

Wholesale Central Bank Digital Currency: A survey and analysis of central bank pilot exercises*

Rhys Bidder¹

¹KBS, QCGBF

January 2022

*Views expressed in this document, and any errors or omissions, should not be interpreted as those of Qatar Central Bank or the Bank of England. Nothing in this report should be construed as any form of financial advice.

Contents

1	Summary	3
2	Introduction	10
3	Background	11
3.1	The changing financial landscape	11
3.2	Motivations for exploring wCBDC	13
3.3	Existing pilots	16
4	Key functionality in wCBDC pilots	19
4.1	Liquidity and the security lifecycle	19
4.2	DvP settlement and ledger interoperability	23
5	Performance, privacy and resilience	28
5.1	Performance	28
5.2	Privacy	32
5.3	Resilience	36
6	Legal and policy issues	40
6.1	Legal matters	40
6.2	Monetary Policy	46
	Appendices	53
A	Blockchain features	54
A.1	Permissioned and permissionless blockchains	54
A.2	Consensus and finality	54
A.3	Tokens and accounts: Fixing terminology	55
B	Terminology: DvP, Atomic Settlement, Immediate Settlement...	58
C	Project Khokha - Phase 2	59
D	Providers	61
D.1	Conda by R3	61
D.2	Hyperledger Fabric	63
D.3	Quorum by Consensus	64

1 Summary

Background

The private sector is advancing rapidly in its adoption of distributed ledger technology (DLT) and applications based upon it, owing to apparent cost and process advantages. While there are many potential benefits of these innovations there may be a role for policymakers to intervene for reasons of financial stability and market efficiency.

From the financial stability perspective, a key worry is that privately issued moneys (such as stablecoin) might replace central bank-issued money (reserves) in wholesale payments and securities settlement systems. This would be especially concerning in systemically important markets.

Central banks also have an interest in promoting efficient market functioning from day to day. If multiple private moneys emerge, overall liquidity - which previously had been focused in reserves - may fragment, increasing costs and complexity.

Providing wholesale CBDC (wCBDC) could be a response to these concerns as a riskless and perfectly liquid form of ‘central bank money’, usable in the DLT context. It could then, perhaps, ‘crowd out’ damaging, uncoordinated private sector moneys from markets whose stability and efficiency are paramount.

More positively, *desirable* DLT-based innovation and advances in financial market infrastructure may be held back by the absence of a wCBDC.

Given these motivations, many pilots have been run to explore feasibility and desirability of wCBDC. It is common to group these pilots in ‘waves’, as shown in box 1 and discussed in greater detail in section 3.3.

Key functionality in wCBDC pilots

Liquidity and the security lifecycle

Liquidity management is key in large value payments and securities settlement systems. Various mechanisms help minimize the amount of (costly) liquidity participants must hold in order to ensure smooth operation of the payments system.

‘Liquidity savings mechanisms’ (LSMs) introduce queuing of transactions and some limited multi-lateral netting to reduce the gross payments parties must make. ‘Gridlock resolution mechanisms’ handle situations where, despite *overall* liquidity among a group of counterparties being sufficient, no one transaction can settle without others settling beforehand. For more extreme cases where overall liquidity is inadequate, some form of intraday credit may be needed.

In section 4.1 we discuss how pilots have implemented such schemes. Some challenges emerged

Wave 1: Interbank payments using public blockchains with limited functionality

- *Jasper₁* (2017), *Ubin₁* (2017), *BCB* (2017)
- Developing understanding of basic interbank payments systems using DLT, re-implementing basic functionality of legacy reserves and RTGS systems
- Based on public blockchain technology (Ethereum) with limited realism, especially in the lack of transaction finality and confidentiality

Wave 2: Interbank payments using private blockchains and more advanced mechanisms

- *Stella₁* (2017), *Ubin₂* (2017), *Jasper₂* (2017), *Khokha₁* (2018), *Inthanon₁* (2019)
- Permissioned blockchains with enhanced confidentiality and immediate finality, more suited to real world financial transactions
- More realistic auxiliary functions beyond simply settlement (liquidity savings mechanisms and gridlock resolution)

Wave 3: Securities settlement and ledger interoperability

- *Stella₂* (2018), *Ubin₃* (2018), *Inthanon₂* (2019), *Helvetia₁* (2020), *Khokha₂* (2022)
- Delivery-vs-Payment settlement of tokenized securities for wCBDC
- Methods to enable interoperability of ledgers - so securities and money could be on different blockchains
- Further realism added through post-trade security lifecycle events (coupons, repo)

Wave 4: Exploring a DLT ecosystem

- *Ubin₅* (2021), *Khokha₂* (2022), *eAUD* (2022)
- Broadening class of participants/applications
- Greater private sector input on use cases
- Incorporating wCBDC in a more elaborate interoperable network of systems

Box 1: Waves of domestic wholesale CBDC pilots.

depending on the particular platforms used in the pilots and how the central bank traded off different characteristics of the systems. Gridlock resolution in particular presented trade-offs in the DLT context between a desire to decentralize the system (reducing reliance on the central bank), confidentiality of transactions and speed of execution.¹ Nevertheless, the systems have been successfully implemented with performance comparable to and even beyond legacy systems.

With the use of smart contracts, systems have also implemented standard events in tokenized securities' lifecycles (coupon and principle payments, for example).

DvP settlement and ledger interoperability

A 'delivery-vs-payment' (DvP) structure to settlement is typically thought desirable, all else equal. Thus, pilots have aimed to design systems where (tokenized) securities are delivered if and only if payment occurs. If the security and the settlement asset are on the same ledger, then smart contracts can implement DvP in a straightforward manner. If wCBDC is on one ledger and the security is on another, then matters are somewhat more complex and the interoperability of ledgers must be addressed. We examine these issues in section 4.2.

If a framework aims to have wCBDC issued directly on the same ledger as a security (an 'on-ledger' approach) then there may be benefits in terms of simplifying and coordinating transactions among the securities on that particular ledger. In many cases instantaneous atomic settlement should be possible and liquidity may therefore benefit as assets are not locked up for any extended settlement window and intermediate steps are reduced. Several pilots have indeed trialled such a framework and shown its feasibility.

An alternative, 'interoperable-ledger', approach, is for a central bank to issue wCBDC on a single ledger and then allow for interoperability between this ledger and any other ledgers where securities reside. This might allow the central bank to focus its resources and administrative energies on a single ledger, and delegate the creation of an appropriate connection between ledgers to other, perhaps more expert parties. This approach has also been experimented with in various pilots, and can be implemented in various ways (see section 4.2.2 for some of the techniques).

Pilots that adopted the interoperable-ledger approach noted some additional complexity (legal and technical) involved, with some approaches featuring delays and greater difficulty of achieving atomic exchange, relative to the on-ledger approach. However, enabling interoperability among blockchains is an active research area, motivated by growing industry demands for interoperability, well beyond CBDC.

¹In fact, these trade-offs repeatedly emerge throughout the wCBDC pilots in some guise or another. See section 5 for extended discussions of this trio of features.

Performance, privacy, and resilience

Performance, privacy and resilience are key high-level characteristics of wCBDC systems that must be assessed. In some respects they are in tension and a wCBDC will ultimately have to trade them off according to country-specific costs and benefits. Nevertheless, over the recent years of wCBDC experimentation, technological advances have mitigated these trade-offs. We discuss these issues in section 5.

Performance

Permissioned DLT systems, designed for enterprise usage, have shown that *performance* levels of legacy payments systems can be attained and even exceeded, as discussed in section 5.1. Importantly, this performance is consistent with a large degree of decentralization and privacy protections. Much of the ability to improve performance over the earliest pilots came from blockchain consensus mechanisms and network structures that allowed transactions to be rapidly committed to the ledger with complete finality (see appendices A and D).

Some pilots are more rigorous than others, but the systematic assessment of performance still seems something of an afterthought. This is especially the case in the later waves of pilots, focusing on settlement and ledger interoperability.

Privacy

Privacy - or confidentiality - is vital in wholesale systems, given the sensitivity and value of transactions. Early trials using simple Ethereum implementations (knowingly) ignored this issue while basic DLT knowledge was acquired and simple transactions were executed. Soon, however, ensuring confidentiality became - and remains - a key requirement for wCBDC pilots.

Some of the most prominent DLT platforms relied in the first instance on dividing networks into sub-networks of participants to enhance privacy, though this can still allow information leakage, depending on the particular platform used. It can also introduce layers of complexity for participants due to the need to coordinate behavior across sub-networks. A whole array of complementary techniques can be used to close (or narrow) the gaps remaining after simply utilizing sub-networks. We discuss several in section 5.2 (and appendix D). Some approaches - such as those using ‘zero knowledge proofs’ - may enable a high degree of privacy, even without an elaborate sub-network framework.

It is possible to shield information from the central bank, as well as from other participants. If desired, central banks (or some other regulatory authority) operating administrative nodes in a blockchain network can be restricted to validating transactions without seeing any interpretable details of the transaction. Partial revelation of information, such as to allow checks for KYC/AML, has also been demonstrated.

Resilience

Various characteristics of a system might influence *resilience* but decentralization and the avoidance of a single point of failure (SPF) are emphasized in the wCBDC pilots. Reducing reliance of processes on the central bank has been a particular focus.

The platforms used for wCBDC pilots typically exhibit significant *centralization* relative to well known public blockchains. Partly this stems from the influential position held by a central bank in terms of issuance, rule-setting and oversight. Nevertheless, much progress has been made in limiting reliance on the central bank being online. Resilience to failures of a single or small set of market participants has also been demonstrated (see section 5.3).

The avoidance of a single point of failure (SPF) is a key benefit of a decentralized system. In designs where the ledger is recorded by multiple parties, even if one node goes offline, the information saved by the offline node is not trapped, but can be reliably sourced from elsewhere, enabling continuity for active nodes and rapid recovery for the failed node.

Even some complicated multilateral processes (such as gridlock resolution) which in legacy systems are highly centralized, have also been decentralized to a significant degree. However, this - and the degree to which privacy and performance are sacrificed - is platform dependent. Some areas, however, remain difficult to decentralize, such as arbitration in the case of disputes, or where technical failure of counterparties has disrupted transactions, requiring third party intervention.

While pilots have made great progress in reducing SPF-risks in several dimensions, SPF risks may re-emerge ‘through the back door’ depending on the DLT implementation chosen. The benefits of decentralized record keeping might be undermined by certain features of DLT systems designed to enhance privacy and performance. For example, if data is stored off-chain to enhance confidentiality, it may be more difficult to reliably reconstitute a complete record after a system or node failure. Alternatively, different consensus mechanisms may imply varying degrees of centralization.

Legal and policy issues

Legal matters

Pilots have focused on the practical feasibility of introducing a wCBDC and exploring core *technical* mechanisms. As such, legal details have received lesser attention. As the sophistication and extent of pilots continues to grow, the time will soon come when the legal framework will need to be formalized (see section 6.1).

The legal status and even the terminology around tokens are not settled, owing to the novelty of the technology. In most jurisdictions’ existing legal frameworks, the status of digital assets is extremely unclear, if they are acknowledged at all. As such, tokens may fail to be covered by

any reliable legal structures, or interact in unexpected ways with laws not designed with them in mind.

Tokenized securities that relate to an underlying asset seem likely to raise more legal difficulties than native tokens, which are themselves ‘the asset’. Tokenized securities connect one asset with another and the nature of that connection (in terms of ownership, control, finality in transactions and behavior in bankruptcy...) is at this point especially ambiguous. However, both types of token seem to lack a fully worked out legal basis in many jurisdictions.

Many of these issues apply not only to tokenized securities issued by private parties, but also to wCBDC itself. Different models of issuance have been considered in different pilots, with apparently different levels of legal complexity were they to be codified in a production-ready launch. In particular, wCBDC’s use in interoperable ledgers raises subtle issues, especially with regard to legal vs. technical finality of settlement. The interpretation and legal standing of ported wCBDC on non-native ledgers (not the ledgers on which the wCBDC was initially issued) will require careful consideration.

None of the pilots suggest that formalizing a legal framework will be impossible, but simply that it must be done. Some countries (Switzerland, for example) have made more progress than others, and offer a model that could be informative for other countries’ approaches - even if solutions will need to differ according to countries’ idiosyncratic characteristics. For a start, in some countries there is even a question over whether the central bank has authority to issue CBDC.

Monetary policy

Debates over monetary policy are far more common in the *retail* CBDC domain. Nevertheless, some issues arise in the wCBDC case that are relevant for monetary policy debates (see section 6.2).

In the absence of intraday liquidity provision, all pilots seem to have designed their wCBDC issuance in a way that leaves the balance sheet size unchanged. Composition may change (fewer reserves, more wCBDC) but the overall amount of central bank money remains fixed. If wCBDC is offered as part of an intraday repo liquidity facility, then the money supply would vary. But the same would apply in legacy systems with *reserves* being exchanged for collateral and these operations are typically not thought of as indicating a monetary policy stance, *per se*.

In some pilots wCBDC did not have to be redeemed at the end of the day, requiring a decision on the rate at which it would be remunerated overnight. Some pilots featured positive remuneration and others left the balances unremunerated. Clearly, if reserves are being remunerated then the relative rates will have implications for what sort of central bank money participants would want to hold.

There may be some risk of fragmentation of liquidity arising from the co-existence of two forms of

wholesale central bank money - reserves and wCBDC. One could envisage banks separating into groups on the basis of which of the two 'money markets' they would predominantly operate in, with the interaction between the two markets subject to frictions. If this were to substantially affect money-market efficiency then *conceivably* the transmission of policy could be affected.

A broader debate over the optimal split between reserves and wCBDC, given an overall amount of central bank money will at some point need to be had, if wCBDCs begin to be introduced, but this is beyond a pilot's scope.

2 Introduction

In recent years, the revolution in distributed ledger technology (DLT) has disrupted many industries and business processes. Financial markets and payments systems in particular have seen especially rapid change. Central banks and financial regulators have thus had to adapt and formulate policy in response, while simultaneously learning about the new paradigms.

In this context, the debate over Central Bank Digital Currency (CBDC) is especially prominent. Many central banks have launched pilots, and in some cases deployed fully fledged systems. Many more are examining whether a CBDC may be desirable and, if so, what form it should take. In terms of form, two simple distinctions can be made: first, between CBDCs designed for only domestic use versus those that are part of a multilateral, cross border system and, second, between CBDCs designed for ‘wholesale’ use (among banks and large financial institutions versus ‘retail’ or ‘general purpose’ use (among households and non-financial firms).

This report is a survey of the current state of knowledge regarding *domestic, wholesale* CBDC. We draw on the experience central banks across the world now have from running pilots. Apart from describing these pilots - which in itself is useful - we elicit insights about how CBDC could be implemented in the future. Nevertheless, there are significant overlaps between the topics we discuss and those that are relevant for retail or cross-border CBDC. Throughout the report, an emphasis is put on tying the analysis back to practical examples from existing pilots.

First, we describe the background to our analysis by discussing the (rapidly) changing financial landscape, how it has motivated interest in wCBDC, and what pilots have already been launched by central banks to explore it.

We then discuss some key functionality of payments and security settlement platforms and how they can be implemented in the wCBDC context. These comprise liquidity management, the administration of securities’ lifecycles, delivery-versus-payment settlement, and the ledger interoperability. We then discuss slightly higher-level issues of performance, privacy and resilience. We group these issues together because, not only are they individually important, but there can be trade-offs and tensions between the three.

Finally, we take a less technical perspective and consider some of the legal and monetary policy issues that are relevant to wCBDC.

3 Background

3.1 The changing financial landscape

Distributed ledger technology - or DLT - applications are spreading rapidly within the private sector in general, and financial markets in particular.² There are several interrelated factors driving these rapid developments:

- **Cost:** Cost reduction through via the elimination of duplicated record keeping and reconciliation between parties using different (and often incompatible) back-office databases.
- **Consensus:** A single, synchronized and mutually agreed-upon record offers enormous potential benefits in avoiding confusion, disagreement and legal wrangling.
- **Permissioning:** While public systems, such as Bitcoin and Ethereum, have captured much publicity, recently developed ‘permissioned’ or ‘private’ systems that are more suited to the confidential nature of financial transactions.
- **Resilience:** Failures of individual parties can be tolerated and quickly recovered from, allowing resources to be released from ensuring the reliability single points of failure.
- **Speed:** Reforms aimed at shortening settlement times, advocated especially since the financial crisis, can be taken even further.
- **Automation:** Smart contracts may allow more complex business and regulatory logic to be incorporated into transactions.
- **Consolidation:** Intermediate stages in transactions and various market infrastructure platforms can be compacted or eliminated.
- **Decentralization:** Reliance on central authorities and intermediaries can be reduced or even eliminated.
- **Tokenization:** Tokenization of existing assets can unlock assets that are difficult to trade on traditional platforms, or allow transactions involving assets that *can* already be easily traded through legacy systems, to be traded in novel ways.

While all the above factors have played important roles in driving private sector DLT adoption, perhaps the key has been the ability to tokenize assets. This, combined with DLT implementations, offers a revolutionary path for the financial sector. With tokenized assets issued and recorded on distributed ledgers, new models of exchange and settlement are possible. In fact, they are not only possible but are already emerging.

²While blockchains are a particular case of DLT, we use the terms interchangeably throughout. See appendix A for discussion of core blockchain features.

Let us briefly consider a few illustrative examples focusing on banks and financial market infrastructure, given their obvious relevance to central banks and regulators:

- Project [Spunta](#), instituted by the Italian Banking Association, with [R3](#) as technological partner, is a private sector-led initiative to modernize interbank payments, using DLT, and has been operational since 2020.
- In Australia, CBA helped the World Bank to issue ‘bond-i’ and then enable [secondary market trading](#) in it. CBA was also involved in the [Lygon](#) consortium that digitized and automated bank guarantees.
- J.P. Morgan now operates a blockchain unit ([Onyx](#)), having been heavily involved in the initial development of the [Quorum](#) blockchain protocol, sold to ConsenSys in 2020. Among the products supported by Onyx, the bank provides [intraday repo](#) services on a DLT platform, with tokenised collateral exchanged atomically against [JPM Coin](#), the stablecoin the bank provides for settlement among its clients.
- In some jurisdictions, entirely ‘digital banks’ are now emerging - with [Sygnum](#) holding a banking license in Switzerland.
- In Switzerland, the SIX Digital Exchange ([SDX](#)) provides a regulated platform for the issuance, trading and custody of tokenized assets. Their services are already being used - such as in the recent [digital bond issuance](#) by UBS. Notably, SDX offers a 100% backed Swiss franc ‘stablecoin’ that is a liability of SDX and which offers the promise of becoming a reliable settlement asset for large volumes of securities trading.
- In the US, the influential Depository Trust and Clearing Corporation ([DTCC](#)) has launched [Project Ion](#) - a DLT platform for settlement of cash-equity trades - as part of its broader goal of modernizing and accelerating settlement in financial markets.

While many banks and other financial institutions with sufficient expertise may be able to develop their own systems, there is a growing set of specialist technology companies and consortia that are helping to design and deliver many of the technical solutions currently being adopted. Notable among these are [R3](#), [ConsenSys](#), the [Hyperledger Foundation](#), and [Digital Asset](#). These companies provide permissioned blockchain frameworks that are adaptable to particular use cases and designed for enterprise applications. As their experience has grown, their products have been refined and the fixed costs to companies of launching blockchain systems has declined substantially, contributing to their rapid spread.

Another important trend within payments and financial markets is the emergence of companies and platforms focused on connecting disparate systems, providing interfaces that can tie together the myriad ledgers, allowing tokenized assets - securities and digital coins - to be exchanged. Companies such as [Fnality](#) and [HQLA^X](#) have emerged to help provide the glue between what has

hitherto been a somewhat dispersed system of ledgers. For example, these two companies (aided by Santander, Goldman Sachs and UBS) recently executed a [cross chain repo swap](#), between R3's Corda and an Enterprise Ethereum implementation - that is, between two distinct DLT. Tools and software such [Cosmos](#) and [Polkadot](#) are also gaining increasing prominence in achieving interoperability.

Many more examples of the spread of DLT in wholesale markets emerge each day, but the subset above are representative of some of the key dimensions in which the private sector is advancing rapidly.

3.2 Motivations for exploring wCBDC

Along with the potential benefits of these private sector advances, come several potential risks. These risks relate to financial stability, payments efficiency and industrial structure. Together they justify the exploration of wholesale CBDC (wCBDC) and development of expertise within central banks that will allow them to better engage with crypto issues in general.

3.2.1 Financial stability

Recent developments have made clear that private digital money could play the role of settlement asset in large value transactions in systemically important systems - interbank payments and the trading of securities. In particular, the emergence of private stablecoin, ostensibly 100% backed by fiat or high quality liquid assets, is a striking development. Provision of a reliable settlement asset has traditionally been thought of as the domain of the central bank, with this money taking the form of reserves. Private settlement money - such as stablecoins - offer a form of liquidity that could conceivably compete with central bank money. However, movement away from using central bank money in settlement may open the door to significant systemic risks ([BIS \(2012\)](#)).

Central bank money in the form of reserves is unparalleled in its liquidity, its promise of finality for those transactions settled in it, and its absence of any credit risk. This reflects reserves' backing by a central bank that can print its own money ([BIS \(2003\)](#)). They therefore do not add to principal and replacement cost risks in trades settled with them. In contrast, private moneys may exhibit credit and liquidity risks arising from their 'sponsors' or the assets (or algorithms) that back them. These imperfections may be tolerated by private agents, however, when traded off against superior *private* benefits of usage (such as in smart contract integration and broader access), which reserves may lack. But the *social* costs from systemic crises might not be appropriately internalized, leading to socially inefficient fragility in key markets.

Memories are still fresh from the '[run on repo](#)' and the problems in the last financial crisis arising from private sector provision of money-like assets to satisfy demand for liquidity. As such, some have argued that provision of liquid assets by public authorities could be desirable in 'crowding out'

the private sector. The private sector could then concentrate on other areas of digital innovation where they have comparative advantage.³ Providing wCBDC could be a modern manifestation of this idea.

Arguably there are *some* inherent advantages to DLT and tokenization in achieving a safe system and some of the existing guardrails for legacy systems may not be necessary in the new model. For example, the sort of atomic settlement that is feasible through the combination of DLT, tokenization and smart contracts likely reduces the need for protections against principal risk. As another example, smart contracts might be able to encode regulatory requirements into transactions, rather than having them checked *ex post* or at arbitrary times (Auer (2019)).

However, allowing new systems to emerge in a regulatory vacuum and without official engagement is clearly inadvisable. Exploring the issuance of wCBDC as a reliable settlement asset is one step towards enhancing safety (BIS (2012)). More broadly, pilots should enable broader interaction and communication between the private sector and the central bank/regulators. Mutual understanding and shared expertise are vital to establish appropriate guardrails while the DLT ecosystem is - though rapidly growing - still of manageable size and limited systemic importance.

3.2.2 Payments efficiency

Setting aside the financial stability aspects of private tokenization, its spread also raises issues regarding payments efficiency. As private moneys emerge it is possible that their issuers may obtain a competitive advantage such that market power could be exploited and future innovation could be suppressed. For example, the provision of a private settlement asset may permit an informational advantage through the monitoring of transactions, or can be bundled with other services to cross subsidize and undermine competition in other business areas. Alternatively, tiering of access could lead to market power concerns and even a possible re-emergence of a form of correspondent banking.

Another risk is the possibility of uncoordinated proliferation of currencies, which may impose costs through the fracturing of overall liquidity. A failure to internalize network externalities from issuing private moneys could lead to inefficiency. Enormous strides have been made in reducing liquidity fragmentation in various jurisdictions in recent years and the fact that so many reforms - notably those around the Euro's TARGET systems - have been necessary should be cautionary.

It is not automatic that tokenized assets on DLT will avoid the same pitfalls that these reforms have been designed to fix. Far better for policymakers to be involved in coordination and oversight from an early stage - *prevention is better than a cure*. Thus, there may be a role for wCBDC as a settlement asset with no concerns over the exploitation of market power or information advantages, and with a coordinating role to maximize the size of the 'liquidity pool'.

³See Krishnamurthy and Vissing-Jorgensen (2012), Krishnamurthy and Vissing-Jorgensen (2015) and Gorton and Metrick (2012) for discussions of these issues in a more general setting, and Gorton and Zhang (2023) for explicit reference to stablecoin.

3.2.3 Promoting innovation

The absence of an engaged central bank may also hold back desirable activity in the private sector. The private sector is, in fact, keen for a wCBDC to be introduced, in order to ‘close the circle’ of settlement. At the moment central bank money cannot be used for settlement on ledgers where tokenized assets are traded, in contrast to traditional systems where central bank money is the principal payment option. Thus, parties must rely on stablecoins and other cryptocurrencies, or delivery-versus-delivery of one type of security against another. There is no way to complete the entire life cycle of exchange purely on DLT frameworks, ending with parties ultimately holding central bank money, or using it in intermediate steps.

Clearly, parties frequently require conversion to central bank money, for various reasons, and this step partly defeats the object and undermines the benefits of transitioning to a DLT framework. Indeed, in *Jasper₂* (2017), providing ‘the last mile’ of settlement is referred to as ‘*likely to be the primary benefit of introducing a DLT interbank cash payment platform*’. The role of central banks as enablers of further innovation is explicitly recognized by important parties within the private sector:

The use of a high-quality, safe and stable CBDC for the settlement of digital asset transactions would unleash the potential of the emerging digital assets ecosystem while reducing the overall risk profile of such transactions.

[David Newns](#), Head of SDX, May 2022

and

Without digital cash to facilitate true [Delivery-versus-Payment] on and across blockchains, we will not realise the full potential of the digital revolution in financial markets. This is where central bank digital currency or CBDC comes into the equation.

[Tom Phillips and Paul Pirie](#), JP Morgan Chase, October 2022

Demand for a CBDC comes also from outside the set of institutions that traditionally have access to reserve and large value payment systems. It is plausible that innovation by such firms may be a key benefit of the spread of DLT solutions. This may be another channel through which introducing wCBDC could enhance the financial system and broader welfare. However, there is the additional issue of what set of participants should have access, and care must be taken to ensure the safety that has come to be expected in legacy systems is ported to newly emerging frameworks with - perhaps - broader participation.

3.3 Existing pilots

Progress on wCBDC so far has been made in waves.⁴ We here briefly discuss these approximate stages of development. In following sections we delve into greater detail on the technical and conceptual insights derived from these pilots.

3.3.1 Waves 1 & 2: Interbank payments

In the first wave, pilots focused on developing competence and understanding of elementary interbank payments systems using DLT. The focus was on settlement between financial institutions of the sort that is traditionally implemented through a central bank’s balance sheet and legacy settlement systems - systems typically based on reserve accounts and RTGS systems. *Jasper₁* (2017) and *Ubin₁* (2017) were trailblazers in this respect, along with early work by the Bank of Brazil (*BCB* (2017)). These schemes first established the feasibility of replacing settlement on the books of the central bank (debiting and crediting reserve accounts) with settlement through a decentralized ledger. In terms of the technology, these ‘first wave’ pilots were generally based on blockchain systems designed for public usage - typically the Ethereum protocol.

The second wave of pilots extended the early interbank payments experiments, but took more seriously the functions that would be required for any real-world system. Key among these was the requirement for confidentiality and finality, as befits the nature of financial transactions settled over a wholesale network. As such, this wave saw increasing use of architectures designed for enterprise implementations - permissioned DLT technologies with more elaborate privacy enhancing techniques and consensus mechanisms consistent with immediate finality, as discussed in greater detail below. In order to make the systems scalable and practical for real world use, Liquidity Savings Mechanisms (LSM) were also frequently considered.

Examples of such ‘second wave’ systems were *Stella₁* (2017), *Ubin₂* (2017), *Jasper₂* (2017), *Khokha₁* (2018), and *Inthanon₁* (2019).⁵ While our focus is on domestic pilots, it should also be noted that through the establishment of their innovative *shared* ledger system, UAE and Saudi Arabia also advanced the interbank payments frontier considerably in *Aber* (2020).

Any wCBDC system, regardless of its extent, will have at its core an interbank payments system of the sort tested in these early pilots. Indeed, as later cohorts of countries have set up their first pilots, the establishment of a payments system has typically been incorporated into a broader system, rather than in a separate pilot. Any payment system introduced as part of a new pilot going forward should also incorporate up-to-date approaches towards liquidity savings mechanisms, confidentiality, and performance.

⁴See *Khokha₂* (2022) for a compact survey of the early waves of CBDC development.

⁵While not testing LSMs, *Khokha₁* (2018) featured various elements that took the pilot beyond the introductory level of the first wave pilots, such as the incorporation of sophisticated encryption techniques and improved consensus methods.

3.3.2 Wave 3: Securities settlement and ledger interoperability

Once countries had developed experience with DLT-based payment systems, the emphasis turned to securities settlement - a third wave (see *Helvetia₁* (2020) for example). Although the precise nature of the securities being traded has varied across pilots, a general theme has been the desire to achieve a ‘Delivery-versus-Payment’ (DvP) structure to transactions (see appendix B for a brief discussion of the term ‘delivery-versus-payment’).

Another key aspect of this third wave - and beyond - was interoperability. As discussed below, interoperability takes various forms. In most pilots there is some interaction between legacy and DLT systems, enabled by manual work or bespoke APIs. But in the third wave there was a particular focus on examining how multiple distributed ledgers might interact, with tokenized securities and wCBDC perhaps being on different ledgers. *Stella₂* (2018) and *Ubin₃* (2018) are particularly prominent examples of pilots exploring these issues while *Khokha₂* (2022) also featured an elaborate multi-ledger system (see appendix C for a more detailed analysis of *Khokha₂* (2022)).

While securities *settlement* was emphasized in the third wave, the pilots were not simply limited to the transfer of securities. Several pilots explored post-trade activities (coupon and principal payments, for example) and auxiliary systems such as collateralized lending and trading on margin (see *Inthanon₂* (2019), for example). Any pilot exploring securities settlement would now be expected to exhibit at least some of this additional functionality.

3.3.3 Wave 4? Exploring a DLT ecosystem

A key *recent* trend that appears to be emerging - constituting perhaps a ‘fourth wave’ - is the tendency of pilots to go beyond the involvement of standard wholesale market participants and use cases to involve a broader class of participants and applications, even to the extent of inviting proposals for use cases offered by the private sector. These pilots put more emphasis on exploring a DLT ‘ecosystem’.

Stimulating innovation in the broader financial system has always been a stated goal, even of early pilots, but we now see this manifested in tangible actions as the ‘fourth wave’ gains momentum. There has long been the sense that much of the potential for DLT, tokenization and smart contract technologies can only be unlocked at scale and through the interaction of multiple systems. It is in these complicated scenarios where DLT’s comparative advantage over legacy methods is likely greatest, where traditional back-office reconciliation and manual coordination of multiple non-DLT systems puts a limit on how elaborate the collection of interlocking systems can be. Innovations such as smart contracts may allow more complex interaction to be automated. However, these benefits are very difficult to test *ex ante* with a government-run trial or trials involving only a narrow set of banks. Hence the gradual shift towards encouraging private-sector driven use cases.

At the frontier, we are now seeing central banks and regulators inviting collaboration with the

broader fintech community as part of their pilots. For example, *Ubin₅* (2021) is explicitly aimed at promoting development of a broad class of use cases - including trade and supply finance. Following on from Project Atom, the Reserve Bank of Australia and partners are explicitly [inviting submissions](#) for proposed use cases to be built on top of a CBDC platform. The ‘eAUD’ platform - under the management and oversight of the RBA - will then operate as a centralized ledger around which can be built a variety of interoperable systems - wholesale or retail (*eAUD* (2022)).

It is perhaps difficult to make a stark division of pilots into third and fourth waves but early hints of this fourth wave include the presence of client wallets on ledger in the *BdF* (2021) trial of secondary market trading of tokenized sovereign bonds. *Khokha₂* (2022) explicitly allowed for the creation of a ‘stablecoin’ by pilot participants and, in a sub project, allowed participants to develop their own tokens and dApps (decentralized apps) to be operated on top of the base layer of systems developed within the main pilot. Another example of broadening the scope of wCBDC pilots beyond simply payments and security settlement was the syndicated lending applications explored in *Atom* (2021).

4 Key functionality in wCBDC pilots

4.1 Liquidity and the security lifecycle

A key element of wCBDC frameworks is the implementation of a large value payments system using DLT, undertaking the tasks traditionally carried out by legacy (typically RTGS) systems. This was the focus of the first two ‘waves’ of wCBDC pilots.

RTGS systems embed mechanisms to allow payments to occur smoothly despite enormous liquidity demands (Norman (2010)). These ‘liquidity savings mechanisms’ (LSMs) typically entail some sort of queuing system, allowing payments to be delayed and netted against each other in a way that enormously reduces the amount of liquidity required to fund a given sequence of payments during the day.⁶

Without such mechanisms, ‘gridlock’ may occur where payments cannot be settled *in sequence* even though the banks involved have adequate *overall* funds (Bech and Soramäki (2002)). Banks involved in transaction A may need funds from transaction B before settling A, and *vice versa*. As such, some coordination and the use of simultaneous settlement - typically with multilateral netting - may be required.

In addition to LSMs, facilities that provide intraday credit from the central bank also commonly appear in large value settlement systems (BIS (2003)). Appropriately designed intraday credit can help ease ‘deadlocks’, which are cases where not only are there gridlocks but there may *also* be an aggregate lack of funds across queued transactions, even after netting. In fact, the allowance for interbank credit may also enhance the efficiency of settlement, beyond any lending by the central bank.

These systems, naturally generalized, also can apply in the case of securities settlement and not simply in pure payments systems. Multi-asset netting may be applied in these cases. There may also be synergies with the establishment of a DLT system for tokenized securities if it supports the use of collateral in obtaining central bank money on credit.

Additional realism can also be brought to systems featuring securities, by taking seriously the need to implement standard events in securities’ lifecycle (coupon payments, dividends etc.) and several pilots have made progress in these dimensions.

4.1.1 Liquidity demands and DLT

An important question to consider is what effect the introduction of wCBDC and broader DLT-adoption will have on liquidity demand. Qualitatively, one could envisage effects in both directions.

⁶As stated in Norman (2010), ‘[i]n many cases, these payments still settle legally gross, but the economic effect is to net the liquidity requirement’.

Proponents of wCBDC and broader adoption of tokenization emphasize the possibility of *immediate* settlement - perhaps bringing to fruition the recent trend in shortening settlements to T+1 or even, in some form, T+0 (see DTCC (2021)). Separately, *atomic* settlement is often cited as one of the main promised benefits of DLT, whereby smart contracts ensure that a transaction either succeeds or fails in its entirety, with no partial execution.

All else equal, shortening settlement and ensuring atomicity are desirable in that they reduce various risks or costs of protecting against them through margin/collateral requirements. This may free up assets for alternative uses in other transactions. Also, immediate atomic settlement would mean that assets are not locked up during a ‘settlement window’ with perhaps an unknown end date. The improved efficiency in transactions that *do* settle successfully may reduce liquidity demands for a given volume of transactions.⁷

However, any attempt to move to an immediate settlement model may confront the issue that, for some time, sections of the legacy financial system will put constraints on how quickly assets can be delivered for settlement. Demanding immediate settlement may then restrict the set of transactions that can actually occur, or necessitate a greater degree of pre-funding and asset inventory accumulation (Lee et al. (2022)). If DLT-based systems ultimately induce higher volumes and greater complexity of transactions, there may again be extra demands for liquidity. Furthermore, mechanisms that allow assets on two different distributed ledgers to be exchanged may require time delays *as a feature*, limiting the scope to release and re-use liquidity rapidly, even if the transaction is atomic.⁸

4.1.2 Liquidity mechanisms in existing pilots

Various pilots have explored LSMs and, in particular, gridlock resolution. The mechanisms are traditionally quite centralized and inherently multilateral - involving information about many parties’ transactions being communicated to the central bank. Decentralizing this system (reducing or eliminating the role of the central bank) created challenges and trade-offs, particularly in relation to preventing information about parties’ transactions being inappropriately shared.

Jasper₂ (2017) allowed banks to submit payments either for immediate settlement or to a centralized queue, from which an algorithm would determine sets of payments to be matched and netted. The implementation followed the basic approach of a legacy systems used in the UK’s CHAPS RTGS system and retained the central bank as a coordinating agent. *Stella₁* (2017) also examined legacy gridlock resolution mechanisms as found in the the BoJ and ECB’s RTGS systems, focusing primarily on performance.

⁷As ledger interoperability improves, it is also possible that a centralized pool of collateral may be more easily deployable across multiple markets (Bundesbank (2020)).

⁸See *Ubin₃* (2018) (pp. 42) for concerns over the liquidity implications of Hashed Timelock Contracts, which we discuss in more detail below when we discuss ledger inter-operability. While primarily focused on cross-border payments, *Stella₃* (2019) provides a ranking of different cross-ledger approaches in terms of their liquidity implications.

In fact, there appears to have been relatively little innovation in the fundamental structure of the netting algorithms used. Instead pilots have typically re-implemented existing algorithms in the new DLT context. An interesting question is whether new algorithms can be developed that are especially suited to DLT systems and the desire for privacy-respecting decentralization. One example of such an innovation appears to be the ‘Cycle-solver’ algorithm used in the Corda stream of *Ubin₂* (2017), described in detail in *Furgal et al.* (2018).⁹

Ubin₂ (2017) was an important pilot in showing that gridlock resolution could be decentralized to a large degree, while preserving privacy. However, this success was dependent on what DLT architecture was used. All the platforms considered (Corda, Hyperledger Fabric and Quorum) were able to provide an operational mechanism (see appendix D for some discussion of the details of these platforms). However, Hyperledger Fabric’s implementation could not function if the central bank was offline during the process.¹⁰ Also, although the report does not emphasize them, there appeared to be some information leaks associated with the gridlock resolution process under the Corda and Hyperledger Fabric implementations. Quorum - using zero knowledge proofs - discussed further below - apparently provided complete decentralization and confidentiality, though at some cost in terms of speed.

Of course, it may be acceptable to have a centralized system for gridlock resolution. *Inthanon₁* (2019) was based on the Corda platform and concerns over privacy raised in *Ubin₂* (2017) were solved by establishing an ‘oracle node’ run by the central bank. This node played an information-gathering and planning role in guiding the resolution mechanism. In fact, centralization under the authority of the central bank may also help address another concern alluded to in *Ubin₂* (2017) - that allowing the gridlock resolution cycle to be initiated by banks could lead to undesirable strategic behavior and manipulation of the timing of the netting cycle.

Various features of *Inthanon₁* (2019) are relevant to issues of liquidity. The pilot considered queuing mechanisms, multiple gridlock resolution algorithms and automated liquidity provisioning *via* repo between participants and the central bank, using an on-ledger tokenized bond as collateral. Indeed, in *Inthanon₂* (2019) the establishment of *interbank* repo further expanded the liquidity management capabilities of the system, as did an extension of the LSM to allow for multi-asset netting. Also relevant are the insights from *BdF* (2021), which (among other things) explored a cross-platform repo framework with the pledging of collateral occurring in the traditional **T2S** systems, and the ‘lending’ of wCBDC occurring on the DLT platform. The pilot also explored auto-collateralization and (though the documentation is not entirely clear) inter-bank repo.

⁹*Ubin₂* (2017) also expressed a need for innovative implementations of LSMs, perhaps involving distributed the workload across clusters of nodes in the context of very large networks. It is unclear whether such innovations have yet occurred.

¹⁰More recently, *Aber* (2020) came to a similar conclusion over the difficulty of full decentralization of gridlock resolution in its Hyperledger Fabric implementation.

4.1.3 Security lifecycle

Setting aside the outright exchange of securities, DLT systems have been used to administer key events in securities' lifecycles. Some pilots may still exclude such processes from their scope, in order to focus on other elements. However, they appear necessary for any system that approaches the level of production readiness, given their central place in financial market infrastructure and activity.

Two pilots that have arguably made the most progress in establishing such systems are [BdF \(2021\)](#) and [Inthanon₂ \(2019\)](#). Smart contracts were used to automate and streamline actions such as coupon and principal payments in wCBDC, with obvious scope to generalize to dividend payments and other disbursements in different contexts.¹¹ Such disbursements can be tied to almost arbitrary conditions, allowing enormous flexibility to automate business logic. Such implementations may also reduce the risk of non-compliance with agreed obligations, in turn reducing the need for resources to be devoted by investors to keeping track of disbursements ([Bech et al. \(2020\)](#)).

The sophistication of [Inthanon₂ \(2019\)](#) also saw the automation of margin calls within the repo framework and the correct handling of coupon payments while a bond is held by the lender as collateral. It is emphasized in [Inthanon₂ \(2019\)](#) that the margin system's efficiency is enhanced by being executed on ledger because of the constant availability of a consistent and mutually accepted record of positions.

¹¹[Khokha₂ \(2022\)](#) also has illustrated how principal and interest payments may be made in wCBDC, administered *via* smart contract.

4.2 DvP settlement and ledger interoperability

Pilots exploring securities settlement have differed in various respects, but one key aspect is whether they adopted what we will term an ‘on-ledger’ or ‘interoperable-ledger’ format. By ‘on-ledger’ we mean that a wCBDC is separately issued on each ledger associated with a different tokenized security. While pilots would typically only use ‘one ledger’ in this case, the extension of this format to a multiple ledger case world would imply many, siloed ledgers with both wCBDC and securities issued on each. In the ‘interoperable ledger’ approach, the central bank would administer a single ledger on which wCBDC would be issued, and this ledger would then interface with the multiple securities ledgers, without wCBDC strictly being issued on them.

The distinction between these two models is apparently referred to in *Helvetia₂* (2022). The pilot had focused on an ‘on-ledger’ solution in its DLT experiments (in fact only one ledger was used). However, in its conclusions, the pilot considered the possibility of an ‘interoperable ledger’ format (our emphasis added):

wCBDC may be issued onto a single DLT-based payment platform - potentially operated and owned by the central bank - that is interoperable with one or potentially multiple tokenised asset platforms. Compared with the wCBDC issuance on multiple platforms, this...

4.2.1 On-ledger vs inter-operable ledgers

DvP settlement is especially simple to achieve where both the security and the wCBDC are recorded on the same ledger. In this case many complexities of implementation can be avoided and smart contracts can enable atomic settlement. Indeed it may be feasible to make the settlement instantaneous and not just atomic.¹² Only a single transaction is required and no additional interfaces need be designed, beyond the native protocols of the ledger on which the trades occur. Notable also is the fact that the securities and digital cash need never leave their initial owners’ wallets until the instant that the trade is completely settled, without any third parties or auxiliary protocols/systems being involved.

Functions such as liquidity savings mechanisms, netting of positions and collateral management, may benefit from being jointly managed across multiple asset types in a unified system with simple smart contracts (*Stella₂* (2018)). Overall, liquidity may be enhanced in transactions involving assets kept together on a single ledger.

The on-ledger approach was adopted in *Jasper₃* (2018) where the tokenized asset was an equity claim. The pilot achieved in a single step what traditionally had involved distinct settlement processes on a Centralized Securities Depository (CDSX) and in the Canadian large value payments

¹²Note that instantaneous settlement may not be desirable for various reasons (see *Lee et al.* (2022) and references therein). However, all else equal, it is natural to think that faster settlement would be preferred, as illustrated by recent reforms to accelerate settlement in many legacy systems.

system, LVTS. Another pilot adopting this approach was *Inthanon₂* (2019), which explored trading and repo of government bonds using a single ledger. The Bundesbank’s *Blockbuster* (2018) also allowed for settlement of digital coins and digital assets within the same ledger. In fact, it also explicitly allowed for free of payment (bonds vs bonds) and only-payment (coins vs coins) modes of exchange. More recently, *BdF* (2021) has apparently explored on-ledger settlement of primary and secondary market trade scenarios in tokenized sovereign bonds and wCBDC.

Now, ideal DLT specifications (such as node structure, consensus mechanism, privacy enhancement) will vary by asset, even among assets that are commonly traded against each other - necessitating (many) separate blockchains. Indeed, the proliferation of DLT platforms in recent years has been dramatic. It is perhaps implausible to think that a central bank would be willing and able to issue wCBDC directly into many diverse ledgers.

More plausibly, central banks may prefer to focus on a single wCBDC ledger - or use an ‘on-ledger’ approach only for a small set of key markets. They may then follow a *complementary* policy of encouraging (standardized) interfaces to emerge to support a broader ecosystem of interoperable ledgers, delegating implementation to other parties who may have a comparative advantage.¹³ This arguably is the path that the ‘fourth wave’ of domestic wCBDC pilots are taking.

Since ledger interoperability is a key area of research in the DLT community - far beyond CBDC applications - it is likely that the methods to allow interoperability and cross-chain business logic will improve over time. We now describe some of the most prominent that have been used in wCBDC pilots.

4.2.2 Methods for ledger interoperability

Interoperability of blockchains can be achieved in various ways.¹⁴ Among domestic wCBDC pilots *Hashed Timelock Contracts* (HTLCs) have featured prominently. *Stella₂* (2018) and *Ubin₃* (2018) used HTLC to allow DvP exchange of securities for wCBDC, while *Atom* (2021) used the approach as part of its syndicated loan application.¹⁵

A HTLC relies upon two key functions in enabling cross-ledger transfers in the absence of an explicit institutional connection between the ledgers or a trusted intermediary. These functions are a ‘hashlock’ and a ‘timelock’. The hashlock enforces the requirement that only the appropriate parties have access to the wCBDC and securities at particular stages in the exchange, using cryptographic techniques. The timelock limits how long it is permitted for the trade to remain

¹³*BIS* (2021) alludes to issues of comparative advantage and the danger of central banks becoming over-involved in the development of systems that might be better developed in the private sector.

¹⁴See *World Bank* (2021) for a thorough survey of techniques not limited to CBDC applications. Also, see [here](#) for a succinct and accessible summary.

¹⁵HTLCs also feature in cross-border applications, such as *Jasper-Ubin* (2019), where the need for ledger interoperability commonly arises, reflecting different jurisdictions having their own DLT systems and yet wishing to interact efficiently.

unsettled before assets are returned to their initial owners.¹⁶ Simplifying somewhat and ignoring implementation details, the process by which these interact is as follows (taken from [Bech et al. \(2020\)](#)):

1. The seller generates a secret (X) and its corresponding hash ($Y = f(X)$). The seller uses this hash to lock the security tokens on its ledger with a specified time limit (eg four hours). The seller creates an instruction to either send the securities to the buyer using the hash (Y) or, if the time limit expires, return them to the seller.
2. The buyer locks the cash token on its ledger with a shorter time limit (e.g. two hours). The buyer creates a conditional instruction to either send the cash token to the seller using the hash (Y) or return it to the buyer after two hours.
3. The seller reveals the secret (X) to unlock and retrieve the cash token (key here is that only the seller is able to reveal the secret - only the seller knows the X that hashed to Y).
4. The buyer can then use the secret (X) to unlock and retrieve the security tokens.

Implicitly, some communication system must be available so that parties may convey the various messages (containing hashes etc.). How exactly such a communications network is established has received relatively little attention in pilots. However, it seems reasonable to wonder if there may be a coordinating role for regulators or industry consortia to avoid duplication and to establish broadly accepted messaging standards that could apply to transactions across a wide variety of disjoint ledgers.

Multiple distributed ledgers were used in [Khokha₂ \(2022\)](#), a project with a very rich structure. The pilot features interoperable ledgers, implemented without HTLCs. We here discuss elements of the its approach and refer the reader to appendix [C](#) for more detail on this pilot.

First, they exploit some of the machinery developed in the [Cosmos](#) ecosystem. Cosmos provides a ‘notary’ scheme in which a set of trusted nodes manages interoperability of ledgers ([Koens and Poll \(2019\)](#) and [World Bank \(2021\)](#)). They intermediate between the ledgers on the basis of logic captured in decentralized applications, or in Cosmos terminology, modules. The structure adopted is one of ‘hubs and zones’, building upon the Inter-blockchain Communication Protocol ([IBC](#)), with hubs connecting multiple zones. In the case of [Khokha₂ \(2022\)](#) zones were associated with DLT platforms operated by commercial banks, plus an additional DLT platform (the ‘CBDC zone’) operated by the central bank, on which wCBDC was issued. The hub that connected these zones was referred to as the ‘Khokha hub’.

The second way in which the pilot demonstrated ledger interoperability was through the use of software bridges. A collection of APIs was employed to enable the ‘porting’ of wCBDC from the

¹⁶It also enforces asymmetric time limits for different parties to execute their designated actions, preventing nefarious reversion of the first leg, with the intent to collect the proceeds of the second. See [Stella₂ \(2018\)](#) pp. 6, footnote 16, and pp. 18 for details.

CBDC zone, to the Khokha hub, where it could then be used as a settlement asset among banks. The ported wCBDC could also be redeemed against a private stablecoin (wToken) issued by banks in the Khokha hub. The bridge was operated off-chain by the central bank, in contrast to how the [Cosmos](#) processes were handling other aspects of the ledgers' interaction.

Setting aside the technical details of [Khokha₂ \(2022\)](#), a key insight is that the system resembles in some respects the framework alluded to at the end of the previous section. The central bank issues wCBDC natively on a given ledger and then permits connections to be made to this ledger in order to allow a safe settlement asset to be used more widely, in a broader DLT ecosystem. This provides scope for the ecosystem to feature devolved responsibilities - such as the issuance of stablecoins by other parties. Of course, there should not be *complete* devolution of authority. A central bank with monetary and financial stability mandates presumably will always have a say in acceptable standards and guidelines to be followed by those running other ledgers and issuing various tokens. Indeed, Khokha hub governance entailed a Khokha Council with the central bank a member of that council. How such governing bodies should be implemented in any production-ready system is an important question.

Beyond HTLCs, notary systems and software bridges, there are several other approaches to enabling interoperability, though it can be quite difficult to identify if they have been used in domestic wCBDC pilots, based on the public documentation. Again, we refer the reader to [World Bank \(2021\)](#) for a broader discussion. A more formal treatment can be found in [Koens and Poll \(2019\)](#).

4.2.3 Interoperability with legacy systems

Before leaving the theme of interoperability, we note that it is not simply blockchains that need to be stitched together. Interoperability between DLT and legacy (non-DLT) systems is also important, and likely to be so for some time ([WEF \(2020\)](#)). Indeed, it could be that there will always remain some systems not suited to DLT but with which DLT systems will need to interact.

To some extent this form of interoperability features in essentially all the domestic CBDC pilots discussed so far. Almost without exception, the pilots involve existing reserve account systems in some way. The interfaces vary from manual work, to sophisticated APIs but they are there nevertheless.

Some trials have explicitly addressed the issue. For example, in [Helvetia₁ \(2020\)](#), two proofs of concept were attempted. One, where tokenized assets and wCBDC were both on the same ledger, and another where the ledger for tokenized securities was connected to the existing RTGS and reserve balance system (SIC). The latter approach thus featured one 'traditional ledger' for central bank money, and one DLT-based system for securities. Appropriate interfaces were constructed so that settlement instructions in terms of tokenized assets resulted in the 'blocking' of the assets on the ledger, and also the sending of a message to the SIC system, to trigger the associated cash payment. On completion of that payment, a further instruction was sent back to the ledger, to

trigger the unblocking and transfer of the tokenized securities.

This model is somewhat akin to that adopted in [recent experiments](#) carried out by the Bundesbank and Deutsche Börse, in which primary and secondary market transactions in a tokenized government bond were settled, with the *cash* leg using the existing [TARGET2](#) large value payment system. A system of ‘triggers’, implemented through various APIs enabled this interaction. In discussing this pilot it was stated:

Following successful testing, the Eurosystem should be able to implement such a solution in a relatively short space of time - at least in far less time than it would take to issue central bank digital currency, for instance.

[Burkhard Balz](#), Bundesbank Executive Board, March 2021

This quote emphasizes that the timeliness of a rollout may also influence where a central bank devotes its resources. If there is a concern that central banks may be ‘left behind’ legacy-to-DLT interoperability may be an important element of any pilot, to guard against that.

More positively, given the significant [enhancements](#) already implemented and planned in [TARGET2](#) and associated systems it is natural to question whether the *incumbent* systems are able to do what a wCBDC is capable of, at least in the near term (see [Panetta \(2022\)](#) for recent comments in this regard). For countries without access to such efficient *incumbent* technology, this is perhaps less of an argument against wCBDC. But the general need for legacy-to-DLT interoperability still will apply in most jurisdictions, however.¹⁷

Legacy-to-DLT connections can be implemented in various ways, such as with a trusted third party and using existing messaging technologies and interfaces - details we do not discuss here. However, we note that ‘[oracles](#)’ may also play an important role in connecting legacy and DLT systems. While usually discussed in the context of bringing *general* information onto a blockchain, necessary for smart contract execution, they are emphasized in [WEF \(2020\)](#) and [World Bank \(2021\)](#) for the *specific* case of connecting legacy business processes with the new DLT infrastructure.

¹⁷It is also possible that DLT and legacy systems may ultimately be used in parallel, or as mutual back-ups in the medium to long run. In this case might need to interact in such a way that end users are unaware of which implementation is at play in the background at any given time. The Bank of Brazil explicitly trialed a wCBDC system with a view to establishing a backup - what they referred to as the Alternative System for Transactions Settlement or ‘SALT’ ([BCB \(2017\)](#)).

5 Performance, privacy and resilience

5.1 Performance

Most pilots have not delved deeply into measures of performance, beyond identifying key qualitative issues and broad insights. This reflects, first, that there have been other more pressing matters to be considered, such as simply building understanding or establishing confidentiality. Second, several of the platforms used in the trials have been in a state of rapid development, making it difficult to apply performance comparisons beyond the very short term. Finally, the ability to assess performance is constrained by the limited scope of the pilots. Many of the benefits of wCBDC may emerge only in the context of broader DLT adoption in financial markets. However, a belief that wCBDC will be introduced may be necessary for some of this innovation to take place - leading to a ‘chicken and egg’ situation.

Nevertheless, pilots have generated many useful insights, which should be considered and built upon in future trials.

5.1.1 Test frameworks

‘Throughput’ and ‘latency’ are two key metrics commonly adopted in assessing DLT performance. ‘Throughput’ is typically calculated in terms of numbers of transactions per second (TPS) for a given network size. ‘Latency’ is typically calculated using a metric such as the time between a transaction request being sent and when it was actually written to a block by all nodes, or by whatever subset of nodes required by the consensus protocol used.¹⁸

The data used in pilots is commonly inspired by (or even taken from) analogous legacy systems. That is, the data will typically be historical or simulated to capture transaction volumes in ‘normal’ times, times of ‘stress’, and in a hypothetical ‘future’ environment that allows for trend growth in demands on the system. In an interesting recent development in the U.S., the DTCC’s [Project Ion](#) currently runs ‘in parallel’ with their legacy systems so that it receives transaction requests in real time ([DTCC \(2021\)](#)).

Overall throughput and latency are useful, but coarse, measures of performance. Pilots should also distinguish the contribution of different steps and sub-processes to overall performance ([Stella₁ \(2017\)](#) is a good demonstration of this approach). For example, distinguishing the time taken to execute a smart contract, validate a transaction, or to commit a block provides deeper insight into the performance of the framework. Of course, what exactly is reported will depend on the pilot but a careful profiling of the performance of sub-systems will become increasingly important as the pilots approach production-readiness. Identifying bottlenecks at an early stage will allow

¹⁸For an interesting - and largely platform-independent - discussion of how to assess the performance of a blockchain in more detail, see [Hyperledger \(2018\)](#).

resources to be focused on eliminating particular pain points.

If transaction-level performance is discussed, it should not simply be the average (across transactions) that is reported. Then, even if two systems are similar on average, one can distinguish which system has more ‘extreme’ events. *Khokha₁ (2018)* documentation includes the entire distribution of block propagation times, for example. While it is less commonly reported among existing pilots, it is also useful to record the resource demands implied by the system. In *Blockbaster (2018)*, for example, CPU utilization, memory and storage usage, and network traffic were all considered.

Of course, if such technical measures of performance are being reported, it is important to be clear about the underlying specification of the test environment. Otherwise, it may be difficult to assess whether one pilot’s findings are due to, say, the choice of consensus protocol, as opposed to simply the processing power of the computers used. Existing pilots do report some information regarding their hardware and environment, but it is typically something of an afterthought and difficult to compare across systems in a way that allows insight into what particular parameters of the setup might be most influential for performance. Ideally the networking setup should also be benchmarked with non-CBDC operations (general operations unrelated to DLT) as the network architectures often vary dramatically across pilots.

Pilots should also consider how frameworks perform under ‘exceptions’ - cases where there are errors and failures. Trials have considered some limited cases of these - such as submitting intentionally incorrect data (*Stella₁ (2017)*), allowing particular nodes to fail (*Stella₁ (2017)*, *Ubin₂ (2017)*), or experimenting with transactions requested by parties with insufficient balances (*Helvetia₁ (2020)*). As noted in *Hyperledger (2018)*, when calculating performance metrics it is important to adjust for the presence of invalid transactions. More generally, performance assessment in trials should not simply be based on ‘happy path’ execution but should take into account plausible rates and types of exceptions.

5.1.2 Findings from existing pilots

Performance evaluation was quite prominent in early pilots - those that primarily dealt with payments systems. Since then, priorities have been such that performance assessment has become somewhat secondary to building understanding and proving feasibility.

Early pilots, such as *Jasper₁ (2017)* and *Ubin₁ (2017)*, established that while Ethereum-based protocols were capable of comparable throughput to legacy LVTS systems in ‘normal’ times, their ability to scale with surges or in anticipation of trend growth in transaction demands was inadequate. To a large degree, the limitation on scalability derived from the *Proof-of-Work* consensus mechanism used. With the move to permissioned systems, consensus mechanisms were adopted that allowed larger volumes of transactions to be handled at acceptable levels of performance. Consensus was typically based on some form of voting model, requiring designated (possibly singleton) subsets of nodes to validate and commit transactions (see appendix A for more on

consensus mechanisms).

In a stretch ‘sub-project’ of *Ubin*₁ (2017), a Quorum implementation was deployed, with a voting protocol that resulted in much higher throughput. Also using Quorum, *Khokha*₁ (2018) was able to match and even surpass legacy systems’ performance in terms of throughput and latency. This was even under their most elaborate version of the system, with maximal decentralization and privacy enhancement. *Jasper*₂ (2017), employed an early version of the Corda platform and was able to accelerate processing dramatically, owing in large part to its notary-based consensus mechanism. *Stella*₁ (2017) also explored performance of a Hyperledger Fabric system featuring LSMS and found promising results. More recently, *Aber* (2020) confirmed (again, with Hyperledger Fabric) the ability of DLT-based systems to offer high levels of performance while implementing LSMS, even subject to additional requirements such as maximal decentralization and protection of privacy.

It should be noted, however, that some of the fastest consensus mechanisms may expose the system to crashes or malfeasance among nodes - there is a robustness-speed trade-off. Clearly, righting a system after failures may imply a hidden performance cost if not accounted for carefully in trials, which is why pilots sometimes report ‘time taken to restart a node’, and similar metrics.

If the use case is such that speed is not the primary concern and trust among participants is not assured, one might wish to sacrifice raw ‘performance’ for greater safety. For example, *Atom* (2021) emphasizes that the speeds required for a syndicated lending system may be far lower than those required for, say, trading shares. While not strictly a wCBDC pilot, the DTCC’s recently launched *Project Ion* implements a DLT system for bilateral trading of shares (DTCC (2021)). *Ion* is currently built on a Corda platform, and it is plausible that Corda’s default ‘single party notary’ solution is a good match for a setup where there is an obvious candidate for a trusted and reliable party to run the notary node(s) - the DTCC.

Not all of the improvements in performance and scalability relative to the early Ethereum-based trials were purely due to changing the consensus mechanisms. Improvements are also due to platforms allowing different steps in computation (execution of smart contracts, validation, consensus,...) to be operated concurrently by subsets of participants. Storing and operating on some data ‘off-chain’ also helps reduce communication and computation loads for the network. This approach may also enhance privacy, as discussed in section 5.2. We suspect that familiarity and expertise with smart contracts has also grown though these ‘algorithmic’ improvements are not explicitly mentioned.¹⁹

The aforementioned findings relate primarily to the waves of pilots dealing with payments. Among later securities-settlement pilots it is common for performance not to be addressed in great detail. Frequently, performance analysis is explicitly excluded from scope (see *Jasper*₃ (2018), *Inthanon*₂

¹⁹*Blockbaster* (2018) refers to complicated chaincode slowing down transaction processing and parallel processing seems still appear to be a challenge for certain classes of blockchains. *Blockbaster* (2018) also, in fact, refers to research by Hyperledger into enhanced concurrency in Fabric.

(2019), BdF (2021) and *Helvetia₂* (2022)). Some pilots - such as *Blockbuster* (2018) and *Stella₂* (2018) - do refer to throughout/latency targets or provide some formal profiling of bottlenecks, but typically qualitative insights are offered.²⁰

One important performance-related insight is that DvP with DLT can eliminate intermediate steps and compress the settlement cycle, possibly eliminating the role for certain classes of intermediary and speeding up the system dramatically. *Inthanon₂* (2019) and *Ubin₃* (2018) both note the streamlining of intermediate steps in single-ledger and cross-ledger DvP solutions, respectively.

If settlement occurs on a single ledger then trading and settlement can in many cases likely be merged, with an immediate atomic structure. If multiple ledgers are involved then, as noted in *Stella₂* (2018), there can be significant latency in some inter-ledger solutions, due to the locking of assets on one or more ledgers while settlement completes. However, the question then is whether that delay is meaningful, set against the possibly multiple intermediate steps that legacy systems would have to go through to execute the same exchange.

This perhaps points to why performance was not evaluated so mechanically after the first waves of payments pilots - the benchmarks to compare against for securities settlement and broader applications are much more complicated, likely involving multiple legacy systems. Additionally, there is clearly greater complexity to DvP pilots - especially those with multiple ledgers - so the resources of the pilot are possibly better spent in understanding the essential mechanisms rather than benchmarking. At some point, however, more formal performance assessment will have to be done, especially as the updating of legacy systems has set a high performance bar in some jurisdictions (*Panetta* (2022)).

²⁰We are focusing on domestic wCBDC pilots, but nevertheless note that *Stella₃* (2019) offers a qualitative ranking of different ledger-interoperability solutions, on the basis of their implications for liquidity efficiency. Other cross-border / multi CBDC pilots may limit their analysis because the orders of magnitude of improvement in transaction times over legacy systems renders a precise comparison moot. For example, the Inthanon-Lionrock to mBridge pilot suggested transactions may take between 2 to 10 seconds, rather than 3-5 days. Of course, the question is whether the elimination of legal and compatibility barriers that account for much of the delay in legacy systems can truly be avoided in a real world DLT implementation.

5.2 Privacy

Privacy is a key concern in the design of wholesale systems, owing to the value of payments executed, to the business disadvantages that could emerge from disclosure of information to unauthorized parties, and to client data protection requirements imposed upon many financial intermediaries. Early trials used the standard Ethereum protocols which rendered information transparent to all nodes on the network (see *Ubin*₁ (2017) for example). While these trials established useful experience with DLT designed for public or ‘permissionless’ systems, from a confidentiality perspective they were clearly inadequate.

Establishing a permissioned system is a first step towards enhancing privacy (see appendix A for further discussion of permissioned and permissionless systems). However, even *within* a permissioned system there will typically be information that should not be shared among all participants. For example, within a network of banks, the details of bilateral transactions should, in the first instance, only be known to the involved parties. These details might be identities of the counterparties, values of exchanged funds and securities, counterparty balances, smart contract logic involved in the transaction or many other sensitive details. There are various interrelated approaches that can be used to enhance confidentiality beyond simply permissioning and we here discuss some of the main techniques used so far in the wCBDC domain. Note that *Stella*₄ (2020) provides a thorough analysis of many issues regarding confidentiality in DLT and a broader survey of privacy-enhancing techniques than we provide here.

5.2.1 Sub-networks of participants

A basic method of enhancing confidentiality is to establish ‘sub networks’ within a broader network of participants. Different platforms refer to, and implement, such methods in different ways. For example, Corda adopts a ‘need to know’ approach to peer-to-peer communication. Transactions are, in the first instance, accessible and known only to the involved counterparties - a particularly stark example of a subnetwork. In comparison, Hyperledger Fabric exploits a ‘channels’ approach to establish a more persistent concept of a sub-network that doesn’t simply exist for a given transaction. Within a given channel, participants can conduct multiple private and confidential transactions and a private blockchain is associated with that particular channel.

Quorum arguably does not use sub-networks, or at least uses them in a less extreme form. Rather than limiting communication regarding a transaction to involve only counterparties, *what* is communicated differs across parties. ‘Private transactions’ refer to a particular set of participants directly involved in the transaction, who receive encrypted information (via private transaction managers) that only they can decrypt and interpret. Other participants are aware of the transactions but only in a form that contains hashed versions of the encrypted information.²¹

²¹See appendix D for more detail on the various platforms.

5.2.2 Disguising identities

One problem with the Corda approach, in its early forms, was the necessity of validators to ‘walk the chain’ when assessing the legitimacy of a transaction (see [Appendix D](#) for additional detail). Given the [UTXO](#) nature of Corda’s implementation, this implied that details of historical transactions - in particular the identities of counterparties - could be observed by agents who were not originally among those counterparties. To deal with the revelation of transaction lineage, Corda now uses ‘confidential identities’ (see [Jasper₃ \(2018\)](#) and [Ubin₂ \(2017\)](#) for example), which associate unique public keys and certificates with each transaction. Identities are transparent to those involved in a given transaction. However, they are shielded from participants in *future* transactions that are connected to the earlier transaction through the [UTXO](#) chain.

It should be noted (as in [Ubin₂ \(2017\)](#)) that using this approach still may not completely secure the parties’ identity as transaction quantities remain unshielded. As such, information from chains of transactions might expose counterparties to *inference* of identity through machine learning techniques and cross referencing of, say transaction times and amounts, with other information. Nevertheless, additional techniques can complement confidential identities, to render the transactions even more opaque to all but the involved parties.

In Hyperledger Fabric there is a similar option of using an ‘[Identity Mixer](#)’ which enables anonymity and unlinkability, while demonstrating that the sender of the transaction has the authority to be undertaking that transaction. The concept of unlinkability means that multiple transactions from the same party cannot be identified as coming from the same party, whoever it is.²²

5.2.3 Off chain information

Not all information need be encoded within the transactions stored in the ledger. As aforementioned, Quorum is designed to keep confidential information off chain, with only a hashed representation of encrypted transaction details made available to network participants not directly involved in a given transaction. Similarly, even *within* a Hyperledger Fabric ‘channel’, often only a subset of participants should be privy to certain information. As such, HLF offers ‘[private data collection](#)’ functionality, where, again, elements of the transaction are only kept on chain in hashed form for others to view, but not interpret. Corda participants’ private ledgers are designed to contain only information related to each participant’s own transactions.²³

²²Identity Mixer was apparently employed in [Aber \(2020\)](#).

²³Recall though in Corda that in the simplest implementation, information from other transactions not directly involving a given participant may be revealed during the ‘walk the chain’ step of validation.

5.2.4 Zero knowledge proofs and other methods

One of the more advanced approaches used to enhance privacy is the use of [zero knowledge proofs](#) (ZKPs). These proofs allow a party to prove that they know a fact, without disclosing the fact itself. Essentially, the proofs take the place of the underlying sensitive data and can be operated upon and shared within the DLT system as required. In the context of digital currencies, this may relate to proving a valid identity. Alternatively, they may convey information about funds transferred in a transaction and/or that the party has sufficient funds to effect the transaction.

ZKPs have been incorporated for some time in Quorum implementations used in various wCBDC pilots, shielding information about counterparties, amounts and other transaction details. Indeed, as noted in [Ubin₂ \(2017\)](#) the approach taken could allow full decentralization of gridlock resolution, without some of the information leakage problems that were identified for other platforms used in the trial.

Since ZKPs allow validation without revelation it is perhaps natural that Quorum, which features a global ledger stored over all nodes, should have made heavy use of them from an early stage. More recently it appears as if they have been incorporated, perhaps in a somewhat limited way, by other platforms. For example, ZKPs feature in the aforementioned [Identity Mixer](#) system used by Hyperledger Fabric. It is somewhat unclear from Hyperledger Fabric documentation, however, how extensively ZKPs feature in their platform. References have been made to ZKPs in the context of ‘Zero Knowledge Asset Transfers’ for example (see [here](#)). However, in the recently published [Zand et al. \(2021\)](#) it is stated only that ‘*within Hyperledger, work is being performed on zero-knowledge proofs (ZKPs)*’, suggesting that they may still be under research and evaluation.

Corda also appears to regard ZKPs as an important avenue for research. In this recent [blog post](#) regarding privacy problems in CBDCs, ZKPs are mentioned as one of several ‘workarounds’ to deal with wrinkles arising from Corda’s baseline framework (the ‘walking the chain’ problem). Interestingly it is also noted that *hardware* solutions could be used as alternatives to ZKPs. Such hardware-based solutions are embedded in R3’s [Conclave](#) product and rely on ‘[trusted execution environments](#)’ in computers’ processors such as those enabled by [SGX technology](#) offered by Intel.²⁴

As discussed [here](#), it is possible to send encrypted information to a counterparty in a way that the *counterparty* will be unable to decrypt the information, even if their hardware *has* been able to decrypt it (to establish identity or balances transferred, say). This is possible because the decryption would only - and verifiably - be done on the chips in the counterparty’s computers that were running ‘*in a protected mode, that locks out the operator of the computer*’. Thus, the underlying information would not be revealed to the counterparty. As such, the hardware-based solution achieves essentially the same end as the ZKP.

Research into ZKPs is extremely active. The method is elegant, powerful and appears to allow

²⁴[Ubin₃ \(2018\)](#) in fact tested a DLT framework that made use of the SGX technology, in one of the systems trialed (the Anquan-Quorum case).

other aspects of DLT architecture to be simplified (the complexity of Hyperledger Fabric’s channels is commented on in more than one pilot’s documentation, for example). On the other hand, various pilots have noted that ZKPs can be costly in terms of performance - requiring time to construct and verify the proofs, increasing transaction latency.²⁵ Alternative methods, such as the hardware solutions just discussed, may also be substitutes. As such, ZKPs are in a state of flux and pilots should likely consider their costs and benefits carefully, going forward.

5.2.5 Central bank knowledge

The question often arises of what information needs to be, or should be, known by the central bank or some other key administrative participant. Answering this requires a decision that goes beyond the narrow technology of privacy enhancement as attitudes to privacy and oversight likely vary substantially across countries.

Central banks may be given administrative roles (such as running notaries or orderers in Corda or Hyperledger Fabric, respectively) though this is consistent with varying degrees of information visibility, according to how the privacy enhancing methods discussed above are used. It may be the case that the central bank has full visibility of everything in the system. However, the degree of transparency can be reduced.

Legacy *payments* systems featuring settlement via the central bank balance sheet provide a benchmark of complete transparency. It could be that this is desired for DLT-based alternatives (see [Aber \(2020\)](#) for example). Where securities exchange occurs, it is perhaps less obvious that full transparency is desirable, even if wCBDC is used as the settlement asset. Indeed, there is also a question of what transaction details should be seen by (possibly private) operators of DLT-based exchanges. In [Helvetia₁ \(2020\)](#) it is emphasized that the notary node run by the digital exchange, SDX, would not see or validate the business content of transactions. This is much like how in [Inthanon₁ \(2019\)](#) the central bank’s notary node only operated in a ‘non-validating’ mode - not knowing the contents of the transactions even if it could check for problems such as double spends. Thus, some privacy is clearly feasible, even when there is significant centralization of governance.

Finally, we note that some pilots have featured *partial* decryption of information on a transaction for KYC/AML oversight by other regulatory bodies ([Khokha₁ \(2018\)](#)). That is, there could be intermediate degrees of transparency for monitoring and audit, based on splitting up information within a transaction.

²⁵However, [Khokha₁ \(2018\)](#) used ZKPs and still equaled or surpassed legacy systems’ performance.

5.3 Resilience

Most pilots tend to associate the resilience of a wCBDC system with its degree of decentralization. Of course there are other dimensions that might influence resilience - the ability to restart failed nodes quickly, the robustness of the consensus mechanism to errors or malfeasance, the quality of hardware and various other factors. But decentralization and the avoidance of a single point of failure is especially important and we will focus on it here.

Decentralization is a key feature of some of the most prominent blockchains. In almost every respect, Bitcoin is maximally decentralized: all nodes possess a (possibly [pruned](#)) copy of the ledger, any node can undertake any activity (validation or mining) and the system does not rely on any particular nodes to function. DLT *permits* this extreme case, but it is also consistent with designs featuring a greater degree of centralization.

The platforms used for wCBDC pilots typically exhibit significant centralization, in one way or another. This is unsurprising. There is inherently likely to be a greater degree of centralization arising from the influential position held by a central bank (or associated regulatory bodies) in terms of issuance, rule-setting and oversight. Furthermore, preparedness to accept a degree of centralization is higher among wholesale participants, relative to many of the communities involved in the establishment and operation of public blockchains like Bitcoin, where a desire for decentralization is almost a point of dogma.

Nevertheless, central banks must consider what the appropriate level of decentralization should be - trading off resilience against other properties such as day-to-day performance and privacy protections.

5.3.1 Decentralized storage of records

The avoidance of a single point of failure (SPF) is a key benefit of a decentralized system. In designs where the ledger is recorded by multiple parties, even if one node (or a relatively small subset of nodes) goes offline, the information saved by the offline node is not trapped, but can be reliably sourced from elsewhere. Thus, any activities that rely on accurate information on the ledger can continue. Furthermore, if the node comes back online it can quickly recover the full record of transactions, including updates since the time of its failure.²⁶

Thus, the saving of data in multiple locations not only helps avoid errors, disagreement, and costs of reconciling incompatible databases, but it also helps avoid SPF in storage. Clearly this is especially attractive if the risk of cyber-attacks on a single storage point is high, either with the intent destroying the record or manipulating it in some way. The hope, therefore, is that a DLT system might allow less stringent requirements for securing a given node (as its constant

²⁶Indeed, if there were to be a catastrophic system failure of multiple nodes, restarting the system should also be made easier by an agreed upon record saved in multiple locations.

availability is not so vital), ultimately reducing costs of maintaining a node, and possibly reducing barriers to entry among participants. Of course, care should be taken in trading off resilience arising from decentralization against resilience derived from securing each node.

It is worth noting that these benefits of decentralized record keeping might be undermined by certain features of DLT systems designed to enhance privacy and performance. For example, if data is stored off-chain it may be more difficult to reliably reconstitute (for recovery or perhaps for audit) a complete record after a system or node failure. Also, in the channels approach of Hyperledger Fabric with limited information being broadcast globally and in Corda's peer-to-peer design with only counterparties recording their transaction, it might also be more difficult to recover from system or node failures, or for an overarching audit to be undertaken.

5.3.2 Avoiding transaction disruption

A decentralized DLT design should also allow other activities to continue in the case of a node malfunctioning, beyond simply *retrieving* accurate information. An obvious case to consider here is whether the central bank is required for the system to function effectively. It is difficult to envisage *all* processes surviving such a situation. As noted in [Aber \(2020\)](#), [Atom \(2021\)](#) and other pilots, the issuance of wCBDC seems intrinsically centralized, while exchanging reserves for wCBDC (such as in the pledging and redemption steps of a DDR model) also appear difficult to insulate from central bank downtime. Nevertheless, various pilots have explored how transactions and other processes might continue during central bank downtime.

[Aber \(2020\)](#) was especially focused on exploring the degree of decentralization that could be achieved with DLT. Subject to transaction correctness and settlement finality the pilot succeeded in achieving '*payment and settlement between commercial banks...without the central bank(s) nodes being available*', though the counterparties of the transaction *did* need to be online. The key to this was a form of transaction validation that permitted a quorum of participants not necessarily including the central bank, making use of Hyperledger Fabric's flexibility over consensus mechanisms. Similarly, using Hyperledger Besu, Project Atom also showed the feasibility of a DLT system that allows transactions to continue without the central bank (RBA) node being online. Furthermore, [Khokha₁ \(2018\)](#) implemented a system where the SARB could *observe* transaction details without being required for transaction *validation*.

Note, however, that if the central bank goes offline, there will typically be a reduction in participants available for validation. Thus, while the central bank is not strictly necessary (the minimal number of validators may still be attained without the central bank), there is perhaps *some* increase in risk to transactions proceeding, given risks to other validators' availability. A similar point is raised in [Atom \(2021\)](#) where it is noted that depending on the consensus protocol used, the loss of a candidate validator node - such as the RBA node - raises the probability of the system being unable to verify transactions. Essentially, protocols that rely on a certain quorum

of participants as part of the consensus/validation mechanism exhibit a nonlinearity: the loss of *some* participants may not degrade the system immediately, but further losses beyond a critical point can lead to a shutdown (see [Stella₁ \(2017\)](#)).

As in [Ubin₂ \(2017\)](#)'s *Hyperledger Fabric-based* experiments, [Aber \(2020\)](#)'s gridlock resolution mechanism was not entirely decentralized. However, it is worth noting that [Ubin₂ \(2017\)](#) experimented with other platforms also. Under both Quorum and Corda frameworks gridlock resolution was implemented in a way that did not require the central bank to be online, though the Corda implementation at the time required *transaction participants* to remain online during the gridlock resolution.

While it may be *feasible* to avoid central bank involvement for a transaction to be processed, whether it is *desirable* requires a broader view of the aims of the wCBDC system. Decentralizing can lead to extra complexity or performance degradation. These must be set against other design priorities and the ability to make a SPF as reliable as possible. Not all pilots - even in recent years - have sought to remove central agents from the transaction process. For example, in [BdF \(2021\)](#) the central bank and central securities depository are required in validating transactions involving CBDC and securities movements, respectively.

5.3.3 Novel sources of centralization

It is important to note that while pilots have made great progress on decentralizing away from central bank involvement, other examples of SPF may emerge, depending on the DLT implementation chosen. Centralization may re-emerge through novel channels.

Within the Corda framework, for example, the [notary service](#) appears to be a candidate for SPF ([Ubin₂ \(2017\)](#)). Similarly, the [orderer service](#) in HLF, depending on how it is implemented, can lead to some centralization. In [Ubin₂ \(2017\)](#) the pilot was run with a single node ordering service, though it is noted that multiple node implementations would also be possible.²⁷ Generally, it should be understood that different consensus mechanisms may imply varying degrees of centralization.

Administrative systems, such as network mapping services and certification services that underpin communication and identification within DLT also should be examined for whether they imply systemic risk of failure. For example, [Stella₁ \(2017\)](#) considered what would happen if the HLF Certificate Authority failed - showing that transactions would be rejected until the authority came back online.

It is less obvious that Quorum suffers from such problems to the same degree as it is built to be essentially a permissioned version of Ethereum, augmented with additional features. Quorum retains Ethereum's property that all nodes may carry out the same roles, implying a highly

²⁷In Corda it is also apparently possible to run a notary service as a cluster, if required, as noted in [Jasper₂ \(2017\)](#).

decentralized structure as default. Nevertheless, it is important to be alert to where elements of centralization may arise - perhaps through permissioning or the consensus protocols used.

In situations where something has ‘gone wrong’ it may be difficult to avoid centralization. For example, in *Ubin₃* (2018), which studies interoperability of distributed ledgers, a central authority plays the role of an arbitrator when a dispute between counterparties arises. Again, it is a matter of debate as to whether this sort of centralization is desirable or whether it is important to find some more decentralized approach to resolving disputes.²⁸

While it is natural for arbitration to involve some centralization, there may be some hope for smart contracts to reduce the rate of disputes occurring in the first place, such that the centralization of dispute resolution is less likely to expose the system to a damaging SPF. It is a fascinating open question how far disputes and ambiguity can be reduced by automation of business and regulatory logic through smart contracts and other DLT features such as ‘oracles’. Indeed, in *Ubin₃* (2018) one long term goal is stated as:

This project seeks to explore whether smart contracts can act as a replacement for compliance requirements due to the inherently integrated rulebook by, for example, having arbitration built into the smart contracts’ conditions to validate transfer instructions. Furthermore, error margins that define an erroneous trade can be pre-programmed into the trade to prevent market participants from partaking in such trades.

²⁸The arbitrator may also be able to intervene when a transaction is troubled not because of disputes but because one counterparty goes offline during the execution of the cross ledger process.

6 Legal and policy issues

6.1 Legal matters

The focus of most pilots has been on the practical feasibility of introducing a wCBDC and simulating core *technical* mechanisms. As such, legal details have received lesser attention, though as the sophistication and extent of pilots continues to grow, the time will soon come when the legal framework will need to be formalized. In this section we summarize some legal issues raised in the various pilots, beginning with those that refer to tokens generally, before focusing specifically on wCBDC.

6.1.1 Tokens

The legal status and even the terminology around tokens are not settled, owing to the novelty of the technology. Starting with the terminology, it is useful to distinguish different categories of token. A simple taxonomy is found in [Bundesbank \(2020\)](#), which we repeat in slightly abridged form here:

Information token

Securities are stored in conventional legacy systems (i.e. a custodian-controlled database, a vault, etc.) and corresponding tokens just carry specific information on these securities. The information stored in the legacy system would take precedence over the information stored on the DLT. And the factual control over the token, or more precisely, its transfer from one participant to another, would not entail a transfer of any rights or obligations with regards to the corresponding security.

Extended token

The same information stored on the DLT would take precedence over the information stored in the legacy system. Control over the security token would provide evidence of rights/obligations regarding the corresponding securities (like for example ownership of the security). However, while control over the token would only indicate eg. ownership, it indirectly affects it through an additional legal construct like for example a trusted third party which ensures that ownership of the security and control over the corresponding token do not diverge. This is because the main legal reference object would still be the security existing outside the DLT Layer.

Native token

The security in total is issued in the form of a token and therefore carries all information as well as all rights/obligations to the security. The legal reference object would be the token and the DLT system itself; any change in token ownership would also automatically trigger legal consequences.

In our reading of the wCBDC pilots, the ‘information token’ model does not appear to be the aim of any of the systems considered. Thus, we focus on the ‘extended token’ and ‘native token’

models. In fact, these two cases seem to be essentially what *Khokha*₂ (2022) refers to as ‘tokenized securities’ and ‘security tokens’, respectively (see appendix C for more details):

The tokenisation of securities can take one of two forms: (i) security tokens, which as with the debenture tokens entail issuing the security directly on DLT where the token is the security; or (ii) tokenised securities, where an existing security is tokenised or encapsulated in a token wrapper. The implications between the two may differ in that, tokenised securities will have to provide verification of the legal right to the underlying asset and the fact that it is ring-fenced for purpose of tokenisation.

As we will see below, among the pilots there are cases of wCBDC being designed to take the form of a native token and also cases where the wCBDC is, effectively, a tokenized security, with the underlying asset being some balance of reserves.

The quotes above both hint at the somewhat more complex legal issues surrounding tokenized securities, relative to native tokens. Nevertheless, the legal standing for both forms of tokens appears still to be rather unclear in many jurisdictions. Native security tokens may not even be referred to in existing legislation, while the connection between underlying assets and the tokenized securities associated with them also requires careful analysis. As noted in *BIS* (2019):

Traditional wholesale payment arrangements, such as real-time gross settlement systems, are based on long-established and well founded legal structures and arrangements. Protections under existing legislation, including payments law, contract law, settlement finality provisions and conflicts of law regimes in their local jurisdictions, were not written with wholesale token arrangements in mind, and may not necessarily extend to such arrangements, leading to possible legal uncertainties and risks.

6.1.2 What form should the wCBDC take?

Clearly, a wCBDC pilot requires *some* manifestation of central bank money, but it is unclear what form. Various approaches have been taken in existing wCBDC pilots. While the choice of approach may not dramatically influence *technical* features of the blockchains on which the wCBDC exists, the differences in approach are important to understand. Specifically, we contrast the Digital Depository Receipt (DDR) model, with a model where wCBDC is issued as a direct liability of the central bank.

A Digital Depository Receipt (DDR) model was adopted by early pilots (*Jasper*₂ (2017) and *Ubin*₁ (2017)). Speaking somewhat loosely (consult *Jasper*₂ (2017) pp. 43 for tortured legalese) the DDR is a claim on reserves held in accounts with the central bank. As such, only institutions with access to traditional reserve accounts can receive wCBDC *directly*. Under this approach, there is a pledge-generation (PG) stage and a redeem-return (RR) stage at the two ends of the DDR lifecycle, both of which entail two steps.

In the PG stage a (commercial) bank pledges a certain value of reserves to a hypothecated account at the central bank²⁹ It receives in exchange the corresponding DDR, sent to it by the central bank on the distributed ledger. Note that only the central bank can issue DDR.

The RR stage occurs when a bank requests redemption of a certain amount of DDR by sending the appropriate amount to the central bank, on ledger. The central bank then allows the corresponding value of reserves to be removed from the hypothecated account, and credited to the bank's standard reserve account. Note that reserves can only be extracted from the hypothecated account if an equivalent value of DDR is extinguished from the bank's holdings on ledger.

From our reading, it is not entirely clear if DDR is *strictly*, central bank money in the sense of being a direct liability of the central bank. A holder possesses a claim to reserves, which *are* direct liabilities. Nevertheless, it is not outlandish to regard DDR as 'central bank money' in a slightly looser sense given an appropriate legal framework tightening up the connection between DDR and the underlying reserves (discussed further, below). Indeed, in *Ubin₁ (2017)* the 'Jasper model' is referred to, and it is suggested that DDR appear explicitly on the central bank balance sheet (see *Ubin₁ (2017)*, pp. 29-30).

Helvetia₁ (2020) explored the implementation of wCBDC under two distinct legal frameworks. First, under existing Swiss law and second, anticipating a [new legal framework](#) that explicitly acknowledged 'ledger-based securities'. In both cases the wCBDC was designed to be '*a distinct central bank liability that would emerge on the central bank balance sheet*', without the extra complexity of being a claim on reserves, as in the DDR model.

While both legal approaches to wCBDC considered in the Helvetia pilot implied the creation of wCBDC as a direct liability of the central bank, the distinction between the two is worth considering. In the first approach, an elaborate set of civil law provisions are invoked such that the wCBDC is regarded as a 'payment instruction' and one which has force due to SDX/SNB establishing a system - as part of the pilot - that executes these instructions. This is both complicated and, likely, somewhat Switzerland-specific in its details. Nevertheless, it is an example of what sort of legal gymnastics may be required in any country that attempts to issue wCBDC in a legal framework designed prior to the emergence of tokenization and DLT.

In contrast, the issuance of wCBDC as a 'native' token in Helvetia's second legal approach implies the identification of a claim against the SNB with ownership of the token, without any reference to the presence of an intermediary.³⁰ The token intrinsically embodies the holder's right, as opposed to simply capturing a 'payment instruction'.

²⁹This hypothecated account may be a single pooled or 'omnibus' account to which all banks contribute or, in some pilots, an account solely associated with the particular bank, with each bank having such an account. See [BoE \(2021\)](#) for a discussion of omnibus accounts in the British context.

³⁰Again, the semantics in the Helvetia pilot documentation are very subtle. In the legal appendix it is stated 'a ledger-based security can be transferred peer-to-peer on the ledger without the involvement of an intermediary'. However, throughout the pilot, the SDX DLT platform is used, The point (perhaps) being made is that the DLT operates as a technological, rather than legal, enabler. That is, one could envisage the security being transferred without reference to SDX whatsoever, even though transfers were made exclusively via SDX in the trial.

Other central banks have issued wCBDC natively as direct liabilities, such as the SARB in *Khokha₂* (2022) and Banque de France in *BdF* (2021). In these cases wCBDC is not a ‘claim’ on reserves - it is purchased in exchange for them.

6.1.3 Settlement finality: Technical vs legal

The distinction between *legal* and *technical* is frequently raised in wCBDC pilots. If a wCBDC is to be adopted as a settlement asset in systems that satisfy the PFMI, it is vital that the formalization of legal finality is secure. Again the early phases of Jasper, which used a DDR approach, point the way (our emphasis):

*To ensure legal settlement finality, Project Jasper was structured such that exchange of DDR between platform participants would be equivalent to a full and irrevocable transfer of the underlying claim on central bank deposits. **If an appropriate legal structure were in place to support this**, one could argue that settlement of value exchange on the Jasper platform is ultimately achieved in central bank money (a key requirement of the PFMI). For all intents and purposes, DDR appears to function as central bank money in the system. **Nevertheless, the strength of the legal basis for settlement finality on the Jasper platform warrants further discussion with legal experts in the payments field.***

Similarly, in *Inthanon₁* (2019) it is emphasized that settlement finality ‘has to be clearly defined in both operational and legal aspect before pushing the PoC towards a production-grade system’.

Technical finality can be assured by using the sort of consensus mechanisms incorporated in private blockchains, and by protecting the system from any nefarious or erroneous attempts to retroactively amend the record of transactions (the ledger must be immutable). The issue remains however as to when (or if) finality occurs in a legal sense.

We have already noted the legal subtleties around the DDR model. Similar issues seem to be at play in the framework of *Khokha₂* (2022) when the wCBDC is ported from the ‘CBDC zone’ to the ‘Khokha hub’ (see appendix C). In the Khokha hub, the ported wCBDC appears to be a tokenized security. Like DDR in *Jasper₂* (2017), it is treated for all intents and purposes as the same as the un-ported CBDC but, in their conclusions, SARB express concern over the ‘*split between when technical/operational settlement and legal settlement*’ that results from porting.

Discussions of finality may also arise in relation to bankruptcy proceedings. Clearly a legal structure would also need to address issues over ownership of underlying reserves or assets underpinning tokenized securities in the case of bankruptcy. For example, in the DDR case, a participant may have accumulated net positions in DRR that are not yet reconciled in the underlying reserves settlement system at the time of bankruptcy. A legal stance on who has superior claim to the underlying reserves will be needed.

6.1.4 Powers of the central bank

Central banks may not (yet) have legal authority to issue a production version of a CBDC (see [IMF \(2020\)](#), [Kosse and Mattei \(2022\)](#) and [here](#) for further discussions and evidence on this point). This obstacle may be especially problematic in the retail domain, where there is typically no precedent for digital central bank money. However, legal impediments may also arise in the wholesale case, despite digital currency in the form of reserves having a long history.

Some pilots have explicitly suggested areas where legal authority must be sought or where adjustments to legislation may be necessary. For example, [Inthanon₁ \(2019\)](#) refers to how the Thai ‘Currency Act’ would need to be revisited to underpin the certification of wCBDC as legal tender. Similarly, [Khokha₁ \(2018\)](#) notes the importance of adjusting the legal framework to allow wCBDC to satisfy legal requirements of the [National Payments System Act](#) over settlement finality.

In some cases, however, legal authority may already exist. For example, in [Helvetia₂ \(2022\)](#), it is argued that ‘*provided that the central bank retains necessary wCBDC control and monitoring functions*’ the SNB would be legally permitted to issue a wCBDC on a third-party-operated platform such as SDX. Indeed, the claim is already that the wCBDC *per se* comes within the existing authority of the central bank’s because it can be classed as ‘*an alternative representation of traditional reserve balances*’.

It is important to consider whether the Swiss approach - in which the wCBDC can be legitimized purely as another form of reserves - is applicable elsewhere. It is not at all obvious that it will be. For example, a country may have a much narrower definition of reserves, perhaps tied to some design element that could be inconsistent with how wCBDC is implemented.

6.1.5 Updating the legal architecture

In response to the ambiguities around the legal treatment of tokenized assets and the novel demands arising from rapid private sector adoption, regulatory and legislative strategies do appear to be emerging. In Europe, the EU’s [Digital Finance Strategy](#) has been advancing in recent years, with recent notable developments such as the creation of a [pilot regime](#) for market infrastructures based on DLT.³¹

In the US, the recent [Executive Order](#) on Ensuring Responsible Development of Digital Assets signals the start of a coordinated strategy that will encompass, among other things, the role of DLT in financial markets. Until now, the fragmented regulatory structure in the US has seen the SEC take the lead in many respects. In this regard, the SEC appears to have attempted to repurpose existing structures designed for traditional securities but to be applied to tokens. Attempts are afoot in congress to codify concepts and roles more clearly. Notably, the work of Senators Lummis

³¹The developments are not only in regard to wholesale markets. The European Market in Crypto-assets ([MiCA](#)) regulations are quite broad in scope.

and Gillibrand - captured in the ‘[Responsible Financial Innovation Act](#)’ - attempts to clarify the nature of tokens, as securities or as commodities, and assign regulatory responsibility clearly on that basis.

In the UK, after a [consultation period](#), the UK government [appears](#) set to introduce extensive legislation related to digital assets, as part of a broader financial services and markets bill.

There are several other jurisdictions where progress on these matters is more advanced and which provide interesting insights to build upon. Germany, [Switzerland](#) and Singapore, in particular, have constructed elaborate legal frameworks. Detailed analysis of these frameworks is beyond the scope of this report. However, we note that the 2021 passing of the Swiss ‘[Blockchain Act](#)’ marked an important formalization of the crypto legal environment and represents an interesting reference point for other countries in the future. In dealing with definitions of ownership, treatment of assets in the case of bankruptcy, and in codifying licenses for DLT-based exchanges, the legislation goes a long way to providing the legal architecture that both regulators and the private sector have been seeking.

All these legislative schemes cover topics far broader than CBDC. However, they are no doubt relevant, given the integration of wCBDC into a wide variety of market structures. Furthermore, to the extent that introducing a wCBDC stimulates further innovation in DLT and tokenization it behoves central banks to be intimately involved with the establishment of an appropriate legal framework. As pilots increasingly engage with the broader private sector, there is an obvious opportunity for central banks to influence industry best practice, with a view to it ultimately being codified in legislation. Indeed, in the recently launched [Project Guardian](#), MAS explicitly refers to one of its aims being ‘*develop[ing] rulebook, governance model, and reviewing legal and regulatory frameworks for tokenized assets*’.

Even as new legal frameworks are established, complexity may also arise from the reach of *existing* legislation. For example, in [Inthanon₁ \(2019\)](#), it was noted that a wCBDC payments system could fall within the scope of the extant ‘Payment System Act’. Important KYC/AML and data protection legislation may also interact in unexpected ways with the introduction of wCBDC, given the record-keeping elements that are fundamental to DLT, combined with its novel characteristics of decentralization and cryptographic techniques.

6.2 Monetary Policy

Much debate on CBDC’s implications for monetary policy has focused on the retail case and, almost without exception, pilots have been run to help understand payments systems and other aspects of financial infrastructure, and *not* to explore monetary policy. But there are some monetary policy-relevant issues raised in the wholesale domain also.

6.2.1 Size of the balance sheet

Pilots have been structured such that wCBDC is explicitly ‘associated’ one-for-one with corresponding balances held within the central bank’s legacy accounts system.

This ‘association’ may be *via* the DDR model, where the wCBDC is a representation of underlying claims on reserves ‘locked’ in a technical account at the central bank. If one regards the DDR as only *representations* of central bank liabilities then trivially the balance sheet is unaffected by their issuance. If instead - as suggested in *Ubin₁* (2017) - they are regarded as liabilities, then the size of the balance sheet would not change, but the composition would.³²

Alternatively, the ‘association’ may be such that the wCBDC is *itself* the claim on the central bank. If they are then ‘bought’ with reserves at a 1:1 exchange rate, this removes reserves (one type of liability) from the system, even as wCBDC balances (another type of liability) increase. Thus the overall amount of central bank money remains flat, but the composition changes (see *Helvetia₁* (2020) for explanations of the effect on public, bank and central bank balance sheets).

Variation in the balance sheet size is observed in pilots that introduce repo facilities, such as *BdF* (2021) and *Blockbuster* (2018). In *BdF* (2021) cross ledger repo operations were allowed where banks sold (traditional) securities to the central bank to receive wCBDC. Note that such variation in the size of the balance sheet isn’t a peculiarity of wCBDC systems - the same would apply in a repo with *reserves* exchanged for collateral. The fact that the operations are intraday also insulates the economy from a persistent monetary impulse.

6.2.2 Remuneration

The issue of remuneration has been avoided by several pilots by only using intraday wCBDC. Any such wCBDC issued during the day is ultimately redeemed/destroyed at end of day (see *Blockbuster* (2018)) and (*Jasper₂* (2017), for example). However, in *Helvetia₁* (2020), for example, the wCBDC did not have to be redeemed at the end of the day, and the decision was taken to allow it to earn interest overnight at the same rate offered on traditional reserves. In *Ubin₁* (2017), banks could hold CBDC balances overnight on the ledger though in this case the balances were unremunerated.

³²The legal and semantic distinctions here are very subtle and any country launching a pilot would need to consult legal specialists to tune the precise wording and concepts appropriately.

We briefly cite *Aber* (2020), despite our main focus on domestic wCBDC, because it raised issues related to remuneration. The wCBDC issued was unremunerated but since one central bank in the pilot paid interest on reserves, it implied there would be an opportunity cost of holding wCBDC overnight. Arguably it was the *difference* in rates between the two countries' reserve systems that was the key complexity, rather than remunerating CBDC *per se*. Nevertheless, the solutions proposed in the pilot are worth considering, especially if interoperability of wCBDC systems internationally becomes a prominent issue.

Clearly, if reserves are being remunerated then the relative rates (wCBDC vs reserves) will have implications for what sort of central bank money participants would want to hold. This, in turn, could perhaps influence money market functioning if the two types of digital money are not perfect substitutes in their suitability for various market activities.

6.2.3 Liquidity and money markets

There is perhaps some risk of fragmentation of liquidity arising from the co-existence of two forms of wholesale central bank money - reserves and wCBDC. One could envisage banks separating into groups on the basis of which of the two 'money markets' they would predominantly operate in, with the interaction between the two markets subject to frictions. Alternatively there might be additional complexity and cost for a single bank coordinating its activities in both markets.

While these are concerns that should be considered carefully, they seem more appropriate for latter stage trials. Also, while there may conceivably be some risk of fragmentation, it seems that this risk is plausibly greater in the absence of wCBDC if that absence allows the excessive proliferation of privately provided settlement assets that wCBDC might otherwise crowd out.

It is noted in *Atom* (2021) that if the amounts of existing balances backing the wCBDC were to become substantial (that is, orders of magnitude beyond the scale at which pilots typically operate), they might inject volatility in the supply of unencumbered reserves in the interbank markets. Since rates in these markets are targeted as part of standard monetary policy, such fluctuations could conceivably make the day-to-day implementation of policy (through repo operations) somewhat more complicated.³³ The scale of wCBDC balances ultimately will be a design decision for a production-ready system and a full discussion is beyond the scope of this report.

³³This effect would be akin to how 'autonomous factors' (cash conversion and government treasury account activities) affect such rates through more familiar channels (*Ihrig et al. (2020)*).

References

- Aber*, “Project Aber: Saudi Central Bank and Central Bank of the U.A.E. Joint Digital Currency and Distributed Ledger Project,” Technical Report, Saudi Central Bank and Central Bank of the United Arab Emirates 2020. Available at https://www.sama.gov.sa/en-US/News/Documents/Project_Aber_report-EN.pdf.
- Androulaki, Elli, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick**, “Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains,” in “Proceedings of the Thirteenth EuroSys Conference” EuroSys ’18 Association for Computing Machinery New York, NY, USA 2018.
- Atom*, “Project Atom: Exploring a Wholesale CBDC for Syndicated Lending,” Technical Report, Reserve Bank of Australia 2021. Available at https://www.rba.gov.au/payments-and-infrastructure/central-bank-digital-currency/pdf/project-atom-report_2021-12.pdf.
- Auer, Raphael**, “Embedded supervision: how to build regulation into blockchain finance,” BIS Working Papers 811, Bank for International Settlements September 2019.
- Bains, Parma**, “Blockchain consensus mechanisms: A primer for supervisors,” FinTech Notes, IMF 2022. Available at <https://www.imf.org/en/Publications/fintech-notes/Issues/2022/01/25/Blockchain-Consensus-Mechanisms-511769>.
- BCB**, “Distributed ledger technical research in Central Bank of Brazil,” Technical Report, Banco Central do Brasil 2017. Available at https://www.bcb.gov.br/htms/public/microcredito/Distributed_ledger_technical_research_in_Central_Bank_of_Brazil.pdf.
- BdF**, “Experimenting settlement of French government bonds in Central Bank Digital Currency with blockchain technology,” Technical Report, Banque de France 2021. Available at <https://www.euroclear.com/content/dam/euroclear/news%20%20insights/Format/Whitepapers-Reports/settlement-french-government%20bonds-in-cbdc-with-blockchain.pdf>.
- Bech, Morten and Kimmo Soramäki**, “Liquidity, gridlocks and bank failures in large value payment systems,” *E-money and Payment Systems Review*, 2002, pp. 111–126.
- Bech, Morten Linnemann, Jenny Hancock, Tara Rice, and Amber Wadsworth**, “On the future of securities settlement,” *BIS Quarterly Review*, March 2020.
- BIS**, “Delivery versus payment in securities settlement systems,” Technical Report, IOSCO Technical Committee and BIS Committee on Payment and Settlement Systems 1992. Available at <https://www.bis.org/cpmi/publ/d06.pdf>.
- , “The role of central bank money in payment systems,” Technical Report, Bank for International Settlements 2003. Available at <https://www.bis.org/cpmi/publ/d55.htm>.
- , “Principles for financial market infrastructures,” Technical Report, IOSCO Technical Committee and BIS Committee on Payment and Settlement Systems 2012. Available at <https://www.bis.org/cpmi/publ/d101a.pdf>.

- , “Central bank digital currencies,” Technical Report, Bank for International Settlements - Committee on Payments and Market Infrastructures, Markets Committee March 2018. Available at <https://www.bis.org/cpmi/publ/d174.pdf>.
- , “Wholesale digital tokens,” Technical Report, Bank for International Settlements 2019. Available at <https://www.bis.org/cpmi/publ/d190.htm>.
- , “Central bank digital currencies: system design and interoperability,” Technical Report, Bank for International Settlements March 2021. Available at https://www.bis.org/publ/othp42_system_design.pdf.
- Blockbaster**, “Blockbaster: Final report,” Technical Report, Deutsche Börse and Deutsche Bundesbank 2018. Available at <https://www.bundesbank.de/resource/blob/766672/2e2ccde1855071cb55ae97c7b025da8d/mL/2018-10-25-blockbaster-final-report-data.pdf>.
- BoE**, “Central Bank Digital Currency: Opportunities, challenges and design,” Discussion paper, Bank of England 2020.
- , “Bank of England Omnibus Accounts - Access Policy,” Technical Report, Bank of England 2021. Available at <https://www.bankofengland.co.uk/-/media/boe/files/payments/boeomnibusaccounts.pdf>.
- Bundesbank, Deutsche**, “How can collateral management benefit from DLT?,” Technical Report, Deutsche Bundesbank 2020. Available at <https://www.bundesbank.de/resource/blob/823072/4d14afd4b6dbffa94a46ee52f46e99bd/mL/how-can-collateral-management-benefit-from-dlt-data.pdf>.
- DTCC**, “Building the settlement system of the future,” Technical Report, Depository Trust and Clearing Corporation 2021.
- eAUD**, “Australian CBDC Pilot for Digital Finance Innovation,” Technical Report, Reserve Bank of Australia and Digital Finance CRC 2022. Available at <https://www.rba.gov.au/payments-and-infrastructure/central-bank-digital-currency/pdf/australian-cbdc-pilot-for-digital-finance-innovation-white-paper.pdf>.
- Furgal, Adam, Rodney Garratt, Zhiling Guo, and Dave Hudson**, “A Proposal for a Decentralized Liquidity Savings Mechanism With Side Payments,” Technical Report, R3 2018.
- Garratt, Rod, Michael Junho Lee, Brendan Malone, and Antoine Martin**, “Token- or Account-Based? A Digital Currency Can Be Both,” Liberty Street Economics 20200812, Federal Reserve Bank of New York August 2020.
- Gorton, Gary and Andrew Metrick**, “Securitized banking and the run on repo,” *Journal of Financial Economics*, 2012, 104 (3), 425 – 451. Market Institutions, Financial Market Risks and Financial Crisis.
- Gorton, Gary B. and Jeffery Zhang**, “Taming wildcat stablecoins,” *University of Chicago Law Review (forthcoming)*, 2023, 90.
- Helvetia*₁, “Project Helvetia: Settling tokenised assets in central bank money,” Technical Report, Swiss National Bank and BIS Innovation Hub and SIX 2020. Available at <https://www.bis.org/publ/othp35.pdf>.
- Helvetia*₂, “Project Helvetia Phase II: Settling tokenised assets in wholesale CBDC,” Technical Report, Swiss National Bank and BIS Innovation Hub and SIX 2022. Available at https://www.snb.ch/en/mmr/reference/project_helvetia_phase_ii_report/source/project_helvetia_phase_ii_report.en.pdf.

- Hyperledger**, “Hyperledger Blockchain Performance Metrics,” Technical Report, Hyperledger Performance and Scale Working Group 2018. Available at <https://www.hyperledger.org/learn/publications/blockchain-performance-metrics>.
- Ihrig, Jane E., Zeynep Senyuz, and Gretchen C. Weinbach**, “The Fedâs âAmple-Reservesâ Approach to Implementing Monetary Policy,” Finance and Economics Discussion Series 2020-022, Board of Governors of the Federal Reserve System (U.S.) February 2020.
- IMF**, “Legal Aspects of Central Bank Digital Currency: Central Bank and Monetary Law Considerations,” IMF Working Papers 2020/254, International Monetary Fund 2020.
- Inthanon*₁, “Inthanon Phase 1: An application of Distributed Ledger Technology for a Decentralised Real Time Gross Settlement system using Wholesale Central Bank Digital Currency,” Technical Report, Bank of Thailand 2019. Available at https://www.bot.or.th/Thai/PaymentSystems/Documents/Inthanon_Phase1_Report.pdf.
- Inthanon*₂, “Inthanon Phase 2: Enhancing Bond Lifecycle Functionalities & Programmable Compliance Using Distributed Ledger Technology,” Technical Report, Bank of Thailand 2019. Available at https://www.bot.or.th/English/FinancialMarkets/ProjectInthanon/Documents/Inthanon_Phase2_Report.pdf.
- Jasper-Ubin**, “Jasper-Ubin Design Paper: Enabling Cross-Border High Value Transfer Using Distributed Ledger Technologies,” Technical Report, Bank of Canada and Monetary Authority of Singapore 2019. Available at <https://www.mas.gov.sg/-/media/Jasper-Ubin-Design-Paper.pdf>.
- Jasper*₁, “Project Jasper Primer,” Technical Report, Bank of Canada and Payments Canada and R3 2017. Available at <https://payments.ca/insights/research/project-jasper-primer>.
- Jasper*₂, “Project Jasper: A Canadian Experiment with Distributed Ledger Technology for Domestic Interbank Payments Settlement,” Technical Report, Bank of Canada and Payments Canada and R3 2017. Available at https://payments.ca/sites/default/files/2022-09/jasper_report_eng.pdf.
- Jasper*₃, “Jasper Phase III: Securities settlement using distributed ledger technology,” Technical Report, Bank of Canada and Payments Canada and R3 2018. Available at https://payments.ca/sites/default/files/2022-09/jasper_phase_iii_whitepaper_EN.pdf.
- Khokha*₁, “Project Khokha: Exploring the use of distributed ledger technology for interbank payments settlement in South Africa,” Technical Report, South African Reserve Bank 2018. Available at https://www.resbank.co.za/content/dam/sarb/quick-links/fintech/SARB_ProjectKhokha_20180605.pdf.
- Khokha*₂, “Project Khokha 2: Exploring the implications of tokenisation in financial markets,” Technical Report, South African Reserve Bank 2022. Available at <https://www.resbank.co.za/content/dam/sarb/publications/media-releases/2022/project-khokha-2/Project%20Khokha%20%20Summary%20Report%206%20April%202022.pdf>.
- Koens, T. and E. Poll**, “Assessing interoperability solutions for distributed ledgers,” *Pervasive and Mobile Computing*, 2019, 59, 101079.
- Kosse, Anneke and Ilaria Mattei**, *Gaining momentum â Results of the 2021 BIS survey on central bank digital currencies* number 125. In ‘BIS Papers.’, Bank for International Settlements, November 2022.
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The Aggregate Demand for Treasury Debt,” *Journal of Political Economy*, 2012, 120 (2), 233–267.

- **and** —, “The impact of Treasury supply on financial sector lending and stability,” *Journal of Financial Economics*, 2015, 118 (3), 571–600.
- Lee, Michael Junho, Antoine Martin, and Benjamin Müller**, “What Is Atomic Settlement?,” Liberty Street Economics 20221107, Federal Reserve Bank of New York November 2022.
- Norman, Ben**, “Financial Stability Paper No 7: Liquidity Saving in Real-Time Gross Settlement Systems - an Overview,” Bank of England Financial Stability Papers 7, Bank of England May 2010.
- Panetta, Fabio**, “Demystifying wholesale central bank digital currency - a speech at the Symposium on Payments and Securities Settlement in Europe â today and tomorrow hosted by the Deutsche Bundesbank, Frankfurt am Main, September 26, 2022,” Speech, European Central Bank September 2022.
- Sheikh, J.**, *Mastering Corda: Blockchain for Java Developers*, O’Reilly Media, 2020.
- Stella*₁, “Payment systems: liquidity saving mechanisms in a distributed ledger environment,” Technical Report, European Central Bank and Bank of Japan 2017. Available at https://www.ecb.europa.eu/pub/pdf/other/ecb.stella_project_report_september_2017.pdf.
- Stella*₂, “Securities settlement systems: delivery-versus-payment in a distributed ledger environment,” Technical Report, European Central Bank and Bank of Japan 2018. Available at https://www.ecb.europa.eu/pub/pdf/other/stella_project_report_march_2018.pdf.
- Stella*₃, “Synchronised cross-border payments,” Technical Report, European Central Bank and Bank of Japan 2019. Available at <https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.miptopical190604.en.pdf>.
- Stella*₄, “Balancing confidentiality and auditability in a distributed ledger environment,” Technical Report, European Central Bank and Bank of Japan 2020. Available at <https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.miptopical200212.en.pdf>.
- Ubin*₁, “Project Ubin: SGD on Distributed Ledger,” Technical Report, Monetary Authority of Singapore and Deloitte 2017. Available at <https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin--SGD-on-Distributed-Ledger.pdf>.
- Ubin*₂, “Project Ubin Phase 2: Re-imagining Interbank Real-Time Gross Settlement System Using Distributed Ledger Technologies,” Technical Report, Monetary Authority of Singapore and Deloitte 2017. Available at <https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin-Phase-2-Reimagining-RTGS.pdf>.
- Ubin*₃, “Project Ubin: Delivery versus Payment on Distributed Ledger Technologies,” Technical Report, Monetary Authority of Singapore and Deloitte 2018. Available at <https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin-DvP-on-Distributed-Ledger-Technologies.pdf>.
- Ubin*₅, “Project Ubin Phase 5: Enabling broad ecosystem opportunities,” Technical Report, Monetary Authority of Singapore and Temasek 2021. Available at <https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin-Phase-5-Enabling-Broad-Ecosystem-Opportunities.pdf>.
- WEF**, “Bridging the Governance Gap: Interoperability for blockchain and legacy systems,” Technical Report, World Economic Forum, Centre for the Fourth Industrial Revolution December 2020. White paper available at <https://www.weforum.org/whitepapers/bridging-the-governance-gap-interoperability-for-blockchain-and-legacy-systems/>.
- World Bank**, “Blockchain Interoperability,” Technical Report, World Bank - Technology & Innovation Lab March 2021. Available at <http://documents1.worldbank.org/curated/en/373781615365676101/pdf/Blockchain-Interoperability.pdf>.

Zand, M., X. Wu, and M.A. Morris, *Hands-On Smart Contract Development with Hyperledger Fabric V2*, O'Reilly Media, 2021.

Appendices

A Blockchain features

A.1 Permissioned and permissionless blockchains

The blockchain systems that most people think of when discussing DLT are Bitcoin and Ethereum. These systems can be joined and used by anybody who sets up client software and connects to the network - there are no gatekeepers. Information about all ‘transactions’ is transmitted to all participants and, ultimately, recorded by all participants in the synchronized ledger. The earliest wCBDC PoCs often adopted technologies designed for such systems (standard Ethereum clients), while developing basic operational experience and understanding. The permissioning was thus not a feature of the DLT *per se*, with access being separately controlled in the basic sense of a restricted set of computers or virtual machines being used to communicate with each other. Examples of such an approach are [Ubin₁ \(2017\)](#) and [Jasper₁ \(2017\)](#).¹ Nevertheless, it was explicitly acknowledged from the start that tailored solutions, designed explicitly for private networks and with inbuilt permissioning, would be more appropriate for wholesale systems.

In a permissioned system, parties must be approved to join the ledger and access is closely controlled, in line with agreed policies. In the CBDC and financial markets domain, such policies will likely be set by governments, regulators, central banks and - possibly - approved industry bodies. Once granted access, users prove their identity on the basis of digital signatures. These signatures can be confidently accepted as being operated by approved parties, whose identities are jointly registered with the keys underpinning the signatures by a trusted [Certificate Authority](#). On joining a permissioned network, a party would typically be required to obtain a registered certificate.²

A.2 Consensus and finality

In blockchain systems, there is a need for network participants to agree on what constitutes the ‘true’ history of valid transactions - what transactions should be added to blocks and what sequence of blocks is the accepted ‘state’ of the system. How this comes about is referred to as the ‘consensus mechanism’ of the blockchain.³ The term ‘consensus’ connotes a distributed form of agreement, reflecting the fact that most (public) blockchains - notably Bitcoin and Ethereum - do not rely on a centralized trusted single party to impose this agreement.

Bitcoin and early versions of Ethereum (notably those used in first wave wCBDC pilots) relied upon a ‘Proof of Work’ (PoW) consensus mechanism, whereby nodes would compete to solve a time consuming computational problem to win the right to add a block of transactions to the chain. The benefits of PoW were that it was robust to an absence of trust among a maximally decentralized set of participants. The protocol is highly secure despite its lack of a trust assumption. Attacks aimed at manipulating the ledger would require more than half of the network’s computational resources to be mobilized to be successful - something that is typically thought to be essentially infeasible.

Blockchains based on PoW consensus admit the possibility of ‘forks’, whereby two nodes may ‘simultaneously’ solve the aforementioned computational problem and broadcast valid blocks to the network. In a fork, alternative histories of transactions - captured in alternative chains stemming from an initial block - may co-exist for some time (and possibly continue to have blocks added to them) as awareness of the competing blocks spreads through the network. Ultimately the ‘[longest chain](#)’ rule will see to it that only one branch of the fork eventually is adopted across the network. In such a situation, a transaction included on one of the competing chains may ultimately be

¹[Ubin₁ \(2017\)](#) also included a secondary experiment with a version of Quorum explicitly designed as a permissioned implementation of Ethereum.

²For an introduction to Digital Certificates and Certificate Authorities, see [Sheikh \(2020\)](#).

³See [here](#), [here](#) and [Bains \(2022\)](#) for accessible discussions.

reversed if the chain ultimately adopted does *not* include that transaction. The removal of blocks that had been incorporated into the rejected chain is referred to as a ‘reorganization’.

Reorganizations are unlikely to affect a given transaction once enough blocks have been built ‘on top’ of the block in which the transaction was originally included. This is why various protocols (such as those run by crypto exchanges) require multiple ‘confirmations’ of the block in which the transaction is contained before the transaction is regarded as ‘final’. However, this concept of finality, based on a finite number of confirmations, *still* leaves some (very low) risk of a reorganization. This probabilistic wrinkle *and* the time taken to reduce the probability to a remotely acceptable level, means that such approaches are inappropriate for financial transactions which must be executed at scale and at speed. The possibility of reversal of transactions must be entirely eliminated from the system, in line with widely accepted Principles of Financial Market Infrastructure (BIS (2012)).

The earliest wCBDC pilots used Ethereum, which at the time relied upon PoW.⁴ However, permissioned blockchains are now exclusively used. The key aspect of these frameworks is that they are typically built on existing relationships (perhaps with auxiliary regulatory oversight) that confer a higher level of intrinsic ‘trust’ than ‘open to all’ public blockchains. As such, consensus mechanisms need not be quite so robust as PoW.

The preparedness to trade off resilience versus performance (speed of achieving consensus and finality and scalability with network size) has led private DLT to alternative protocols, as the chances of crashes or nefarious activity among nodes are thought to be somewhat lower to begin with. This is because permissioned systems - at least in the cases related to wCBDC - not only feature controls on who may join and how they are identified, but also will frequently feature penalties for bad behavior. Among all but the earliest wCBDC pilots, some form of ‘Byzantine Fault Tolerant’ (BFT) consensus protocol is applied. Within this class, there are different permutations, including Practical BFT (pBFT) and Istanbul BFT (iBFT and iBFT 2.0), as well as Proof of Authority (PoA) approaches. Crash Fault Tolerance (CFT) systems also appear - such as the Raft protocol.⁵

These protocols ensure that once a block is added to a chain there is no chance of removal and they also speed up the process of adding a block to the chain. Their speeds and robustness vary on the basis of features such as how many and which nodes are required to construct, propose, validate or vote on blocks.⁶ For example, the Raft protocol is very fast in terms of time to create blocks. However, it is only crash tolerant, rather than being Byzantine fault tolerant, so that malevolent node behavior could be problematic under it.

A.3 Tokens and accounts: Fixing terminology

Wholesale CBDC, interpreted narrowly as ‘digital money’ has existed for some time, in the form of reserve accounts for a select group of financial institutions (typically banks). One option for wCBDC, as discussed previously in the retail case, is to implement the currency using DLT while retaining the account-based emphasis. Another option is to provide the wCBDC using a token-based emphasis.

We use the word ‘*emphasis*’ as the usefulness of distinguishing between token-based systems and account-based systems in the context of crypto currencies is increasingly being questioned, certainly at the deep implementation level (see Garratt et al. (2020) and discussions [here](#) and [here](#) on this point). In many CBDC discussions it is common to see systems categorized on the basis of whether it is the *identity* of the account-holder or the *validity*

⁴Since ‘The Merge’, Ethereum has adopted a PoS protocol, though this also allows for reorganizations.

⁵In the simplest Corda implementation (see appendix D) a single notary node brings about final consensus on a transaction, subject to bandwidth of the notary, this allows consensus to be attained very quickly, though concerns have been raised over the safety of the method. The balance of risks and benefits will depend on the particular application.

⁶This may come down to requiring super-majorities among voting nodes, rather than simple majorities, or randomly changing which nodes get to propose blocks for other nodes to vote on.

of the token that is key to validating a transaction.⁷ However, as noted in [Garratt et al. \(2020\)](#), using Bitcoin as an example (though for our wCBDC purposes their points still stand):

Bitcoin fits the definition of an account-based system. The account is a Bitcoin address, and the private key is the proof of identity needed to transact from that account. Every time a Bitcoin user wants to spend Bitcoin, that user must verify their identity by using their private key. . .

Bitcoin also fits the definition of a token-based system. When someone wants to spend a Bitcoin, the protocol verifies its validity by tracing its history. The current transaction history is used to verify the validity of the ‘object’ being transferred, as other token-based systems also do.

Thus, one can classify Bitcoin both as an account-based system *and* a token-based system according to familiar definitions of the two. These definitions perhaps permit effective partitioning of legacy systems but appear less useful in categorizing blockchain frameworks for currencies.

In terms of functionality, such as a system’s level of anonymity, it is argued that blockchain systems commonly labeled account-based can implement the same capabilities as systems commonly labeled token-based. As noted in [BoE \(2020\)](#) (and also quoted [here](#)):⁸

We do not see any inherent reason that token-based systems would automatically provide anonymity. Both account-based systems and token-based systems can be configured with various identity solutions, ranging from fully anonymous to pseudonymous and to a fully transparent, identifiable solution. . .

In digital form, neither an account-based approach nor a token-based approach would enable cash-like transfers, where a payment can be made without reference to any third party or intermediary. In an account-based system, the accounts of the payer and payee need to be debited and credited by the operator(s) of the ledger. And in a token-based system, in order to prevent double-spending, ownership of tokens needs to be recorded in a ledger, which will need to be updated to reflect any changes in ownership.

Nevertheless, as argued [here](#), it may *still* be useful to retain the token and account terminology. First, while reasonable criteria defining token vs account may not technically partition blockchains, a blockchain may exhibit more characteristics traditionally *associated* with what are understood as accounts or token frameworks. Or it may emphasize them more in a user interface, or when the framework is explained intuitively.

Second, these additional characteristics which are associated with ‘account-ness’ or ‘token-ness’ may also have implications for the legal or regulatory treatment of the system. Third, while it may strictly be possible to implement some functionality (perhaps an anonymity approach, as previously mentioned) in systems which display a greater degree of ‘account-ness’, it may be much simpler and more natural to do in a system that emphasizes a token abstraction.

Finally, and most simply, it is difficult to avoid using the terms token and account to distinguish wCBDC *systems*, given their prominence in existing documentation (see [BIS \(2018\)](#) for an influential report from the early period of CBDC exploration and [BdF \(2021\)](#) for a more recent example).

Nevertheless, it does seem sensible try to reduce the emphasis put on these terms in explaining the structure of any given wCBDC system. Better simply to explain precisely the implementation of the platform - which may involve notions of both tokens and accounts - without using the token or account shorthand to put the system into some high level category.

⁷In [BdF \(2021\)](#) it is stated: ‘An account-based system requires verifying the identity of the payer, while a token-based system requires verifying the validity of the object used to pay’.

⁸This quote refers more generally to digital platforms, rather than DLT and blockchain - hence the reference to ‘operator(s) of the ledger’ but the thrust of the argument is still relevant.

We emphasize that none of the above means that there is no such thing as a token or no such thing as an account. Simply that it may not be appropriate to define the systems by being strictly one or the other. A system can certainly involve a token and emphasize that aspect in its implementation. For example, [BdF \(2021\)](#) favored the token depiction and refer to their wCBDC as a token, since it was apparently more aligned with their desire for ledger interoperability. In this case - since multiple ledgers are involved, possibly with different regulatory authorities or administrators involved - it is perhaps less natural to set a system up emphasizing account-type characteristics. Instead, designing a system where interfaces use the abstraction of a token moving between ledgers is perhaps more natural.

B Terminology: DvP, Atomic Settlement, Immediate Settlement...

This brief appendix addresses a terminology issue: what is meant by ‘delivery-vs-payment’ (DvP). In our reading, pilots often use the term ‘delivery-vs-payment’ (DvP) in a somewhat loose way. The definition found on the [BIS website](#) is:⁹

A securities settlement mechanism that links a securities transfer and a funds transfer in such a way as to ensure that delivery occurs if and only if the corresponding payment occurs.

This definition is arguably equivalent to what is traditionally thought of as ‘atomic’ settlement, given the phrase ‘*if, and only if*’. Atomic settlement is traditionally taken to mean that either both legs succeed or both legs fail - there is no partial execution. Nevertheless, pilots seem to use the term DvP to describe systems where one leg of an exchange may fail while the other completes - indeed pilots give explicit examples of such failures. In the latter case, the process is ‘DvP’ only along its intended ‘[happy path](#)’.

It is perhaps best to think of a system being DvP if the *intent* is for payment and delivery of securities each to occur if and only if the other occurs, even if failures in the settlement process may prevent this in some pathological cases. While this is perhaps a vacuous definition of DvP, we will continue to use the term in this way as it seems to be standard practice.

Finally, it is important to note that a transaction can be atomic (or DvP), but not immediate - it may take time for the completion of the second leg of a transaction.

⁹See [BIS \(1992\)](#) for discussion of the DvP concept and [Bech et al. \(2020\)](#) is an excellent reference to further explore settlement concepts and how they relate to wCBDC and tokenized securities.

C Project Khokha - Phase 2

The full structure of *Khokha₂* (2022) is beyond the scope of this report to discuss and it is difficult to discern *exactly* how the system is implemented on the basis of public documentation. Nevertheless, we here attempt to delve a little more deeply in our analysis of it because the pilot has many interesting aspects that are worth noting.

A distributed ledger (underpinned by Corda) is used for the issuance of wCBDC - this is referred to as the ‘CBDC zone’. In addition, a second, separate distributed ledger network was created for tokenized assets (digital SARB ‘debentures’) - this is referred to as the ‘Khokha hub’. While the wCBDC is native to the CBDC zone, the system allowed the ‘porting’ of the wCBDC onto the Khokha Hub. While ‘ported’, the wCBDC would be ‘frozen’ in the wCBDC zone while simultaneously an equivalent value of wCBDC was created on the Khokha Hub. The SARB notes that *‘this guarantees a 1:1 mapping and value pegging between the wCBDC tokens in the Khokha Hub and the frozen tokens in the CBDC zone’*.¹⁰

In a sense, while they use the same term ‘wCBDC’ to describe the ported token, one could argue that this is a case of a central bank creating a stablecoin backed by its own wCBDC.¹¹ They refer to both manifestations of the central-bank as wCBDC but, arguably, only wCBDC on its native ledger (the CBDC zone) truly is central bank money.¹² There are perhaps some extra steps in logic to defend ‘wCBDC’ in the Khokha hub as central bank money, even if it is being used as a settlement asset and ‘has the blessing’ of the central bank (this is an example of where the sort of legal issues discussed in section 6.1 appear). Indeed, the pilot documentation thoughtfully discusses this issue at various points - for example:

Porting the wCBDC token to the Khokha hub created a technical challenge, since it created a break between the wCBDC ledger - which in the future could be designated as a settlement system, and the Khokha hub thereby resulting in a split between when technical/operational settlement and legal settlement takes place.

For comprehension, it is perhaps useful to distinguish the CBDC native to the CBDC hub and the ported CBDC. Here, we will refer to them as native CBDC (nCBDC) and ported CBDC (pCBDC), respectively, as opposed to referring to them both as wCBDC (even though that abstraction - treating them as equivalent - is an aim of the pilot, in some sense).

nCBDC is issued on the wCBDC hub and pCBDC is ‘issued’ (though respecting various consistency requirements and connected to nCBDC) on the Khokha hub. nCBDC is unambiguously central bank money in the context of the pilot. pCBDC, however, comes about through a process of connecting the wCBDC hub and the Khokha hub using a ‘software bridge’. This bridge (a collection of APIs and messaging systems operated by the SARB) would, on the initiative of a bank, ‘freeze’ nCBDC in ‘cession wallets’ in the wCBDC zone while creating balances of pCBDC of equivalent value in the Khokha hub.

Rendering the system even more complex is the fact that an explicit stablecoin that isn’t a liability of the central bank, but of the participants, was *also* permitted to be created by the private sector participants in the Khokha hub. This is what the pilot refers to as a wToken (though they apparently sometimes also refer to it as the Khokha Token).¹³ The wToken could be minted by any commercial bank in the ‘Khokha hub’ and was *backed* by reserves

¹⁰Apart from requiring consistency between the transactions involving wCBDC in the Khokha Hub, with balances in the CBDC zone, it was also required that the ported wCBDC only be used by counterparties in the Khokha Hub who also were part of the legacy reserves account systems. That is, ported wCBDC should still adhere to the same terms of use as on its native platform.

¹¹The wCBDC in the CBDC zone is, in turn, backed by reserves in SAMOS

¹²Although, as ever, the terminology is subtle the documentation refers to wCBDC being ‘bought’ by banks using money from their SAMOS (reserve) accounts, which seems distinct from, say a DDR approach (see [subsection 6.1](#)) where the CBDC is manifested simply as claims on reserves.

¹³Indeed, the pilot *also* involved the creation of yet another class of token - a fungible FDM token that built

in the banks' SAMOS settlement accounts. The wToken nevertheless was *redeemable* (distinct from 'backed') for wCBDC on the Khokha hub.

As such the Khokha hub features two settlement assets: pCBDC and the wToken. The CBDC hub features only one: nCBDC. With appropriate legal underpinnings, all of these assets could plausibly satisfy the criteria of being a reliably backed and liquid settlement asset that *conceivably* satisfies Principles of Financial Market Infrastructure (BIS (2012)). Clearly, the biggest leap would be for the wToken to be adopted for large value settlements. Nevertheless the pilot did emphasize ensuring it was *'redeemable by all participants for a predictable, transparent and reliable value to function as an acceptable settlement instrument'* and the *'riskless'* nature of wCBDC was also noted in why it was chosen to be the asset for which wToken could be redeemed.

upon the (non-fungible) SARB debenture token and was traded on a decentralized exchange also established in the pilot. We do not discuss this FDM token in this segment.

D Providers

In this section we briefly describe the platforms offered by technology providers who have been especially prominent in developing solutions for wCBDC (and other DLT applications). Given the immense complexity of the platforms and the fact that the platforms have frequently been updated (and presumably will continue to be so), this summary is inevitably incomplete and should not be regarded as definitive. Extensive communication with providers - and possibly trial use of the platforms - would be needed to make an informed assessment of their features and capabilities.

We focus on the Corda, Hyperledger Fabric and Quorum platforms though note that offerings from [Digital Asset](#) (Blockbuster) and [Elements](#) (Stella 2), among others, have also been used in significant pilots (in parentheses). Furthermore, several pilots have simultaneously used multiple systems to implement the same functional requirements. Indeed, even those pilots that have used a single platform have frequently emphasized that using the platform is not necessarily an endorsement of it as the intended framework for any production system.

Useful discussions of the architecture and capabilities of DLT platforms can be found in [BCB \(2017\)](#), [Ubin₂ \(2017\)](#) and [Stella₂ \(2018\)](#). Again, these references are caveated by the fact that the platforms are constantly evolving, possibly rendering descriptions out of date.

D.1 Corda by R3

[Corda](#), by [R3](#), is a permissioned, ‘blockchain-inspired’ DLT platform that has been designed for enterprise applications in the context of highly regulated markets.¹⁴ A comprehensive reference for the platform is provided in [Sheikh \(2020\)](#).

A standard Corda network consists of business nodes (for example operated by different banks), a notary service, administrative nodes (running network map and identity services) and, optionally, an oracle node(s). Oracle nodes, which are set up to reliably incorporate off-ledger information into transactions, have not yet featured widely in wCBDC pilots (see [Inthanon₂ \(2019\)](#) for an exception) so we set them aside in this summary description of the platform. In fact, we focus on the business and notary nodes.¹⁵

A standard [transaction](#) in Corda updates the ledger in the sense of *consuming input states* that capture the history of transactions and *generating output states* which may then in turn be used as inputs in later transactions. States can only be used once as inputs and, until they are used, the outputs are referred to as ‘unspent transaction outputs’ (UTXOs) - a concept familiar from Bitcoin. For example, UTXOs may capture digital currency holdings that can be aggregated to an overall amount sufficient to execute a proposed ‘asset purchase’ acceptable to a seller - that is the UTXOs capture resources that aggregate to more than the price of the asset.

¹⁴The nebulous term ‘blockchain-inspired’ reflects ongoing (and not especially useful) debate over the precise nature of the system (see discussions [here](#), [here](#) and [here](#), for example). As discussed below, owing to its need-to-know (rather than global broadcast) messaging approach, participants only perceive and store records that relate to their own transactions. The ‘global ledger’ in turn is then *implicit* in the combination of all these individual ledgers, even if it is not stored by any single participant in the standard Corda framework (though see [Jasper₂ \(2017\)](#) for an extension to allow a supervisory node that records the global ledger). In this sense Corda entails a distributed ledger. However, no blocks of transactions are constructed. Instead, the ledgers stored by participants are defined by cryptographically linking each individual transaction to earlier transactions whose outputs (UTXOs) are inputs to the later transaction. The link is to the earlier transactions directly, *via* the UTXO formalism, rather than indirectly through reference to a block including perhaps unrelated transactions.

¹⁵In [Jasper₂ \(2017\)](#) there is also a reference to a ‘supervisory node’ set up to have overarching visibility into all of the individual ledgers - and implicitly transactions - among business participants. This appears to be an extension of the standard Corda framework to help with system oversight and resilience as it ensures that at least one node has recorded the overall state of the global ledger.

The outputs carry information about the transaction that generated them *and* an index that allows each output of the current transaction to be distinguishable from any others (transactions may have multiple outputs). This information is sufficient to construct the ‘backchain’ - that is, the immutable sequence of transactions (operations that consumed particular inputs and generated particular outputs) that yielded the UTXO. This allows nodes to check the provenance of UTXOs referred to by a proposed transaction (known as ‘walking the chain’) and to ascertain if it is contractually valid to consume them as inputs.

Whether or not the transaction is ultimately executed and committed to the counterparties’ ledgers depends on satisfying two types of consensus protocol: validity consensus and ‘uniqueness consensus’. We consider them in turn.

We initially focus on the actions of the business nodes as they can be naturally thought of as ‘counterparties’ in the familiar sense. One node may send another node - or nodes - a proposed transaction, signed by the initial node, but awaiting signatures of the other nodes. The other nodes’ signatures would be provided after they have checked that the transaction is **contractually valid** and - presumably - that they are willing to accept the implications of the transaction from a business perspective. Once all the counterparty nodes have validated and signed the proposed transaction, ‘validity consensus’ is achieved.

At this point the provenance of inputs and various aspects of contractual logic have been checked but the possibility of a double spend remains. As such, before the transactions can be committed to the counterparties’ ledgers, the notary node must check that no inputs would be used more than once. The notary node receives sufficient information to assess whether any inputs in the transaction are subject to a double spend. If all is well, the notary provides the final signature and ‘uniqueness consensus’ is achieved. At this point the transaction is committed to the ledgers of all the involved counterparties, the previously unspent outputs are marked as spent inputs, and a new set of UTXOs will be created, according to the nature of the transaction.

Some other key aspects of Corda are:

- **Need to know communication:** Corda uses a point to point communication model such that only participants in a transaction (and the notary node) are aware of the transaction directly. While any nodes in the network *may* communicate with each other, they will only do so if they are involved in a transaction.¹⁶ Rather than some record of the transaction being broadcast to all nodes in the system, only the parties involved communicate and record the details of that transaction in their ledgers. Consequently, those parties’ ledgers agree on ‘**shared facts**’ related to transactions in which they are both involved (this is part of the consensus protocol discussed above). But their ledgers may - and typically will - differ because they also record other transactions with a disjoint set of counterparties.
- **Notary specifications:** The specification of the notary allows consensus to be achieved quickly. First, the notary need not ‘see’ all aspects of the transaction (only input states need to be checked for double spending) so that the process is relatively lightweight. Second, consensus is attained at the transaction level rather than having to wait for the construction of blocks and their validation. Of course, the question arises as to whether the notary node can process adequate flows of transactions, though this may be enhanced by parallelizing the notary service over a cluster of nodes. Additionally, reliance on a given notary may lead to vulnerabilities. In response, it may be possible to use a service that employs a crash or Byzantine fault-tolerant consensus protocol among multiple notaries (though apparently this is still experimental - see [Sheikh \(2020\)](#) pp. 154).
- **Privacy:** Corda benefits from a baseline level of privacy through its need-to-know communication approach discussed above. Beyond that, cryptographic techniques can be used to further enhance confidentiality. These include **transaction tear-offs** (to hide information from notary and oracle services), **confidential iden-**

¹⁶The network map service permits this ability to reach out directly to another node.

tities (to limit the exposure of information to parties that ‘walk the chain’ in later transactions), and [state reissuance/chain snipping](#) (to limit the length of the backchain to go back only to the issuance of a ‘new’ digital asset). The reason that walking the chain raises privacy issues is that some details of *past* transactions (such as asset ownership) may be revealed to participants in a later transaction for them to verify the provenance of the UTXOs that are supposed to become inputs. Disguising details of transactions, shrouding the identities of the participants and simply shortening the chain may prevent such leakage of information.¹⁷

Finally, we note that Corda appears recently to have established a toolkit called the [Digital Currency Accelerator](#), which appears to gather together various functionalities to ease the implementation of CBDC and stablecoins.

D.2 Hyperledger Fabric

The [Hyperledger Foundation](#) is a group of blockchain developers, focused on promoting cross-industry collaboration on wide variety of technical projects. Among these projects, [Hyperledger Fabric](#) is a permissioned blockchain.¹⁸ A detailed description of the platform can be found in [Androulaki et al. \(2018\)](#) and [Zand et al. \(2021\)](#).

The platform’s success is traditionally thought to stem from its flexibility. It features a ‘modular’ structure that allows different specifications of core elements of the platform to be switched in and out, according to the use case, while retaining the coherence and reliability of the overall system. In particular it supports a models of ‘pluggable consensus’ and ‘pluggable identity management’ to allow case-specific tuning of the system. For example, in a network with a relatively high degree of trust, the consensus mechanism may not need to be Byzantine Fault Tolerant, but only Crash fault tolerant. It also exhibits an open attitude towards supported smart contract languages, again broadening its appeal and enhancing its ability to integrate with other systems.

The Fabric transaction flow follows an ‘execute-order-validate’ pattern. In the execution step, a ‘proposed’ transaction is checked for technical correctness and then endorsed by ‘peers’ according to policies embedded in smart contracts or ‘chaincode’. From this step emerge the transaction results. These results are conveyed to an ‘ordering service’ in which orderer nodes create a block of endorsed transactions according to whatever consensus protocol is being used. The ordering service then delivers the block to peers, who then validate each transaction in the block - adding a flag for whether the validation is satisfied or not - and append the block to the chain. Validation entails two components: first, checking the endorsement policy has been satisfied and, second, checking that no other transactions invalidate it (e.g. by a double spend).¹⁹

This ‘execute-order-validate’ model offers benefits of performance and flexibility, relative to a traditional ‘order-execute’ model as in Ethereum or Quorum, for example. Performance and scalability benefits arise from the fact that only a small number of peers are required to execute the chaincode, leaving the rest of the network’s computational resources to be applied to other steps. Furthermore, since ordering is not yet a concern at the execution stage, they may do this execution in parallel, further enhancing throughput. Importantly, the validation step, executed by a larger set of peers than the endorsement, is much less computationally demanding than execution.

Flexibility is enhanced through the ability to tune consensus policies to a particular step. While the overarching ‘consensus’ required to establish a block (respected by the ordering service) may be implemented in a particular way, consensus among the endorsing nodes in the execute step can be implemented differently. For example, depending on whether there is more inherent trust among orderers or endorsers, one could use more demanding consensus protocols in one step than in the other. Furthermore, there is also additional flexibility on how to implement

¹⁷There may also be performance related reasons for these features, setting aside privacy concerns.

¹⁸Another of the Hyperledger suite of projects, [Hyperledger Besu](#) has recently been involved in [Atom \(2021\)](#) though here we focus on Fabric, given its prominence among CBDC pilots.

¹⁹We use the term ‘peer’ to refer to endorsers and committers. Endorsers are also committers, though not necessarily (or typically - since this would undermine performance) *vice versa*.

confidentiality, with the business logic of chaincode and other transaction details perhaps only being exposed to endorsers, rather than orderers.

The structure also helps enable Fabric's agnosticism with regard to programming languages. In the 'order-execute' case particular languages were required to enforce determinism, so that smart contracts would operate on the current state of the ledger in the same way, and produce the same outputs. In this case, determinism was implicit in the use of a particular type of programming language, from a restricted class of languages. In Fabric's model, non-determinism can be filtered out in the first step as the outputs from execution are being endorsed (or not), before the ordering step begins. The endorsement policy is what enforces agreement across peers executing transaction logic. With this *explicit* agreement over outputs, one need not rely on the *implicit* agreement obtained from determinism-enforcing programming languages (see [here](#) and [here](#) for further discussions on this point). As such, a broader set of languages can be used.

The lightweight nature of the validation step is associated with confidentiality benefits as not all the chaincode logic needs to be shared with validators (as they won't be executing it). Fabric's confidentiality is primarily provided, however, by its use of channels and private data collections. [Channels](#) represent a private sub-network restricted to a particular set of members. The channel has its own distributed ledger. As noted in the Fabric documentation:

[The] isolation of peers and ledger data, by channel, allows network members that require private and confidential transactions to coexist with business competitors and other restricted members, on the same blockchain network.

While some privacy is afforded by the channel structure, additional confidentiality may be required even *within* a channel. As such, [private data collections](#) may be stored off-chain and accessible only by a subset of authorized participants. In particular, the data may not be seen by the ordering service, though an effectively irreversible transformation of the data (a hash) may still be broadcast to all participants within the channel. This allows the execution, ordering and committing steps discussed above to take place, so that the blockchain contains immutable evidence of the data, even if it is not generally interpretable. Hyperledger Fabric may also make use of '[identity mixing](#)' to further enhance the privacy of participants' identities - allowing transactions to be executed anonymously (not just pseudonymously) and in an unlinkable way (so that multiple transactions sent by the same party cannot be identified as all coming from that party).

In the context of wCBDC and tokenized securities, it is worth noting that Fabric has typically been used to implement an account-based approach to characterizing the state of the ledger though it can, in fact, accommodate the UTXO model or some blend of the two (see [Aber \(2020\)](#), for example).

D.3 Quorum by Consensus

Quorum is a blockchain built upon the underlying structure of Ethereum but augmented with additional facilities to allow for permissioning, vote-based consensus and, perhaps most importantly, enhanced privacy. The latter is achieved by allowing for private transactions only interpretable to a subset of parties, even as a cryptographic record of such transactions (a hash) is stored in the global ledger recorded by all nodes in the system.²⁰ The emphasis on privacy reflects the genesis of Quorum being the desire to leverage the Ethereum framework in a way suited to financial services. JP Morgan Chase was the original developer of Quorum though the platform's technology has now been sold to [ConsenSys](#).

The 'Quorum Node' client software is a '[lightweight fork](#)' of the widely used Go Ethereum (geth) client. As such,

²⁰In addition to the private transactions, public transactions are also permitted with the standard Ethereum protocols governing how these are added to the blockchain, though respecting the different consensus and permissioning aspects of Quorum.

Quorum is - and presumably will remain - closely tied to the familiar, well tested and continually updated Ethereum structure. This enhances both reliability and breadth of developer and user appeal. It also enhances compatibility with important tokenization standards. However, owing to the voting-based consensus protocols used in Quorum - particularly the Raft protocol - the throughput of the system can be dramatically higher than on the plain Ethereum framework.²¹

The privacy properties of Quorum are perhaps its main distinguishing feature. Private transactions are encrypted such that interpretable information is sent only to a particular set of nodes who should have knowledge of the transaction data, and who can decrypt the necessary information. While there is a public global ledger, nodes also retain private ledgers containing the details of the private transactions that they have been party to. The private transactions are also manifested on the public ledger but in hashed/encrypted form, so that the global ledger is not a complete record of all details of all transactions, but is nevertheless sufficient to allow validation of blocks by all nodes.

Various wCBDC pilots using Quorum have leveraged Zero Knowledge proofs of various forms in order to implement the privacy properties discussed above, while still allowing validation by nodes not party to all transaction details (*Ubin₂* (2017) and *Khokha₁* (2018), for example). Quorum appears to have been a fairly [early adopter](#) of this privacy enhancing technique though more recent documentation on Quorum appears to make relatively few references to it.

²¹Quorum also permits the use of [sharding](#) for further performance enhancements, whereby each node need only be involved in validate particular subsets of transactions, allowing parallelization and storage benefits.