Paper on new forms of digital money." Discussion paper, 2022-03-24. [online] https://www.bankofengland.co.uk/paper/2022/responses-to-the-bank-of-englands-discussion-paper-on-new-forms-of-digital-money [reteieved 2022-09-28]

[33] Bank of England. "Update on the future of Wholesale Cash Distribution in the UK." Public update on the work undertaken by the Wholesale Distribution Steering Group, 2021-12-15. [online]
https://www.bankofengland.co.uk/paper/2021/update-on-the-future-of-wholesale-cash-distribution-in-the-uk [reteieved 2022-09-28]

[34] House of Lords Economic Affirs Committee. "Central bank digital currencies: a solution in search of a problem?" 3rd Report of Session 2021-22, HL Paper 131, 2021-12-14. [online]
https://publications.parliament.uk/pa/ld5802/ldselect/ldeconaf/131/131.pdf [retrieved 2022-02-17]