

Guo, Sky; Kreitem, Joseph; Moser, Thomas

Article

DLT options for CBDC

Journal of Central Banking Theory and Practice

Provided in Cooperation with:

Central Bank of Montenegro, Podgorica

Suggested Citation: Guo, Sky; Kreitem, Joseph; Moser, Thomas (2024) : DLT options for CBDC, Journal of Central Banking Theory and Practice, ISSN 2336-9205, Sciendo, Warsaw, Vol. 13, Iss. 1, pp. 57-88, <https://doi.org/10.2478/jcbtp-2024-0004>

This Version is available at:

<https://hdl.handle.net/10419/299093>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



UDC: 336.711:004.738.5

DOI: 10.2478/jcbtp-2024-0004

Journal of Central Banking Theory and Practice, 2024, 1, pp. 57-88

Received: 22 February 2023; accepted: 11 July 2023.

Sky Guo *, **Joseph Kreitem** **, **Thomas Moser** ***

* Cypherium,
New York, United States

Email:
sky@cypherium.io

** Cypherium,
New York, United States

Email:
josephkreitem@gmail.com

*** Swiss National Bank,
Zürich, Switzerland

Email:
thomas.moser@snb.ch

DLT Options for CBDC¹

Abstract: This paper provides an overview of the distributed ledger technology (DLT) options available to central banks for issuing central bank digital currency (CBDC). We discuss the main requirements that a DLT solution must fulfill and analyze the various structures for implementation offered by DLT — public, permissioned and private — and the implications that each has for the central bank and the existing financial system. While a CBDC built on an open, permissionless system would provide the full functionality offered by DLT, it is also far more disruptive to the existing financial system and consequently requires more new infrastructure on the part of the central bank.

Keywords: Central Bank Digital Currency (CBDC), Blockchain, Distributed Ledger Technology (DLT)

JEL classification: E42, E51, E52, E58, G2.

1. Introduction

As the prospect of a central bank digital currency (CBDC) moves from research to proof of concepts, decisions on the design become increasingly important. Design choices are inextricably linked to fundamental questions about the nature of money, its utilization, and the government's role in the monetary system. The implications of any CBDC design choice are therefore far reaching, and any discussions of CBDCs require multiple considerations.

¹ The authors would like to thank Reto Bhunjun, Darrell Duffie, Matthias Jüttner, Benjamin Müller, and an anonymous referee for comments and suggestions. The views expressed in this paper are those of the author(s) and do not necessarily reflect those of the Swiss National Bank.

While there are several technical options for a CBDC, distributed ledger technology (DLT) – which includes Blockchains – would offer the greatest potential for a fundamental transformation of the financial system.² This paper provides an in-depth look at the various options available for central banks and the implications for the kinds of CBDC projects that different DLTs enable. We establish some basic criteria for central banks to consider when choosing which DLT to build on. While this paper will not make recommendations for which DLT, if any, to use, it will provide a framework that serves as both context and guidance for central banks in making their decision.

The literature on CBDC and DLT is growing rapidly. BIS (2018) provides a general introduction to CBDC. For a high-level overview of CBDC design choices see Allen et al. (2020), BIS (2020) and BIS (2021). For a comparison of approaches selected in actual central bank CBDC projects see Auer, Cornelli & Frost (2020). For surveys of DLT technologies see Rauchs et al. (2018) and Chowdhury et al. (2019). While there are some papers that provide a decision framework for selecting DLTs for industrial applications (see, e.g. Kubler, Renard, Ghatpande, Georges & Le Traon, 2023), this paper is to our knowledge the only one that attempts to do so specifically for CBDC.

The rest of the paper is structured as follows. After an overview of CBDCs and stablecoins, section 2 covers DLT essentials, and section 3 takes a closer look at different consensus mechanisms. In section 4, we discuss criteria for selecting DLTs for CBDCs, and in section 5, the advantages and disadvantages of permissioned versus public DLTs for CBDCs. Section 6 concludes.

1.1. Background and Purpose

CBDC is not a novel idea. Rather, it is a resurgence of a conversation around digital currencies that began in the 1980s with David Chaum's proposal for digital cash (Chaum, 1983) and James Tobin's proposal to allow individuals to hold deposit accounts in central banks, or commercial banks with access to central bank accounts to offer deposited currency accounts to customers (Tobin, 1985, 1987).

According to Tobin, the government should make available to the public a medium of exchange with the 'convenience of deposits and the safety of currency.' Tobin made his proposal against the backdrop of rapid changes in the structure of monetary and financial institutions due to technological innovations and

² For a non-DLT based CBDC proposal see Chaum & Grothoff & Moser (2021).

private initiatives. His proposal was meant to “protect the system of monetary payments, assure the availability of safe and convenient media of exchange and other assets to the general public, preserve effective macroeconomic monetary control by the Federal Reserve System, and maintain the sovereign power and responsibility of the federal government ... to ‘coin money and regulate the value thereof’” (Tobin, 1987, p. 173). The same motivation is driving central banks today to study CBDCs.

The central question is what role central bank money should play in an increasingly digital and potentially DLT-based economy. At the minimum, central banks must be able to provide financial institutions with access to central bank money so that they can settle their payments with a neutral and risk-free monetary asset. If such settlement occurs on DLT, then the central bank must be able to make central bank money available on DLT. Whether central banks should also make CBDC available to the general public is a different and more controversial question. Arguments in favor include (i) safety, (ii) preservation of currency uniformity, and (iii) the potential for additional functionality.

Safety is Tobin’s argument for a CBDC available to the public, and the argument has even more weight today in an increasingly cashless economy (see, e.g. Fabris, 2019). A CBDC, however, could also foster safety and security beyond Tobin’s concern about credit risk. A CBDC could increase the resilience of the payment system by providing a backup in the event of disruptions to the current electronic banking systems. CBDCs are also a direct response to what many see as the threat from privately issued money and their potential to form a global monopoly on payment systems. A CBDC would be able to introduce competition into such a market and provide citizens with a broader choice, including with regard to data privacy.

The second argument, namely that a CBDC could preserve the uniformity of the currency has been reiterated by Brunnermeier & Landau (2022). By serving as a bridge for converting one private money into another and thereby requiring commercial banks to redeem deposits in central bank money, central bank money guarantees a uniform currency. Regarding the third argument, namely that a CBDC could provide additional functionalities and thus opportunities, the most frequently cited are efficiency gains, financial inclusion, privacy, and more effective monetary and fiscal policies.

1.2. Wholesale versus Retail CBDC

A CBDC designed primarily for the settlement of interbank payments and thus with access limited to selected financial institutions is referred to as wholesale CBDC (w-CBDC), whereas a CBDC accessible to the general public (like banknotes) is referred to as retail CBDC (r-CBDC). The latter has more use cases but is also potentially more disruptive to the existing financial sector and infrastructure.

A w-CBDC operates within the context of the current two-tier banking system without modifying it. Available only to parties that have accounts with the central bank, the w-CBDC would allow for near immediate finality on interbank transfers and cross-border payments. Considering that both banks and selected financial partners already have access to digital central bank money in the form of accounts, a w-CBDC would simply be the tokenization of such money that can be subsequently issued on DLT.³

An r-CBDC would function as a form of digital cash that is issued by the central bank and would have far wider use cases than a w-CBDC. But it would also be more disruptive (e.g. Kaczmarek, 2022). A key argument against r-CBDCs is that they could exert strong competitive pressures on bank deposits, thereby threatening an important source of commercial bank funding, or that an r-CBDC would at least facilitate bank runs. Against this, it has been argued that central banks have a variety of tools to prevent this (e.g. Ben Souissi & Nabi, 2023) or that competition from an r-CBDC would force commercial banks to adapt by offering an improved service to their customers and seek a less risky business model (e.g. Berentsen & Schär, 2018).

1.3. CBDC versus Stablecoins

In many ways, the role a CBDC would play on DLT is currently being filled by stablecoins, which have emerged in response to the volatility of the value of cryptocurrencies such as Bitcoin and Ether. Almost all stablecoins are digital currencies that have value pegged to fiat currency, making them reliant on central banks and central bank money.

³ An example for a w-CBDC that has been tested in a production environment is Project Helvetia, a collaboration between the Swiss National Bank and the financial market infrastructure operator SIX. See Project Helvetia: A multiphase investigation on the settlement of tokenised assets in central bank money. <https://www.bis.org/about/bisih/topics/cbdc/helvetia.htm>

The most widely used stablecoins are known as ‘off-chain asset-backed stablecoins’ and are, as the name suggests, coins that stabilize their value by being backed by highly liquid, conventional assets (hence off-chain) and are redeemable in fiat currency at a fixed price. In this way, they function in much the same way as deposits held at a commercial bank. They are a financial liability of the issuing entity and, accordingly, contain the respective counterparty risk. Since the issuers have a financial incentive to invest a portion of the deposits received for the stablecoins in higher yielding, less liquid assets, stablecoins are typically also subject to liquidity risk. In addition to off-chain asset-backed stablecoins, there are also ‘on-chain asset-backed stablecoins’ backed by other cryptocurrencies or stablecoins, and ‘algorithmic stablecoins’ that attempt to stabilize their value by using algorithms to adjust their supply. However, these types of stablecoins are more volatile and considered less safe.

The collapse of Terra USD (UST) in May 2022 is a clear example of the risk involved in algorithmic stablecoins and the severe effects on the wider market in the event of their collapse. Until its collapse, UST was the largest algorithmic stablecoin in history and widely used across the crypto ecosystem. The collapse sent shockwaves throughout the crypto economy with over \$400 billion in value erased in crypto market capitalization. The fact that the collapse of a stablecoin could have such a devastating effect highlights that stablecoins can pose a huge risk to the markets that they facilitate, and that the scale of this risk is in relation to the clear need that stablecoins fill by serving as a monetary asset. A CBDC — being a risk-free monetary asset — could thus serve a useful role in the crypto economy.

1.4. CBDC versus sCBDC

More relevant to our discussion here is the prospect of what Adrian & Mancini-Griffoli (2019) have called ‘synthetic CBDCs’ (sCBDC), which are stablecoins backed by central bank money. sCBDCs are proposed as a potential middle ground between stablecoins and CBDCs. Central banks could prevent the formation of a monopoly by granting central bank accounts to various stablecoin providers to back their stablecoins. Institutions such as commercial banks that already have access to central bank accounts would be natural candidates for issuing such stablecoins. In Switzerland, such an sCBDC already exists, as the cash leg on the DLT-based, fully regulated SIX Digital Exchange (SDX) is settled via a stablecoin issued by SDX and backed by its account at the Swiss National Bank.

It should be stressed, however, that an sCBDC differs from a CBDC. The difference is that a CBDC is a liability of the central bank, whereas an sCBDC is a liability of the respective private issuer. Consequently, sCBDCs are subject to counterparty risk just like commercial bank money. As part of Project Helvetia Phase III, the Swiss National Bank is currently investigating how a stablecoin backed one-to-one by central bank money can be legally structured in such a way that it has a comparable risk profile to central bank money in the event of bankruptcy of the stablecoin issuer (see Maechler & Moser, 2023).

2. DLT Essentials

In essence, a DLT is a data ledger shared across a network of nodes that store and replicate a consistent state of transactions. Each of these nodes is interconnected and can communicate and share information with one another. Computers, smartphones, servers, or any device that can be configured to communicate with other nodes can function as a node in this network. The decentralized aspect is, of course, that there is no centralized server acting as the intermediary for the communication between the nodes; this allows for a more secure system, as the security of the network does not depend on the integrity of any one node. In this structure, it would take the failure of many nodes at the same time for the integrity of the system to be compromised. The potential of using DLT in central banking comes in its ability to provide the levels of security offered by a closed ledger, with the efficiency and convenience of an open ledger, without compromising on either one.

However, the lack of a centralized player poses a number of problems for DLTs to overcome: how does the system form the consensus required to validate new transactions? How is the validity of these transactions communicated to other nodes? How are the data organized and structured within the network? In a defining response to these issues, the underlying technology that is used by Bitcoin and other cryptocurrencies utilizes a form of DLT known as a Blockchain. A Blockchain functions by storing transaction data in discrete blocks, which are organized in a chain. As the chain grows, each new block is linked to the previous (parent) block via a unique and irreplicable cryptographic hash, which both references the parent block and serves as a unique identifier for itself. Through this system, the verification of each new block also serves to verify both its parent block and subsequently every block in the preceding Blockchain. This means that the security of the Blockchain is strengthened as it grows, making it nearly impossible to tamper with as modifying the data stored in one block would mean having to reconfigure the hash keys of every subsequent block in the chain.

As this entire process occurs without a centralized server or node, DLTs utilize different ways of forming a consensus within their networks. The way any particular DLT achieves this is known as its ‘consensus mechanism’. How consensus is formed is vital to our discussion here, as having varying solutions to the problem of consensus is the primary difference between DLTs. How different DLTs solve this problem will be addressed in greater detail in the next section. First, we must better understand the nature of consensus and why it is so fundamental to the creation of a successful DLT.

2.1. Byzantine Generals Problem

How to form a consensus in a decentralized system was first addressed as a problem in game theory in the form of the ‘Byzantine Generals Problem’ (see Lamport, Shostak & Pease, 1982). This problem is designed to show the difficulty of reaching a consensus in a network without a central party. The problem imagines a scenario where a group of generals and their armies are attacking a city. With the city surrounded, the generals must reach a consensus on when to attack; if they successfully coordinate their attack they will achieve victory, if they do not, they will lose. Crucially, however, some of the generals may be traitors and are trying to send false information to thwart the plan. The loyal generals do not know which or how many of the generals are traitors and can only communicate through messengers that are vulnerable to corruption or interception.

Mapping the analogy onto Blockchains, each general represents a node in the network, the messengers represent signals that are sent between the nodes, and the decision on when to attack represents the validation of a new block. Given that all computer systems must be able to accommodate a limited number of corrupted or failing nodes, the Byzantine Generals problem is foundational to our understanding of DLT. In its simplest form, Lamport et al. (1982) show how consensus can be reached so long as over two-thirds of the generals remain loyal.

2.2. Blockchain Trilemma

The ability of a decentralized network to achieve a consensus under these conditions is known as its Byzantine Fault Tolerance (BFT), and a Blockchain’s consensus mechanism is the way it achieves BFT. However, given the complexity of the processes involved, various solutions vying to solve to different degrees what has been dubbed the ‘Blockchain trilemma’: the preservation of scalability, security and decentralization within the same system.

In brief, the Blockchain trilemma arises out of the number of messages and verifications it takes a Blockchain to achieve consensus. In addressing this, different Blockchain projects typically have to compromise on one of the three vectors. For example, block chains such as Bitcoin and Ethereum are both highly decentralized and secure systems, yet they lack speed, which in turn presents scalability problems, thus making them ill-suited as payment systems in their current form. Other systems that claim to offer speed without compromising security typically compromise decentralization. The pros and cons of the various consensus mechanisms used by DLTs are bound up with how they address this trilemma and will be a key point of focus later in the paper.

2.3. Stress vectors and Vulnerabilities in DLT

As a technology that has been developing at a rapid rate, Blockchain's history of progression and refinement is one of active and dynamic responses to a number of landmark attacks and failures that have since been canonized within the Blockchain space. An understanding of these attacks shines a light on the various vulnerabilities of DLT structures, the ways of safeguarding against them, and a useful way into the nuances of various consensus mechanisms. Furthermore, it is a vital consideration when selecting DLT technology for CBDCs. What follows is an overview of some of the more notorious attacks and vulnerabilities and the ways they have been responded to. However, the vulnerabilities listed below are far from complete. This being the case, it is vital to stress the importance of DLT systems being Byzantine Fault Tolerant, enabling them to tolerate unforeseen errors and attacks

2.3.1. 51% attack

A 51% attack, also known as a majority attack, occurs when an individual or group is able to gain control of 51% of the validator nodes. In this event, the controlling party becomes able to rewrite transaction history as they see fit: confirming transactions that never happened, canceling transactions made in the past, or changing the order of transactions. In the past, this has led to what is known as the double-spending problem, whereby the electronic payment system becomes unable to prove that two or more people did not spend the same digital asset. Such attacks emerge from the decentralized nature of Blockchains, which places pressure on DLTs to compromise on their decentralization.

Since cryptocurrencies using Bitcoin's proof-of-work (PoW) consensus mechanism have suffered 51% attacks⁴, various DLT projects have adapted with adjustments to existing consensus mechanisms. Therefore, DLTs that operate what is known as a dual structure are rendered immune to the attack as they introduce an additional layer through which nodes are verified.⁵ Other Blockchains that were previously vulnerable have grown to such a scale that it is practically impossible for any party to gain majority control. These defenses show that CBDC designs built on fully decentralized DLTs are possible given the appropriate verification infrastructure.

2.3.2. Sybil Attack

The 51% attack is a species of a more common attack within computing known as a Sybil attack that has seen a resurgence with the rise of Blockchains. Named after the case study of Sybil Dorsett, a patient with dissociative identity disorder, a Sybil attack occurs when a malicious actor creates multiple pseudonymous nodes within a system to gain majority control. Once control has been gained, the attacker will be able to manipulate transactions on the Blockchain, much like a 51% attack. The use of protocols such as PoW and proof of stake (PoS) in Blockchain consensus mechanisms are there to prevent Sybil attacks by ensuring that any operator of a node must be sufficiently invested in the system that they have no incentive to corrupt it.

2.3.3. DoS Attack

A Denial of Service (DoS) attack describes any attack designed to prevent legitimate users from accessing an online system. In the case of CBDC, any DoS attack would likely be aimed at preventing users from making transactions and even causing some payments to become lost. In the case of a CBDC becoming a national infrastructure of the scale of a widely used retail payment system, DoS attacks present a route for hostile foreign actors to target a nation's economy.

A DoS attack can occur in a number of ways. In one instance, malicious actors on the inside of the system, such as those operating verification or other permissioned nodes, can abuse their power over the system and deny payments, freeze funds, or make withdrawals without consent. Such functionality is possible with-

⁴ See <https://dci.mit.edu/51-attacks>

⁵ See the DLT overview at the end for more.

in the system to enable compliance procedures, and thus they are at risk of abuse. Alternatively, a more conventional DoS attack entails a group of malicious actors external to the system synchronizing a large number of transactions at the same time, thus overwhelming the system with demand and rendering it unable to process further transactions for a period. Such attacks have been made on large online networks such as Amazon and eBay. DoS attacks can also occur through environmental factors such as flooding or earthquakes affecting a system's servers.

Importantly, any system that is connected to the internet is vulnerable to a DoS attack, and the architecture of any CBDC, whether built on a DLT or centralized server, must safeguard against them. However, the more centralized a system, the more it is at risk. This is obvious in the case of a DoS attack mounted by inside actors, where it is by virtue of the system being centralized that the attackers have privileged access at all. In the case of a more conventional DoS attack, increased decentralization of a system across a large number of nodes prevents the targeting of a centralized point of power. Consequently, a fully permissionless DLT is the most immune to DoS attacks.

2.4. Forks

A fork describes when a Blockchain splits into two separate chains; it is an essential feature in the design of all Blockchains and can occur in a benign way when developers perform updates and maintenance to the technology. There are two categories of forks: 'soft forks' and 'hard forks'. 'Soft forks' typically constitute minor upgrades to the program and are backward compatible with the chain. This backward compatibility means that the network can continue to function as the same network without any major disruption. 'Hard forks' occur when more significant changes to the chain occur, such as an alteration in the underlying chain. In these instances, the chain is no longer backwards compatible and must split, which results in two distinct networks. Historical examples of hard forks in crypto space are the shifts between Ethereum, Ethereum Classic, and Ethereum 2.0. In the case of a hard fork, nodes must then make an active choice to continue validating the new chain.

It is important to understand forking as a necessary and integral part of the operation of any Blockchain, allowing it to perform updates and introduce new features to the network. However, attackers can also abuse forks to conduct a form of digital counterfeiting known as a double-spending attack (see below). In any

case, central banks will have to make use of forks as a mechanism for implementing new regulatory controls, smart contracts and other updates to the system.

2.5. Smart Contracts

One of the major innovations offered by DLT is the development of smart contracts: programs that exist on the Blockchain and are able to automatically control, execute and record predetermined actions that are activated once specific terms or criteria are met. A good illustration of how smart contracts function is the example of vending machines, which are programmed to perform a certain action (dispense confectionery) once predetermined criteria are met (payment). By operating independently on the Blockchain, smart contracts allow for processes that would previously have required mediators or centralized parties to function automatically.

For CBDCs, the use of smart contracts allows for regulatory processes and compliance requirements to be inscribed in code on the Blockchain, allowing central banks more direct control over the assets they issue and to set conditions for their use. This, in turn, allows for significantly improved efficiency in settlements and cross-border payments. As governance requirements are enforced automatically, payments can occur seamlessly with any disputes adjudicated immediately by the code. The benefits of this automation are already being explored in the private sector within the growing DeFi space.

The smart contract options enabled by DLT would also allow for better compatibility of a CBDC as part of a future web3 economy. A clear example of this is the Internet of Things (IoT), a term that encompasses a broad market with potentially millions of use cases as an increasing number of devices, including cars, machines, and even household appliances, become connected to the internet. A DLT-based CBDC would enable interoperability with any markets that also use DLT. With many of these devices set to be connected as part of an electronic payment system that exists on the Blockchain, using DLT is the best way to facilitate interoperability between a CBDC and the nascent IoT market.

Another development facilitated by DLT and anticipated to be the next iteration of the internet is the Metaverse, the name given to a broad swathe of products, applications and hardware that facilitate a continuous, immersive digital world both online, through VR technology, and in the real world through augmented reality. As with the IoT, creating a programmable CBDC via DLT smart contracts is the best way of ensuring interoperability with the Metaverse.

2.6. Virtual Machines

Virtual machines play an integral role in the implementation of smart contracts, as they allow any network participant to access and run a smart contract's code while securing the integrity of that contract. In other words, anyone can execute the contract, but no one can violate its design or change its intended outcomes. Virtual machines are thus a crucial part of any DLT architecture. The design of a DLT's virtual machine has significant consequences for the kinds of smart contracts a particular DLT can support, how well that smart contract functions, who can feasibly program and operate smart contracts. The more coding languages a virtual machine supports, the more accessible the DLT to developers.

Developers must have access to the platform to the extent that it allows them to design, launch and manage functioning programs but not go so far as to be explicitly trusted by the platform. This boundary paradoxically promotes openness and wider accessibility, as it allows untrusted developers into the space without compromising the Blockchain platform's overall functioning or security. One obvious example of the need for such boundaries is the protection of the underlying currency from inflation or counterfeiting. In this way, virtual machines and smart contracts are a vital consideration for central banks in their choice of CBDC, as they work to provide essential boundaries and thus delineate the space in which developers, users and other third parties can adapt and engage with the technology.

2.7. Account-based vs. Unspent Transaction Output (UTXO) Models

The basic process by which DLTs handle the transfer of money between users can be divided into two broad categories: an account-based model and an unspent transaction output (UTXO) model. A Blockchain functions as a state machine in that it records user transactions in the form of blocks. The addition of a new block to the chain constitutes a state transition whereby the balances of the relevant users are updated by the state machine. Both UTXO and account-based models function in this way; the difference occurs in the structure within which the relevant data are arranged and the associated mechanism by which the funds are moved between accounts.

Within an account-based model, states are recorded as a list of accounts that corresponds to a balance. When a transaction takes place, the state change is registered by updating the relevant accounts with the new balances, increasing one and decreasing the other, respectively, much as it would with bank accounts

in the existing banking system. Transactions are triggered only when the paying account presents proof of the authority to do so (e.g. a private key, password, or other proof of identity). In a UTXO system, assets are recorded as individual, set amounts over which a user has ownership. Rather than the more flexible balances in the account-based system, the value of these assets is fixed. Consequently, each new transaction results in the destruction of the original asset and the creation of two new assets, the sum of which makes up the size of the original asset and reflects the amount transferred between users. In this way, where the account-based model reflects the existing structure of bank deposits, a UTXO system is structured in much the same way as cash, whereby a banknote corresponds to a set number that cannot be changed. Just as in the process of paying for a €20 product with a €50 bill, in which the payee hands over the €50 bill and receives €30 in return, transactions in a UTXO model involve replacing the value of the original block with two outputs, one that reflects the value of the payment and another that reimburses the payee with the change.

When comparing the advantages and drawbacks of the UTXO model and account model, it is vital to keep in mind that a UTXO model is a verification model while the account model is a computational model. This means that, where a UTXO model is useful for making simple transactions, an account-based model processes more complex applications. Perhaps most importantly, UTXO models are not ordered, meaning there is the potential for the sequential process of transactions to be manipulated to give some payments an unfair advantage. In contrast, account models can be totally ordered. Outside of these differences, both models offer similar capacities for scalability and privacy.

However, a key advantage of an account-based model for CBDCs is its better support for second-layer functionality, such as smart contracts. As discussed earlier, smart contract functionality is vital for a CBDC looking to facilitate high levels of interoperability. This functionality also makes the CBDC more adaptable to future innovations in the economy, as its more intuitive system allows for better integration with new technologies and forms of payment. This is not to suggest that a UTXO model does not offer smart contract functionality. However, as the account structure allows for a more streamlined system of checking balances and verifying transactions, its simplicity means that the account-based model has better memory storage and exacts less computational power than the UTXO model in checking balances.

3. Overview of Consensus Mechanisms

3.1. Proof of Work (PoW)

PoW — the mechanism first used by Bitcoin — establishes consensus by requiring nodes in the network to compete with one another to solve computational math problems (also known as mining). The first node to solve this problem wins the privilege of storing a set of transactions in the new block, and with it, the financial reward known as the ‘block reward’. This process means that the Blockchain can remain fully decentralized while protecting against attacks by ensuring that participating nodes in the network are sufficiently incentivized to uphold the network.

This form of PoW used by Bitcoin, also known as the Nakamoto Consensus, offered a new way forward for DLT and sparked the recent interest in the potential of digital currencies. However, the Nakamoto Consensus has a number of significant drawbacks, an understanding of which helps illustrate the various innovations and solutions of subsequent DLTs. The core drawback of the Nakamoto Consensus is its lack of finality; Bitcoin transactions are only considered final after six transaction confirmations, after which it is said to have reached ‘probabilistic finality’. This lack of finality plays a key role in Bitcoin’s scalability problem, slowing down transactions to approximately seven transactions per second. Furthermore, not being able to provide absolute finality means that the ledger is at risk of forking, which is when the chain is forced to split into two distinct chains to accommodate unresolvable opinions about a transaction history. Providing absolute finality prevents such ambiguities and safeguards against forks. Another notorious drawback of PoW is its high energy consumption, with the system predicated on miners using large amounts of energy to solve arbitrary computational problems.

3.2. Proof of Stake (PoS)

In the PoS model, a node’s opportunity to act as a block proposer is set in relation to the quantity of the resources they already have invested in the system, as opposed to its ability to solve a computational problem. To some extent, this solves PoW energy consumption, as nodes no longer have to consume large amounts of resources to compete for the block reward. Instead, the system operates a randomizing algorithm to select which nodes will become validators. In traditional models, this algorithm selects randomly, with priority given based on three factors: the amount the node has staked; the age of the coins staked; and a randomized

hash value. Detractors from this system argue that PoS essentially constitutes a centralization of power, as only those already invested in the system can serve as validators. However, given the amount of resources required to competitively mine on some of the larger PoW Blockchains, some have argued that this concern is moot in the discussion around CBDC given the inevitable scale of the system. While PoS does not inherently offer a solution to the scalability problems of PoW, additional functions such as sharding made possible by a PoS architecture propose a new way forward for scalability without compromising on speed.

Despite its resolutions to PoW's energy problem, PoS's architecture exposes it to unique attack vectors such as so-called 'long range' and 'short range' attacks and the 'nothing at stake' problem, which problematizes the way that consensus is formed by having nodes invested in the system. Furthermore, because the majority of tokens must be locked up to protect the network, PoS can cause liquidity crises and a series of unforeseeable security issues. For example, in one instance, the Terra network was forced to halt its operation after its price dropped to nearly zero because it had become extremely cheap to conduct a governance attack. By contrast, PoW avoids this risk because its coins are not part of the consensus mechanism and are consequently irrelevant to the network's safety.

Furthermore, a recent controversy on the PoS protocol Solana highlights many critics' concerns surrounding the PoS governance mechanisms. In this event, Solana's Solend protocol voted to take control of a 'whale' investor's wallet that held \$100 million of the protocol's token, SOL. The reasoning was that the whale investor's huge margin position was dangerously close to causing a crash if they were to liquidate their funds off-chain. In response to this, a governance vote was called to grant Solend Labs the power to liquidate the whale's position through an over-the-counter route. In contrast to the \$100 million held in the whale's wallet, the governance vote was passed with a comparably small \$700k total stake. That the majority of these 'votes' were concentrated in the hands of a relatively small group of investors casts doubt on the extent to which a PoS protocol can be said to be truly decentralized. Furthermore, the takeover suggests that there is an incentive for stakers to launch an attack in the event that there are tokens issued on the Blockchain that are worth more than the native token.

3.3. Delegated Proof of Stake (DPoS)

Typical solutions to the problems of both PoW and PoS involve hybridizing the process to create a consensus mechanism that contains features of both. This involves splitting the process of validation into tiers whereby the process of suggest-

ing state changes to the Blockchain and the subsequent decision to confirm those changes are split between different nodes.

DPOS consensus is one example of this and involves a three-tiered process. The first tier has token holders hold a referendum to choose a minimum requirement of Block Producers. Those elected Block Producers are then delegated to perform the second tier of the consensus by generating blocks. Last, the third tier consists of the block producers conducting Byzantine fault-tolerant consensus, wherein approved blocks can be irreversibly admitted to the Blockchain history. The block may be considered truly confirmed to have completed this process. Proponents of DPOS claim that it can produce a new block every half-second. However, the true confirmation of these blocks cannot be considered final until their consensus has been reached, and that process takes approximately three minutes.

3.4. Proof of Authority (PoA)

PoA bases a node's necessary incentivization on its identity and reputation, rather than the coins it represents or the resources it has invested. In another key difference between PoW and PoS, PoA does not operate a randomized or competitive process for selecting validator nodes but instead relies on a preapproved group of nodes whose job it is to manage state changes and validate transactions. As a consensus mechanism, PoA has proven highly useful in adjudicating private Blockchains, as its small pool of nodes allows for high throughput. However, this again constitutes a far more centralized system, and it is consequently at risk of censorship issues. PoA also forgoes the privacy of its validator nodes as operators in essence stake their identities, a process that arguably places too much power on the validators and puts them and the system at risk of third-party influence. At the extreme, having the identity of the validator node exposed makes a clear point of attack and puts the system at risk of both single points of failure and DDoS attacks. This risk is compounded by the fact that PoA is typically non-transparent and consequently lacks what many see as one of the defining features of Blockchain technology. These pitfalls make it impractical for a CBDC; however, PoA is a viable option for enterprises seeking to operate a localized DLT network in-house.

3.5. Direct Acyclic Graph (DAG)

There are a number of new distributed ledger technologies that do not use traditional Blockchain data structures, favoring instead DAGs, which allow blocks to be authored and accepted simultaneously instead of a single, total-ordered chain.

Hashgraph, Avalanche and Conflux might be the most notable examples of such projects. However, these graphs have weaknesses that limit their utility for CBDC. One shortcoming of DAGs is the recovery of data. If a network participant loses connection to the network for any reason, it is much harder to recover the graph information than it is to retrace the history of a Blockchain. Another central difference between DAGs and conventional Blockchain structures is that DAGs are not linearly ordered, which means that there will always be forks. DAGs thus treat forks as necessary evils to achieve a faster TPS. This creates ambiguities around which is the authentic transaction history. Since transaction confirmation depends on the number of future transactions, there is also no instant transaction finality. Furthermore, the lack of total ordering also presents a significant barrier to support smart contracts. Smart contract executions are supposed to be deterministic, meaning that transaction order is supposed to be unambiguous and does not affect the results. It is difficult for DAGs to meet this requirement.

3.6. Practical Byzantine Fault Tolerance (PBFT)

PBFT marked a significant step forward for consensus mechanisms by providing a solution for BFT that split the decision process into two phases. In the first phase, nodes in the network communicate with one another to confirm that a state transition request has been received. Once a threshold number of nodes has confirmed this, the nodes again communicate with one another to confirm that a second threshold has validated the request. This then allows the network to finalize the state change. PBFT's main innovation is the use of what is called a mesh communication network that has all nodes communicating with one another individually. This process allows for a number of faulty nodes to exist within the network without corrupting the system's integrity as a whole. However, the logistics of this mesh network provide an obvious scalability problem as both the cost and the transaction speed increase as new nodes join the network. Projects that have modified PBFT have also introduced a system whereby the leader node (which creates a proposal for a block and sends it to validator nodes) is rotated after every transaction, rather than only when a problem is detected. This process is known as a view change.

3.7. HotStuff

In a recent innovation, the HotStuff protocol solves what has previously necessitated a trade-off between what is known as the 'responsiveness' of a network and its 'linearity'. In short, linearity is a design feature that means nodes are only charged for linear communication. This economic feature came at the expense

of an integral feature of PBFT systems known as ‘responsiveness’, which refers to the time it takes for a new leader node to be generated by the system, thus compromising the system’s speed. This was the feature of the Tendermint and Caspar projects. The HotStuff protocol resolves this compromise by introducing a third layer into the two-phase system, whereby the process of changing the leader (the view change) is merged with the generation of transaction blocks. This means that nodes can communicate directly with the leader node rather than indirectly via other nodes in the system, greatly reducing the complexity of communicating within the system. HotStuff’s strength as a consensus mechanism comes in its ability to operate a fast and secure system without making the same compromises on decentralization as seen in other mechanisms.

4. Criteria for Choosing DLTs for CBDCs

It is likely that any government-approved digital currency would involve some form of public-private partnership. Indeed, the existing monetary system in most countries constitutes a public-private partnership. When choosing DLT providers, the following basic criteria would have to be considered.

4.1. Transaction Speed

The prospect of a more efficient payment system has been a primary reason for central banks’ interest in DLT options for CBDC. This is particularly the case with cross-border payments, where the obstacles involved in translating between different systems, compliance requirements, and regulations cause significant friction and the need for intermediaries (see, e.g., FSB 2021). Issuing a CBDC on DLT would help to resolve these problems by simplifying the process and providing automated mechanisms in the place of outdated manual processes that are both time-intensive and resource-intensive.⁶

However, while DLT, in theory, offers a vastly more efficient system, transaction speed has proven to be a major obstacle to the scalability and, consequently, the practicality of many Blockchain projects. While it is clear that the technology will soon advance beyond existing legacy systems’ throughput — indeed, some

⁶ Several central banks have collaborated on cross-border CBDC experiments. An example is Project Jura, which was a public-private collaboration involving the Swiss National Bank, the Banque de France, the BIS Innovation Hub Swiss Centre, and a private sector consortium. See: <https://www.bis.org/about/bisih/topics/cbdc/jura.htm>

already have — the problem of transaction speed has meant that the technology necessary for a DLT-based CBDC has lagged behind the concept. To put this problem in perspective, Bitcoin currently has a transaction speed of 7 transactions per second (tps), and Ethereum has a tps of approximately 25. Contrast this with Visa's tps of 1,700 and Mastercard's tps of 5,000, and we obtain a picture of the importance of tps for DLTs to support a viable CBDC. It is important to note, however, that this transaction speed is a problem internal to DLTs, and even with these problems, they still offer a resolution to existing efficiency issues caused by frictions within the existing systems.

In assessing the tps of any DLT, special attention needs to be paid to the specific metrics by which the project has measured its speed and the system by which they achieved it. For example, some projects claim to achieve a high tps, and yet their measurement might only refer to the Blockchain throughput (the number of transactions processed) and not factor in the time it takes to reach finality. This is a poor measure of tps, as only those transactions that achieve finality are considered complete and irreversible. Others will sacrifice processing transactions in a linear order and thus cannot be considered a Blockchain in the strict sense. Knowing the processes behind these metrics is important, as many projects attempt to augment their tps to develop a more viable technology.

4.2. Interoperability

DLT offers significant improvements in interoperability when compared with existing systems. In choosing which DLT to build on, central banks must take into account the high levels of interoperability required of any CBDC. This would, however, be higher for an r-CBDC than a w-CBDC, as an r-CBDC would need to be integrated with not only the existing banking system but also a broader electronic payment ecosystem. Indeed, the level of interoperability required of a fully integrated r-CBDC is set to grow exponentially with the rise of the Metaverse and the Internet of Things, both of which comprise new markets that depend heavily on both electronic payment systems and Blockchain technology. In its broadest sense, interoperability encompasses a future payment landscape that not only includes CBDCs but also other digital assets, such as digital securities, cryptocurrencies and NFTs, and is also essential for enabling the DeFi space.

The degree of interoperability required of a CBDC will also come down to the specific design and policy decisions made by the issuing nations. Domestic CBDCs and international CBDCs both have specific interoperability requirements that are facilitated by DLTs in different ways. Either way, central banks looking to

facilitate interoperability must keep an international perspective in mind when making decisions around design. This perspective starts with key architectural decisions to be made on which DLT to build on, particularly in relation to whether the Blockchain is permissioned or public.

4.3. Cost

The cost considerations of implementing CBDC on DLT occur at three levels. First, there are the internal transaction fees levied by the ledger that vary between different DLT architectures. This is a basic metric and is an easy comparison to make, yet one that should be weighed against any tradeoffs that would occur with other key criteria. The second cost consideration is the extent to which DLT is able to facilitate the economic and efficiency aims of CBDC. With a reduction in the cost of cross-border payments and other transaction frictions in the banking system being a central purpose for the majority of CBDCs, the extent to which a DLT does this is a key criterion. The third cost consideration refers to the necessary infrastructure required by the central bank in the development and maintenance of the CBDC. This is primarily an architecture question and pertains to whether the central bank opts for a fully-fledged CBDC issued and managed by the central bank itself, or another set up in which a central bank-issued CBDC is managed by third parties. At the level of the DLT specifically, the primary factor affecting cost is whether to use a permissioned or public system, the difference between the two is to be assessed on an individual basis between the transaction cost of the specific DLT and the infrastructure costs that come with the increased centralization of a permissioned system.

4.4. Security

Any infrastructure as vital as a CBDC requires the highest levels of security, as protection against domestic and foreign attacks, by both bad actors and potentially hostile regimes. Decentralization as a practice has been held up as a robust security measure, particularly within computational systems, given that the integrity of the system does not hinge on the security of one node but is instead distributed across a network. However, DLT structures have both specific and general vulnerabilities of their own and how various DLT architectures address these is an essential criterion when deciding which to build on. The question of DLT security occurs in terms of policy considerations and at the level of the underlying technology, with both being interdependent on the other. For example, DLT architectures must provide means for administering Know-Your-Customer

(KYC) checks and pathways for implementing regulatory compliance. Ideally, this will be done through smart contracts, which most private DLT projects offer, yet each with its own levels of programmability and accessibility. Consequently, a DLT's smart contract functionality is a key medium through which CBDCs security will be upheld and a key point at which policy requirements intersect with technical design.

5. Permissioned and Public: Two paths for CBDCs on DLT

The options available for issuing a CBDC on a DLT fall under two broad categories: public, which is a fully decentralized system where anyone can operate a node in the network; and permissioned, which is more centralized as access to operating a node is subject to permission by the central bank. In this section, we look at the benefits and pitfalls of issuing a CBDC using public, private and hybrid technologies in relation to broad criteria and examine what each would entail.

5.1. Transparency

While some have argued that permissioned Blockchains protect user privacy over public Blockchains in as far as they conceal specific user data from the majority of users, permissioned nodes and controlling parties still have access to the data, thus providing privileged access to a largely unknown and centralized group. Proponents of public DLT argue that this openness marks the democratization of transparency, as no user has privileged viewing access over any other. This, in turn, allows for greater security, as the ledger can be subjected to a third-party audit at any time from any place.

This transparency is also central to a CBDC's ability to enable faster cross-border payments and better interoperability, as it allows for a simplified approach to compliance that eliminates the need for third-party checks. To achieve similar levels of interoperability offered by public Blockchains, CBDCs built using permissioned Blockchains will likely need to provide privileged access for foreign parties to ensure compliance across regulatory systems. This significantly hampers some of the benefits of DLT.

A crucial caveat to the need for a system that enables transparency is the covalent need for user privacy to be protected. This is a problem faced by both permissioned and public Blockchains and is handled at the level of smart contracts and

digital wallets. Although solving this problem is far from trivial, a number of permissionless Blockchains have proven use cases for operating secure wallets that enable users to view their transactions while ensuring their identity is fully protected, and the same mechanisms can be applied to a CBDC deployment if the system uses a public ledger.

While with private Blockchains, the scale of the threat to user identity is lessened to the extent that it is less transparent to the public, user privacy remains a vital feature to ensure public trust in the system and as a means of protection against potential bad actors on the inside. In either case, both public and permissioned DLTs provide additional user privacy through a range of mechanisms specific to their technology, with most employing a multifaceted system that includes the basic use of pseudonyms, blind signatures, and other advanced cryptographic protocols such as zero-knowledge proofs. While this tension between transparency and privacy exists across all DLTs, it is heightened in CBDCs, given the specific requirements of auditing, regulation and compliance. A solution will either come in the form of cryptographic protocols or in the form of secure hardware.

5.2. Cost

As a means of making existing cross-border payments more cost effective, both permissioned and public DLTs are successful primarily because they simplify existing processes. Within the current system, the complications of enforcing compliance and regulatory requirements necessitate recourse to intermediaries, a process that makes cross-border payments significantly slower and more costly. While both public and permissioned DLTs offer improvements to the existing infrastructure by making cross-border payments more efficient and less expensive, a public DLT is ultimately cheaper.

First, the transparency offered by a public DLT greatly simplifies the process of performing compliance checks, as third-party audits can be conducted at any time from anywhere. For a CBDC issued on a permissioned DLT to achieve the same levels of transparency, it would need to grant privileged access to the ledger for select foreign users so that they can ensure compliance and regulations have been met; however, this relocates rather than solves the problem as it still involves intermediaries. Second, a fully decentralized, permissionless DLT reduces the need for centralized servers required for data management and storage as data and authentication requirements are handled via the participants in the network. A third, associated cost consideration is the extent of the infrastructure that central banks would have to build and maintain given a more managed system

such as a permissioned DLT. Whether outsourced to a third-party provider or managed by the central bank directly, the higher the level of centralization, the greater the managerial cost for the operators. However, these benefits are only possible so long as the transaction fee does not outweigh the cost of operating a centralized system. As such, this advantage is only relevant to public DLTs with low transaction fees.

As with r-CBDCs, CBDCs built on a public DLT are potentially more disruptive to the existing financial infrastructures. Permissioned DLTs offer central banks the option of retaining the traditional system of passing cross-border payments through legacy intermediaries. This way stands as a compromise between a fully public DLT on the one hand and the current system on the other hand by offering some of the efficiency and cost benefits of public DLTs without replacing current intermediaries.

5.3. Trust

As already noted, the extent to which any CBDC aims to serve as a digital replacement for cash has implications for the amount of trust it requires of its users. As it stands, physical cash currently functions as a trustless system whereby recipients of cash have no need to place trust in the payer or a commercial bank to redeem or fulfill their debt. For a CBDC to function as a similarly trustless system, it must be immune to unilateral action by centralized parties, including the central bank itself, intermediaries, or any other node in the network. Public DLT systems enable this by virtue of the fact that no one node in the network has privileged power over any other. Consequently, so long as the system is designed with immunities to the aforementioned threats, such as Sybil attacks and split-view attacks, CBDC users can be assured that the currency they hold is immune from any external actor. A permissioned DLT somewhat compromises this trustlessness, as central banks would have privileged powers over the network and, consequently, unilateral action which, while unlikely, is still a possibility.

5.4. Transaction Speed

As with many of the other factors, the need for intermediaries in permissioned DLTs hampers their ability to enable the near-instant cross border payments that are possible with the technology. Consequently, while on a purely technological basis, public DLTs do not have an advantage over permissioned DLTs in terms of transaction speed, the levels of interoperability enabled by a trustless, transpar-

ent system allow them to reach the full potential of the underlying technology. The efficiency of public networks is further strengthened by having the right to keep records on the ledger distributed across the network, thus avoiding the process of deciding who has record-keeping privileges.

However, it is important to note that the speed advantages of a public Blockchain are limited to cross-border payments. At the purely domestic level, a permissioned Blockchain will always be faster, as its increased centralization allows it to validate transactions without the need for a majority consensus. While it is hard to imagine a viable currency that does not facilitate some form of cross-border payments, permissioned CBDCs have proved particularly effective within closed systems such as exchanges that issue w-CBDCs. Moreover, the control offered by centralization enables more efficient network maintenance and upgrades, as there is no need to establish a consensus throughout the network.

5.5. Security

Public DLTs are generally more secure than both permissioned and private DLTs. This is primarily due to public systems being distributed over significantly more nodes; the fewer the nodes that comprise a network, the easier it is for a malicious actor to gain majority control and manipulate it.

Public DLTs are also censorship resistant, meaning that anyone can transact on the network so long as they comply with its rules and that no transactions can be removed or altered once they have been made. This resistance to censorship means that transactions on public Blockchains are immutable and have been a foundational quality of Blockchains since their inception. CBDCs would need to foster this immutability in as far as they seek to function as a substitute for cash or even to have a payment rail that is insulated against manipulation either by governments or actors within them.

5.6. Interoperability

Public DLTs clearly facilitate greater interoperability when compared with their permissioned and private counterparts and thus allow for a far richer ecosystem of exchange. This is especially pertinent to the issue of cross-border transactions. While it is possible to imagine a purely domestic CBDC that operates on a permissioned DLT and allows developers to operate nodes and build products by granting permissions on a case-by-case basis, this would no longer work when taking into account the sovereignty needs of different nations. As argued throughout

this paper, a permissionless system is the most adaptable and, consequently, the most open to future innovation. However, this shift would also constitute a greater upheaval to the existing financial sector. Interoperability thus highlights one of the tensions that central banks would need to resolve in choosing either a public or a permissioned DLT, where a permissioned system is more accommodating to the legacy banking infrastructure, while a permissionless system would be more accommodating to future innovation.

5.7. Identity

Public and permissionless DLTs offer different capacities for user identity management. When discussing identity and privacy within CBDCs, it is important to clarify it by asking about privacy from whom. Some government and central banks may assume that users are willing to place total trust in them and only want the protections afforded by privacy to apply to other parties. In this case, permissioned and private DLT designs would offer appropriate privacy of accounts and transactions from the wider public yet would be on full display to the privileged parties operating the Blockchain. Such a system would place significant power in the hands of governments and central banks, which, if abused, would enable serious human rights violations. In the instances where the controlling parties can be trusted to act in good faith, the system is still vulnerable in that any attack on the privileged nodes could result in significant data leaks, which would not only allow the attacking parties enormous access to user data but would also constitute a severe corrosion of trust between the public and the central bank.

This being the case, a CBDC that affords the same level of privacy and anonymity as cash (while complying with regulatory requirements) may be preferable. To enable such a system, a CBDC must shield users from the gaze of both the public and the validating nodes themselves. In this instance, public DLTs offer more tools for providing both anonymous and semi-anonymous user identity features. However, given the transparency of the ledger, there have been numerous instances of external parties being able to reveal user identity by tracking spending patterns, with some private entities such as Chainalysis even specializing in deanonymizing users and tracing transactions. Given the vulnerabilities in mere pseudonymization, DLT providers have designed various ways to strengthen user privacy and anonymization, such as randomizing user addresses by using cryptographic algorithms and facilitating offline payments. For sheer convenience, public DLTs have considerable advantages because the various privacy innovations have already been tested and integrated into their technologies.

5.8. Privacy

The BIS report on the future monetary system (BIS, 2022) includes a brief discussion on the uses of DLT for CBDCs. While favoring a permissioned system for what it views as its improved privacy, at the same time, it acknowledges that the privacy offered by a permissioned system is still incomplete, particularly when compared with the anonymity offered by a banknote. This points to the use of innovations such as zero-knowledge proofs, which provide better privacy but can also erode system performance.

However, as noted above, although transactions in permission systems are not publicly visible, permissioned nodes and controlling parties still have access to the data. Permissionless systems can, therefore, potentially offer better privacy protection. For example, zero-knowledge proof architectures provide a solution to both permissioned and permissionless systems. This, in turn, allows a permissionless system to offer greater privacy as it avoids transactions from being viewed by any centralized authority as well as the wider public. Furthermore, recent developments in cryptography such as decentralized ID (DID) solutions provide an alternative solution for resolving the tension between authentication and anonymity. DID has come to be seen as one of the pillars of a future web3 ecosystem for its power to shift the paradigm around data management to one in which users own and are responsible for the management and distribution of their data. Seeking out such privacy solutions within a permissionless architecture is ultimately the more secure option, as it avoids the risks of mass data breaches caused by the single points of failure built into more centralized systems.

5.9. Maintenance

One clear advantage of private and permissioned DLTs is their ability to be updated easily and efficiently by network operators. As discussed throughout this paper, the strength of a permissionless public ledger is its increased decentralization across a high number of nodes, and it is this decentralization that is integral to achieving the full functionality of DLTs. However, this also means that changes and updates to the system require a far higher level of coordination across nodes over which the central bank does not have control. Consequently, in the event of substantial updates that would require a hard fork, updates would be implemented far more easily given the top-down control afforded by centralization.

5.10. Governance

Between public and permissioned DLTs, permissioned DLTs clearly offer the most thorough governance controls, as they allow central banks to vet and monitor validators of nodes to ensure that they uphold the system safely and securely. This provides a necessary level of assurance to customers that an appropriate amount of transactions will be reliably processed, thus bolstering trust in the system. This is more uncertain with public Blockchains, as the processing of transactions is wholly dependent on gas fees (the rewards issued to miners on successful execution of a transaction), and as such, there is no guarantee that transactions will be processed other than the incentivization of the validators.

Another key issue of governance is how to prevent illicit activity on the Blockchain. Much has been said about the need to reign in Blockchain as a favored system for criminal activity, and given the open accessibility of public DLT, there are concerns about who may operate an address on the Blockchain. However, while CBDC issuers cannot prevent someone from establishing an address, they can implement allowed/disallowed addresses in token contracts. This allows central banks the necessary control to bar specific parties from holding and trading CBDCs, if not from operating on the wider Blockchain. Permissioned DLT offers these measures and provides ways to inscribe compliance at the protocol level, meaning that many governance controls can be applied automatically.

5.11. Compliance

Both public and permissioned DLTs provide means through which compliance can be enforced automatically on the Blockchain, thus replacing the need for the considerable human resource power spent on the collection, monitoring and verification of data. In this regard, both systems provide mechanisms that simplify regulatory control by replacing the authority of a centralized institution with economic incentives for verifiers. In both instances, the use of smart contracts plays a key role in automating the role of existing intermediaries.

While much of the literature focuses on mechanisms for enforcing compliance within permissioned models, one advantage of a public Blockchain is its open and uniform standards for all parties. While a 'permissioned' network means greater control on behalf of the operator, the interoperability advantages of a public Blockchain allow for better integration and, consequently, better regulation of adjacent markets. One clear benefit of this is the opportunity to bring the DeFi market under regulatory control.

5.12. ERC Token Standards

An ERC (Ethereum Request for Comments) token standard refers to a set of guidelines or standards in place that smart contracts and applications must adhere to be compatible with each other, thus enabling a viable ecosystem. The most famous of these standards is ERC-20, which sets out the basic guidelines for building smart contracts and applications on Ethereum. Since then, numerous ERC standards have been developed, all establishing additional parameters that enable a specific type of token and that are reverse compatible with the foundational ERC-20 Standard.

Relevant to CBDC design are the ERC-1400 and its extension, the Universal Token Standard. These standards make a CBDC issued on a public Blockchain viable by allowing compliance and regulatory standards to be embedded in the protocol. This, in turn, enables the CBDC platform to be interoperable with other networks using common ERC standards. Functions for regulatory compliance, fractional ownership, fungibility, and issuer-forced transactions are all enabled by ERC-1400. This provides a framework that equips CBDCs with future functionality, making them adaptable to a market in which an increasing number of assets become tokenized.

6. Suggested Framework for CBDCs

The choice between a permissioned or public Blockchain takes the form of a choice between compromising on the functionality of the DLT technology in the case of a permissioned DLT or opting for a more disruptive and less controlled system in the case of a public DLT. Given the newness of the technology, central banks have understandable reservations about the full implications of a CBDC issued on a public ledger. Consequently, the majority of CBDC projects are pursuing permission-based designs. However, given some of the clear advantages offered by public DLTs, central banks may want to transition their CBDCs to public DLTs in the future. Ultimately, any decision around CBDC architecture is a question of government policy and a nation's specific aims and needs.

The ERC 1400 standard makes issuing a CBDC on a public DLT a feasible option for central banks. In addition to the basic provisions set out by the standard, banks will need to ensure that all users comply with KYC and Customer Due Diligence (CDD) requirements in Anti-Money Laundering/Combating the Financing of Terrorism (AML/CFT) regulations. Previous literature has raised concerns that the public key encryption that disguises user identities on most public DLTs

prevents these compliance checks and thus risks enabling illegal trade with unvetted parties. However, some public DLTs in the private sector have shown how compliance can be enabled by using a module that establishes an information channel for exchanging KYC/AML, risk control information, handling fees, exchange rates, and other payment-related information. These can then be adapted by participating banks to ensure compatibility with the participating banks' specific requirements. In addition, some DLTs offer bifurcated identity authentication (for both practical anonymity and retrievable KYC necessary to combat criminality and terrorism) as well as crucial integration with business applications, legacy transaction systems, and parallel Blockchain systems.

To be able to build and maintain smart contracts on the chain, central banks need to run a node as one of the validators. This would then allow a central bank to issue and maintain various features enabled by smart contracts on the ledger, thus providing a contact point between the central bank and its customer base. This also provides an avenue to allow competitive innovation into the space, promising fruitful collaboration with private entities that design useful smart contracts. Furthermore, privileged functions such as mining new tokens or sanctioning addresses can be in the sole control of a central bank through the use of an Admin-Key, which enables a privileged party to build a centralized system at the virtual level while allowing decentralization to be retained at the hardware level.

In setting forth a framework for a CBDC, it is worth considering the use of both public and private DLTs as a way of taking advantage of the benefits of each while avoiding pitfalls. In such a model, the type of technology used will be applicable to the environment to which it is best suited: public DLT for cross-border payments and interoperability scenarios, private DLT for high-frequency, high-privacy transactions, primarily within domestic settings. This hybrid framework ensures that a CBDC will be able to meet the demand, scalability and governance requirements of an r-CBDC in a densely populated country while at the same time being able to take advantage of the full interoperability and transparency advantages of a public DLT for cross-border payments.

In issuing a CBDC on a public Blockchain for cross-border payments, the central bank would operate a node on the DLT network through which it could create and issue smart contracts. Via these smart contracts, the central bank could perform compliance and KYC checks on customers without the need for centralized control over the entire network. Such a model would work for both a w-CBDC and an r-CBDC, allowing for the maximum interoperability benefits offered by Blockchains. The flexibility offered by the public Blockchain further allows for easier integration with other CBDC projects.

7. Conclusion

When selecting DLT for CBDCs, we believe that (i) Byzantine Fault Tolerance, (ii) immediate block finality, and (iii) smart contract support are the most important requirements. Byzantine fault tolerance allows the network to continue to function and reach consensus despite any potential node failure. Immediate block finality requires that only one block will be proposed at any given chain height, which removes the potential for the creation of forks and the possibility that a transaction will have to be undone. Finally, any CBDC seeking to achieve the maximum benefits of the technology will want to utilize smart contracts, and it is unlikely that any CBDC can do without them entirely. Few existing Blockchains can currently meet all these criteria because the Blockchain trilemma is difficult to solve.⁷

Within the scope of DLT, this paper shows that at the technological level, a CBDC issued on a fully public DLT is the best way to achieve the full scope of functionality offered by decentralized technology, allowing for better interoperability, faster cross-border payments, and greater adaptability to an increasingly digitized global economy. However, this must be weighed against the significant disruption that a public DLT poses to both existing financial systems and central banks.

⁷ Existing Blockchains that meet all the mentioned criteria include Ethereum, Cypherium, Hyperledger Fabric, and R3 Corda. Among other existing Blockchains, Solana is not Byzantine Fault Tolerant, Avalanche is not linear but a DAG, which does not support absolute finality, and Stellar does not currently support smart contracts.

References

1. Adrian, T. and Mancini-Griffoli, T. (2019). The Rise of Digital Money, IMF Fintech Note 19/01.
2. Allen, S., Čapkun, S., Eyal, I., Fanti, G., Ford, B. A., Grimmelmann, J., Juels, A., Kostiainen, K., Meiklejohn, S., Miller, A., Prasad, E., Wüst, K. and Zhang, F. (2020). Design Choices for Central Bank Digital Currency: Policy and Technical Considerations. NBER Working Paper no. 27634.
3. Auer, R., Cornelli, G. and Frost, J. (2020). Rise of the Central Bank Digital Currencies: Drivers, Approaches and Technologies. CEPR Discussion Paper No. DP15363.
4. Bank for International Settlements (BIS) (2018). Central Bank Digital Currencies. Joint Report of the Committee on Payments and Market Infrastructures and Markets Committee.
5. Bank for International Settlements (BIS) (2020). Central bank digital currencies: foundational principles and core features. <https://www.bis.org/publ/othp33.pdf>.
6. Bank for International Settlements (BIS) (2021). Central bank digital currencies: system design and interoperability. https://www.bis.org/publ/othp42_system_design.pdf.
7. Bank for International Settlements (BIS) (2022). The future monetary system, Box C. Annual Economic Report, 95-96.
8. Ben Souissi, S. & Nabi, M. S. (2023). Could the Issuance of CBDC Reduce the Likelihood of Banking Panic? *Journal of Central Banking Theory and Practice*, 12 (2), 83-101.
9. Berentsen, A. and Schär, F. (2018). The Case for Central Bank Electronic Money and the Non-case for Central Bank Cryptocurrencies. *Federal Reserve Bank of St. Louis Review*, 100 (2), 97-106.
10. Brunnermeier, M. and Landau, J.-P. (2022). The digital euro: policy implications and perspectives. Publication for the committee on Economic and Monetary Affairs, Policy Department for Economic, Scientific and Quality of Life Policies, European Parliament, Luxembourg.
11. Chaum, D. (1983). "Blind Signatures for Untraceable Payments", In *Advances in Cryptology*, ed. Chaum D & Rivest, R. L. & Sherman, A.T., Springer, Boston, MA.
12. Chaum, D., Grothoff, C. and Moser, T. (2021). How to Issue a Central Bank Digital Currency, SNB Working Papers, 3/2021.
13. Chowdhury, M.J.M., Ferdous, S., Biswas, K., Chowdhury, N., Kayes, A.S.M., Alazab, M. and Watters, P. (2019). A Comparative Analysis of Distributed Ledger Technology Platforms, *IEEE Access*, 7, 167930-167943.

14. Fabris, N. (2019). Cashless Society – The Future of Money or a Utopia? *Journal of Central Banking Theory and Practice*, Vol. 8 (1), 53-66.
15. Financial Stability Board (FSB). (2021). G20 Roadmap for Enhancing Cross-border Payments: First consolidated progress report. October 2021.
16. Kaczmarek, P. (2022). Central Bank Digital Currency: Scenarios of Implementation and Potential Consequences for Monetary System. *Journal of Central Banking Theory and Practice*, 11 (3), 137-154.
17. Kubler, S., Renard, M., Ghatpande, S., Georges, J.-P. and Le Traon, Y. (2023). Decision support system for blockchain (DLT) platform selection based on ITU recommendations: A systematic literature review approach, *Expert Systems with Applications*, 211, 118704.
18. Lamport, L., Shostak, R. and Pease, M. (1982). The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems*, 4 (3), 382–401.
19. Maechler, A. and Moser, T. (2023). Swiss Payments Vision – an ecosystem for future-proof payments. Money Market Event, Zurich, 30 March 2023. Available at www.snb.ch.
20. Rauchs, M., Glidden, A., Gordon, B., Pieters, G. C., Recanatini, M., Rostand, F., Vagneur, K. and Zhang, B. Z. (2018). Distributed Ledger Technology Systems: A Conceptual Framework.
21. Tobin, J. (1985). Financial innovation and deregulation in perspective. *Bank of Japan Monetary and Economic Studies*, 3(2), 19-29.
22. Tobin, J. (1987). “The case for preserving regulatory distinctions,” In: *Restructuring the Financial System, Proceedings of the Economic Policy Symposium*, Jackson Hole, Federal Reserve Bank of Kansas City, 167-183.