BANK FOR INTERNATIONAL SETTLEMENTS

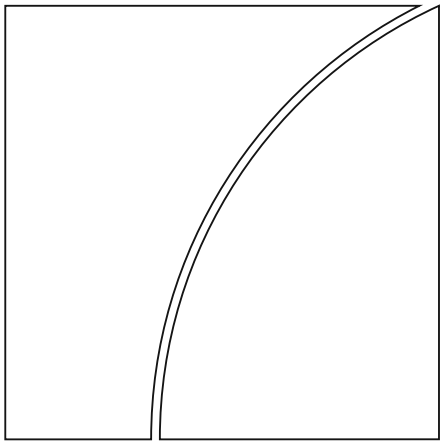

BIS Working Papers

No 948

Central bank digital currency: the quest for minimally invasive technology

by Raphael Auer and Rainer Böhme

Monetary and Economic Department

June 2021

JEL classification: E42, E58, G21, G28.

Keywords: central bank digital currency, CBDC, payments, cash, privacy, distributed systems.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Central bank digital currency: the quest for minimally invasive technology[1]

By Raphael Auer and Rainer Böhme

CBDCs should let central banks provide a universal means of payment for the digital era. At the same time, such currencies must safeguard consumer privacy and maintain the two-tier financial system. We set out the economic and operational requirements for a "minimally invasive" design – one that preserves the private sector's primary role in retail payments and financial intermediation – for CBDCs and discuss the implications for the underlying technology. Developments inspired by popular cryptocurrency systems do not meet these requirements. Instead, cash is the model for CBDC design. Showing particular promise are digital banknotes that run on "intermediated" or "hybrid" CBDC architectures, supported with technology to facilitate record-keeping of direct claims on the central bank by private sector entities. Their economic design should emphasise the use of the CBDC as medium of exchange but needs to limit its appeal as a savings vehicle. In the process, a novel trade-off for central banks emerges: they can operate either a complex technical infrastructure or a complex supervisory regime. There are many ways to proceed, but all require central banks to develop substantial technological expertise.

JEL classification: E42, E58, G21, G28.

Keywords: central bank digital currency, CBDC, payments, cash, privacy, distributed systems.

[1] Bank for International Settlements and University of Innsbruck, respectively. We thank Svetlana Abramova, Claudio Borio, Stijn Claessens, Giulio Cornelli, Sebastian Dörr, Antonio Fatas, Jon Frost, Michael Fröwis, Colin Glass, Patrik Keller, Leonardo Gambacorta, Daniel Woods and an anonymous referee at the BIS working paper series for comments. The views expressed in this article are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

Introduction

Central banks have only recently started to consider issuing a retail digital currency (CBDC) of their own. But the thinking behind it dates back decades. David Chaum set out his vision for anonymous electronic cash 35 years ago (Chaum 1985). Issuance of electronic money by the central bank was also suggested at an early stage (Tobin 1987), although central banks themselves were slow to embrace the concept. Times have since changed, with 86% of central bank respondents to a survey at least researching the topics (Boar and Wehrli 2021), and more than 46 having launched design reports or prototypes (see Auer et al (2020 a,b) and Graph 1). The world's major central banks have joined forces to outline the core principles for issuance, and two retail CBDCs are already in use.[2]

Timeline of central bank activities on CBDC

Graph 1

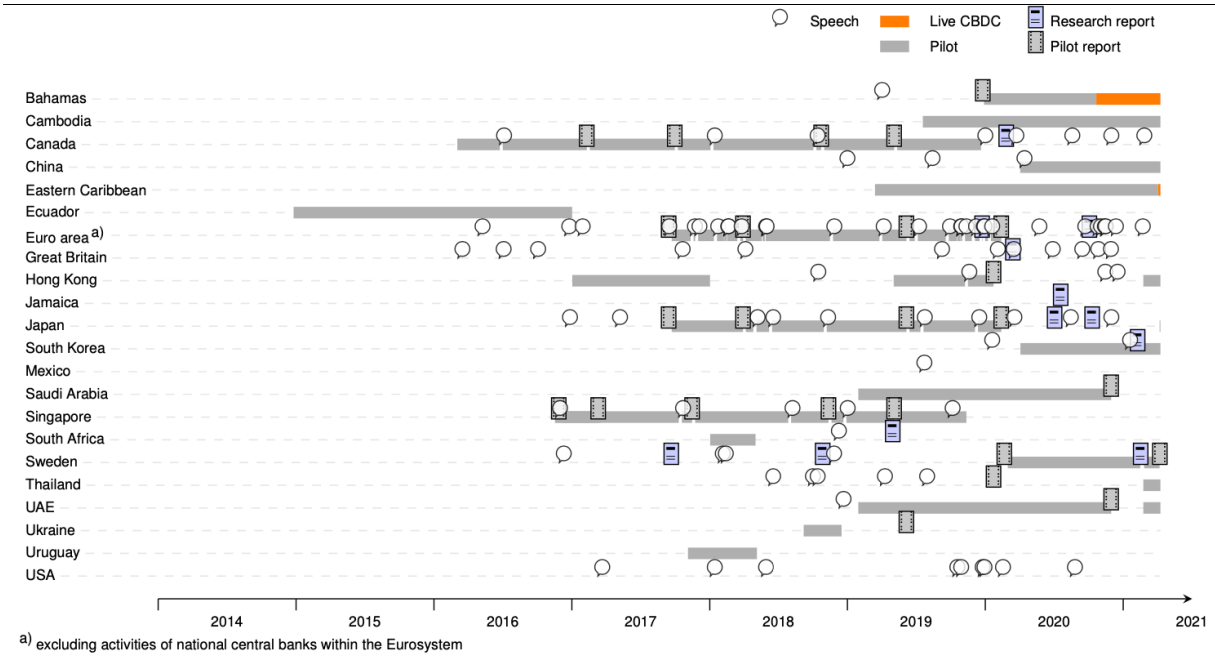

Source: R Auer, G Cornelli and J Frost, "Rise of the central bank digital currencies: drivers, approaches and technologies", *BIS Working Papers*, no 880, August 2020.

Design efforts have to be viewed against the backdrop of central banks' core mandate to provide a resilient and universally accepted means of payment. For centuries, this has been cash. But cash is being used less and less as a means of payment, and the surge of online commerce during the Covid-19 pandemic has accelerated this development (Auer et al (2020c), Alfonso et al (2021)). Should this trend prevail and cash no longer be generally accepted, central banks would have to develop a digital complement, an accessible and resilient means of payment for the digital era.[3]

[2] See Group of central banks (2020), Bank of Bahamas (2020) and Eastern Caribbean Central Bank (2021), respectively.

[3] Promoting inclusion would require low technical access hurdles for the CBDC. Multiple interfaces and physical payment tokens, prepaid CBDC cards or dedicated universal access devices should be considered, together with user-friendly options for children, seniors, and groups with special needs, such as the visually impaired (see Auer, R, J Frost, T Lammer, T Rice and A Wadsworth (2020)).

The key difference between cash and today's electronic retail money is that the latter represents a claim on an intermediary, whereas the former is a direct claim on the central bank. This raises several issues, as the intermediary might run into insolvency, be fraudulent, or suffer technical outages. The collapse of Wirecard and the ensuing impairment of some electronic payment options (Collins (2020), Barba Navetti et al (2020)) foreshadows the importance of these considerations. Looking ahead, a concern is that if the use of cash decreases further, to the point where it loses its universal acceptability, a financial crisis could create havoc by leading to situations in which some financial institutions have to freeze their retail clients' deposits, thus preventing their clients from paying their bills. In Sweden, where cash use has already declined substantially, considerations along these lines have led the central bank to propose to government that digital central bank money held by the general public should also be given the status of legal tender (Sveriges Riksbank (2019)).[4]

At the same time, a CBDC should by no means displace the private sector in financial intermediation or retail payments. The first consideration here concerns the balance sheet. The economic design of a CBDC should not cause a massive reallocation of funds away from commercial banks and to the central bank. While central banks around the world are mandated to provide a universal means of payment (see BIS (2009)), this by no means implies that they should offer savings accounts for the entire economy.

A second consideration – less discussed, but equally relevant – concerns the operational dimension and the efficiency of the payment system. The customer-facing side of retail payments, including onboarding for payment accounts, authorisation, clearing, settlement, dispute resolution, compliance with anti-money laundering (AML) and counter the financing of terrorism (CFT) rules are large operational tasks. These tasks are arguably better handled by the private sector than the central bank (see Borio (2019), Carstens (2019) and BIS (2020)).

In this paper, we hence set out the technical and economic requirements for a "minimally invasive" design – one that upgrades money to the needs of the 21st century without disrupting the tested two-tier architecture of the monetary system. We start by recalling the unique role of cash in today's financial system and highlight the associated economic and operational requirements for a retail CBDC. We then discuss the implications of these requirements for the design of the underlying technology. We argue that there are several options, but a novel trade-off for central banks emerges: a central bank can operate either a complex technical infrastructure or a complex supervisory regime.

The economic design of retail CBDC: cash is the model

To start designing a CBDC, one first has to identify the problems it should solve, as well as what aspects of the monetary system it should preserve. Let us consider these in turn, using the analogy with cash, which already achieves an important balance: it is a secure and useful means of payment, but its use as a savings vehicle is limited. CBDCs could become a digital equivalent.

[4] Other purported advantages of CBDC include lower costs to the consumer than cash or electronic payments, the possibility of encoding novel technological features allowing for programmability, more effective monetary policy implementation, and better privacy. It is noteworthy that not all of these benefits may be realisable, especially not at the same time (see Auer and Böhme (2020a,b)).

On the payments side, cash is unique among all retail payment options, as it is a direct legal claim on the central bank.[5] Everyone can accept cash safely assuming that the received notes and coins will have value in future transactions. Notes and coins are recognised as "legal tender", which typically means that they must be accepted when redeeming debt.

Deposits, by contrast, are legal claims on the respective commercial bank. Bank transfers, check settlements, or debit card charges from A to B merely change one or two commercial banks' promises as to who can withdraw how much cash on request. Every commercial bank backs some of these promises with reserves at the central bank (see Graph 2). This – as well as the bank's equity – increases the depositors' confidence that a surge of withdrawal requests can be fulfilled, but there will always be residual doubt, as such value backing is never full.[6] A commercial bank might run into temporary solvency issues or go bankrupt. In the former case, the payment process might temporarily be interrupted. In the latter case, the claim might not be fully honoured, or even if it is, the legal process to regain funds or compensation from deposit insurance might take time.

Cash, electronic payment instruments, and retail CBDC — Graph 2

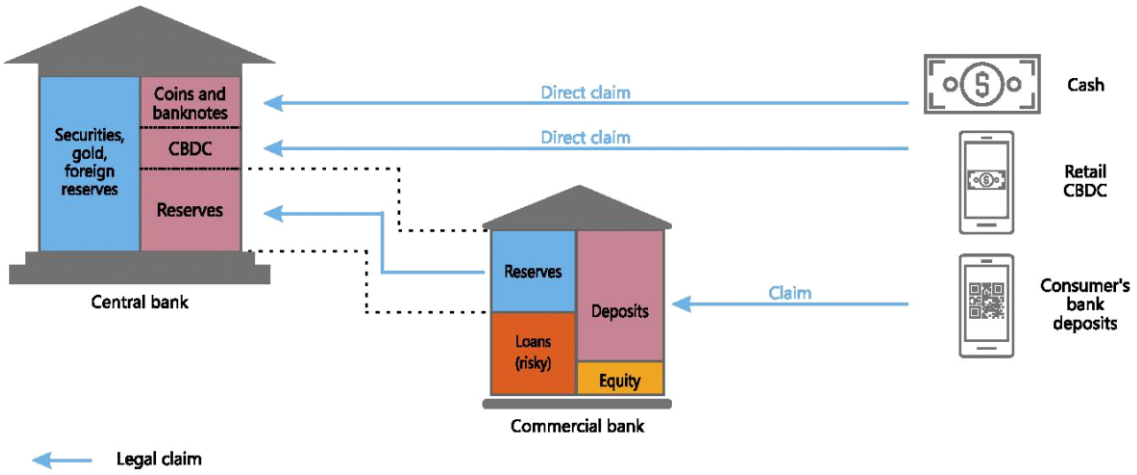

Cash is a direct claim on the central bank, while deposit accounts are claims on the commercial bank. Commercial banks back some of these claims by holding reserves at the central bank and have equity, but the value backing is never full. A CBDC that is unaffected by financial crisis must be a cash-like direct claim on the central bank.

Source: Authors' elaboration.

This showcases the qualitative difference between CBDCs and existing electronic payment instruments. The latter might no longer be accepted in commerce whenever trust in the issuing commercial bank (or other payment service provider) is in doubt. A CBDC, however, would not rely on the soundness of commercial banks, and could thus serve as an anchor for trust, just

[5] See CPSS (2003) for a discussion of the role of central bank money in wholesale payment systems.

[6] The minimum level of funds a commercial bank needs to deposit with the central bank is called the reserve requirement. It is typically equal to a small percentage of customer deposits (and not required in all countries). As in a CBDC, consumers would deposit all of their payment balances directly with the central bank. Hence, from a balance sheet size perspective (but not from other perspectives), a CBDC would be similar to a two-tier system with a 100% minimum reserve requirement.

as cash does today.[7] Therefore, as a legal and economic concept, a CBDC goes far beyond a central bank-operated variant of known electronic payment instruments.

On the other hand, a worry in this respect is that positioning a CBDC as the most secure digital payment instrument could also make it attractive as a savings vehicle. Household investments into a CBDC could substantially increase the balance sheet of central banks, and crowd out deposits at commercial banks. As a result, the business models of commercial banks could be at risk since their source of funds would become more expensive – or dry up altogether. Since commercial banks finance loans with deposits, a CBDC could negatively impact the economy (see Andalfatto (2020), Fernandez-Villarverde et al (2020)).

These considerations underline that cash – despite its usefulness in payments – is of limited appeal as a store of value. This is inherent to physical cash, which carries no interest and is indeed costly to store in large quantities and over long horizons with the danger of damage, loss, or theft. As a consequence, the total outstanding stock of paper currency is moderate, for example USD 5,200 per capita in the US or EUR 3,600 in the euro area. By contrast, private households hold a large share of their wealth in the form of deposits at commercial banks, for example USD 38,000 per capita in the United States and EUR 53,000 in the euro area (see Graph 3).[8]

Cash holdings are small in comparison to deposits at commercial banks [1]

As a percentage of GDP — Graph 3

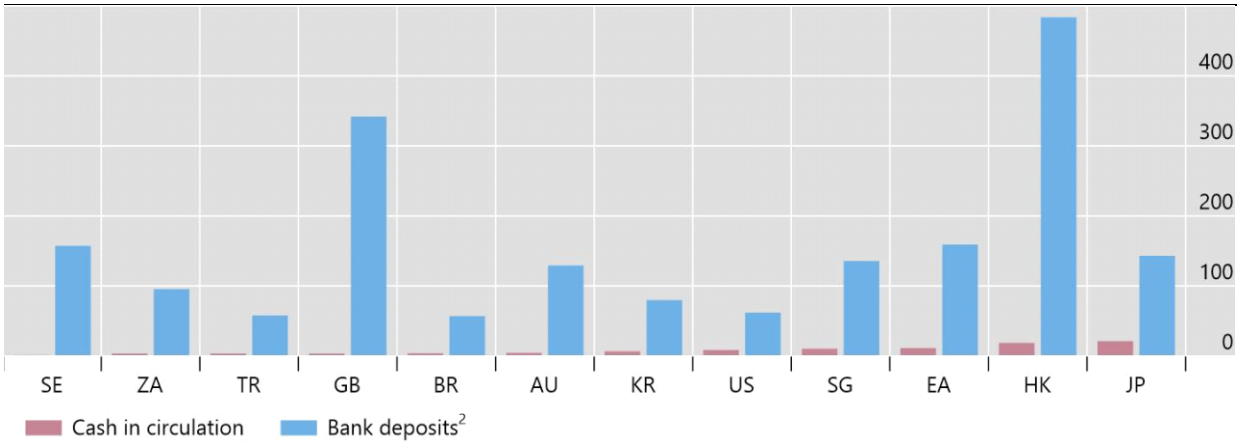

[1] Data for 2019. [2] Closest alternative where data are not available.

Sources: Authors' calculations based on Committee on Payments and Market Infrastructures, *Red Book statistics*, 2019; IMF, *World Economic Outlook*; national data; authors' calculations.

A worry with CBDC design is that consumers could find it too attractive as a store of value: if a substantial share of bank deposits was converted into CBDC, the central bank would receive a massive inflow of funds. In turn, it would have to reinvest these funds, thereby assuming a role as investor it was not set up for.

[7] Of course, cash can lose its value via inflation. However, the key point is that the value of cash is under threat only from inflation, while deposits at commercial banks are exposed to additional insolvency and illiquidity risks.

[8] All mentioned data are end of 2019 values. Currency in circulation is from CPMI Redbook statistics, Deposits from Board of Governors of the Federal Reserve System and ECB, and population counts from IMF World Economic Outlook database.

The current monetary system divides responsibility, with central banks stabilising the core and commercial banks conducting all consumer-facing activities. This two-tiered architecture, depicted in Graph 2, has evolved for good reasons and from a long history of institutional arrangements. Maintaining this convention – in particular, the value of money – depends on sound governance. Its absence has been causally associated with repeated episodes of debasement and economic depression (ie Shin and Schnabel 2018 and Frost et al. 2020). This experience shows that the value of money is best safeguarded by an institution that is accountable to the public rather than to private investors, and focused on its primary task of stabilising the monetary system.[9]

The rationale for the private sector's involvement is the commitment to market-based solutions: good investment decisions often require specific local knowledge, while efficient provision of services requires open and competitive markets. One aspect is that a bank that invests by making loans must know or be able to estimate creditors' solvency to price the associated risk. Public sector institutions might not have this knowledge to the degree that local and specialised private investors do – this is the core hypothesis of Hayek's (1945) case for free markets. Another aspect is that competitive markets are also contestable, and allow many firms to compete, improving economic efficiency. A corollary of the free market idea is that the customer-facing side of money and finance, as well as continued innovation in this field, is best left to the private sector.[10]

As a consequence, the economic design of a CBDC should hence allow commercial banks to keep their intermediation role between savers and investors. An academic literature has focused on how to recycle received funds back into the private sector during calm periods (see Andalfatto (2020), Niepelt (2020), and Fernández-Villaverde et al (2020)) and market turmoil (see Brunnermeier and Niepelt (2019)).

Policymakers have, however, favoured solutions that keep the amount of outstanding CBDC small to begin with (see Carstens (2021 a and b), BIS (2020), and Group of Central Banks (2020)). One option is to remunerate holdings at a zero interest rate, or at least one that is lower than that on commercial bank deposits.[11] Another option is to cap holdings per household at some maximum amount, for example EUR 3,000 per citizen has been mentioned in the context of the euro area (see Lock (2021)). Bindseil (2019) proposes a tiered approach in which an unfavourable interest rate is applied to holdings exceeding a defined threshold.

[9] See ie Carstens (2019). Indeed, in many countries, issuing and safeguarding the value of cash is the reason the central bank was founded. A case in point is the US Federal Reserve System, which was founded in 1913 after a period of privately issued currencies that were often debased (see Giannini (2011) Friedman and Schwartz (1963)).

[10] Free market economies barely exist in pristine form. State-run development banks or „unconventional" monetary policy measures of buying government bonds and other financial assets on a secondary market are examples of risk-taking by the public sector.

[11] Some scholars, such as Bordo and Levine (2017) advocate the introduction of CBDC so that central banks can implement negative interest rate policies more effectively. Central Banks are more cautious in this regard and have come to view the monetary policy implications as second order (Group of Central Banks (2020)). Andolfatto (2017) and Bank of Canada (2019) study how commercial banks will adapt the remuneration of their deposits when the central bank introduces CDBC.

However, in this discussion, the interaction of technology and systemic implications is often overlooked: a fully anonymous CBDC[12] does not allow person-specific limits to be applied, ie caps come at the cost of forgoing anonymity.[13]

A second consideration for a minimally invasive CBDC design regards the operational dimension and the efficiency of the payment system. As the customer-facing side of retail payments is arguably better handled by the private sector than the central bank, the underlying question is how an operational architecture can balance direct claims on the central bank with the operational involvement of private sector payments services providers (PSP).

Yet again, CBDC design efforts should aim for "cash-likeness" in a number of dimensions. The digital record would be tamper-proof and imply a direct claim on the central bank. It could still be exchanged when intermediaries face technical outages or financial stress. The CBDC would be designed to keep the operational burden of the central bank low, and give the private sector an important role in all customer-facing activities, just as in today's system.

From requirements to operational design: CBDC architectures

How do these requirements for a CBDC affect the technical infrastructure? In technical terms, every form of digital money requires a distributed record-keeping system. Here, distributed means that it is implemented in many different places, eg the merchants' terminals as well as consumer devices carrying record are components of this system (see eg Lamport (1978)). This record-keeping system updates a shared state, which encodes how a given number of currency units is allocated to holders.[14] Technical communication ensures that every component of the distributed system is up-to-date at least with the part of the shared state relevant to this component. For example, a consumer's digital wallet needs to know its current balance, and arguably should not know any other wallet's balance. In the case of CBDC, laws must be created to ensure that information encoded in this state is mapped to ownership of claims against the central bank.

The technical architecture of a CBDC is defined by the role of the components of the distributed record-keeping system, their communication relations, and the question who is in control of each component. Graph 4 gives an overview of possible architectures for CBDCs and the alternative of a narrow payments bank, building on our previous work in Auer and Böhme (2020a). These examples differ in terms of the structure of legal claims and the record kept by the central bank.[15] We discuss the extreme options first, which either realise a direct claim or involve the private sector, before moving on to the more promising options that combine both.

---

[12] In economic jargon – and much to the confusion of computer scientist – a fully anonymous design is termed "token-based." This definition of token versus accounts follows Kahn and Roberds (2009), and as put by Kahn (2016) implies: "*In a token-based system, the thing that must be identified for the payee to be satisfied with the validity of the payment is the "thing" being transferred – "is this thing counterfeit or legitimate?" In an account-based system, however, the identification is of the customer – "Is this person who she says she is? Does she really have an account with us?". What remains still ambiguous is whether the "thing" can be a bit string or must be physical.*"

[13] A combination of caps and interest rate polices is the "tiered" approach with CBDC balances below a given level earning a positive interest rate and balances above that level earning a lower interest rate, or even a penalty negative interest rate (Bindseil (2020)). The central bank of the Bahamas has adopted this design in the "Sand Dollar" CBDC.

[14] State machines are a key concept in computer science. This terminology has previously been used in the economic literature, for instance by Rubinstein (1986) to tract the study of repeated games.

[15] All four architectures could be either account- or token-based, and might run on various infrastructures. These choices are discussed in Auer and Böhme (2020a).

Retail CBDC architectures and fully backed alternatives

Graph 4

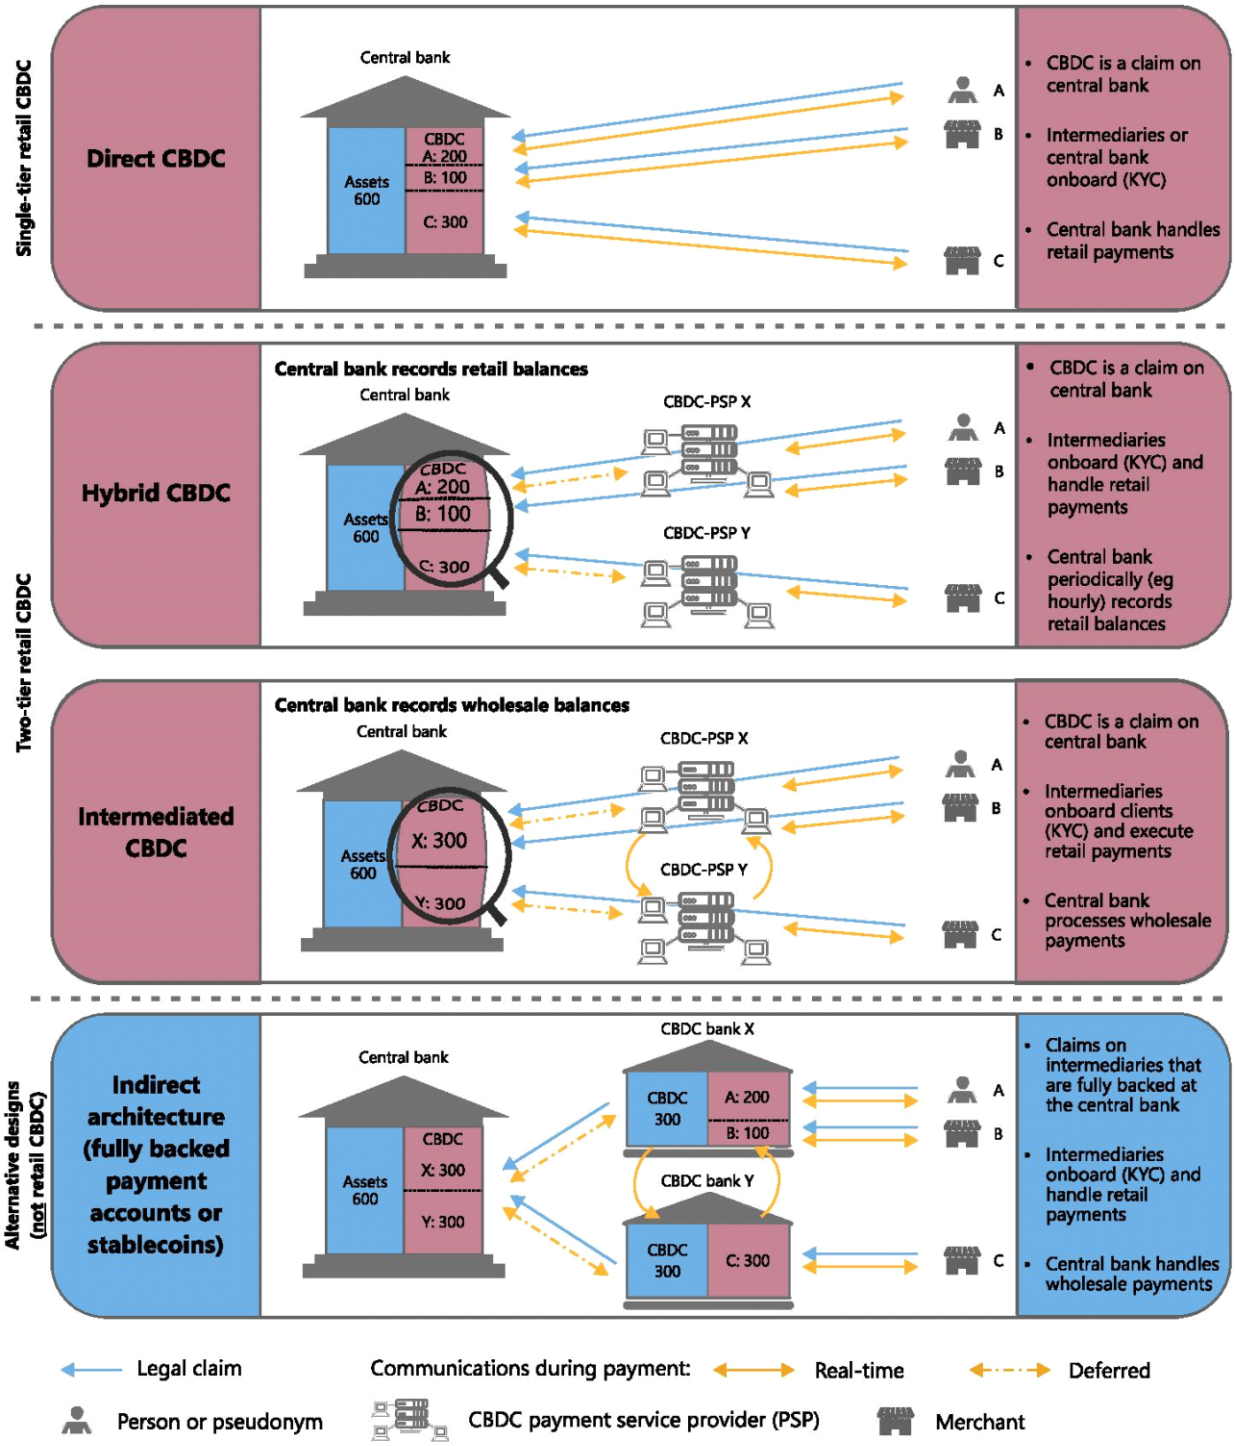

In the "Direct CBDC" model (top panel), the CBDC is a direct claim on the central bank, which also handles all payments in real time and thus keeps a record of all retail holdings. Hybrid CBDC architectures incorporate a two-tier structure with direct claims on the central bank while real-time payments are handled by intermediaries. Several variants of the hybrid architecture can be envisioned. The central bank could either retain a copy of all retail CBDC holdings (second panel), or only run a wholesale ledger (third panel). In the indirect architecture (bottom panel), a CBDC is issued and redeemed only by the central bank, but this is done indirectly to intermediaries. Intermediaries, in turn, issue a claim to consumers. The intermediary is required to fully back each claim with a CBDC holding at the central bank. The central bank operates the wholesale payment system only.

Sources: elaboration based on R Auer and R Böhme, "The technology of retail central bank digital currency", *BIS Quarterly Review*, March 2020, pp 85–100.

In one possible CBDC architecture, the central bank would run this system itself, handling all payments, and directly updating the state after each transaction. Even if a central bank were to build the necessary infrastructure, the resulting CBDC might be less attractive to consumers than today's retail payment systems. This is because real-world payment systems must deal with connectivity outages or offline payments. The CAP theorem in computer science (Gilbert and Lynch 2002) tells us that no distributed system can be available and consistent while parts of it are disconnected. Many existing electronic payment systems work around this technical impasse with an economic solution. They involve intermediaries, such as credit card networks, who take on financial risk resulting from potential inconsistencies in the state update process, and charge fees for this service. If the central bank would run the system itself, it would have to engage in this risk taking. However, this can be perceived as incompatible with the aforementioned concept of a reserving tasks in the monetary system that require local knowledge to the private sector.[16]

Another concern with this "Direct CBDC" architecture is that it might marginalise private sector involvement. An emerging payments technology sector innovates with value-added services, such as automated financial advice, integration with consumer platforms, and connection to other financial products like consumer credit. It is unlikely that the central bank could substitute for the private sector in all these activities.

In contrast, consider an alternative to issuing retail CBDC: the simple requirement to fully back payment accounts with reserves at the central bank (bottom panel of Graph 2). This proposal has been floated under many names; it could be considered a narrow payment bank, or even a "rigid stablecoin" that is backed 100% by reserves at the central bank. Adrian and Mancini-Griffoli (2019) term this architecture "synthetic CBDC," though some central banks have argued that it does not warrant the "CBDC" label (ie Group of central banks (2020)). Here, we follow Auer and Böhme (2020a) and term it the "indirect" architecture.

This model's regulatory and supervisory issues, as well as those pertaining to deposit insurance, are different from those of a CBDC with direct claims. If the intermediary goes bankrupt, determining the legitimate owner might involve lengthy and costly legal processes with uncertain outcomes. Whereas full backing would likely mean that such episodes occur infrequently, the recent example of Wirecard underlines that these concerns have to be taken seriously nevertheless.

On balance, we argue that the most interesting design space combines the credibility of a direct claim on the central bank with the convenience of private sector payment services. One possible architecture is called "Hybrid CBDC" and was described in Auer and Böhme (2020a). A key element is the legal framework that underpins direct claims on the central bank – ie the CBDC is never on the balance sheet of the payment service providers (PSPs) and thus unaffected by bankruptcy. This way, in the event of PSP insolvency, consumers' CBDC holdings would not be exposed to claims by the PSP's creditors.

[16] While exceptions for the sake of relaxing a CBDC's consistency requirements appear conceivable, those would require an expansion of central bank operations, requiring a broader political debate and almost certainly lead to a more invasive solution than envisioned here.

Technological resilience is achieved in the hybrid design via the central bank operating a backup infrastructure (hence the name hybrid – a payment system that can run on either a public or a private infrastructure), as depicted in the upper middle panel of Graph 4. If a PSP fails – financially or technically – there must be a way for the central bank to unambiguously honour claims and, ideally, resume payments for the failing PSP's customers without much delay. This capability depends on the information about retail accounts available to the central bank in such an event. While the central bank does not operate retail payments, it maintains a backup copy of balances which allows it to restart payment should intermediaries run into insolvency or face technical outages. The central bank's technical capability to restore retail balances could be achieved by keeping digitally signed payment confirmations with the intermediaries, the retail account holders themselves, and maybe at a lower frequency with the central bank itself. As digital signatures are non-repudiable,[17] the central bank can to honour digitally signed claims, regardless in which of many places the record was held – including under the matrass (Abramova et al. 2021). This can serve as a credible anchor of resilience.

We note that some central banks may shy away from running a record of all retail data, for example due to the issues with privacy and data security (see ie Powel (2019)). Such central banks might consider an "intermediated" CBDC architecture, in which the central bank records wholesale balances only.[18] An advantage of having less payments data at the central bank is that it becomes less exposed to malicious attacks than in a hybrid (or a direct) architecture. This reduces the risk and impact of data breaches at the central bank. However, the downside of the intermediated CBDC architecture is that the central bank needs to honour claims that it has no record of. It would thus have to rely on the integrity and availability of records kept by third parties. Consequently, to safeguard cash-like credibility, PSPs would need close supervision capable to ensure at all times that the wholesale holdings they communicate to the central bank indeed add up to the sum of all retail accounts. Some of these concerns can be addressed with cryptography, albeit at the cost of higher technical complexity and lower speed when processing payments.[19] This is an area of ongoing technical research, primarily for the use in distributed ledgers, which we have argued are not necessarily the best choice to set up a CBDC infrastructure. It remains to be seen which new results in cryptography withstand the test of time and can be applied fruitfully in architectures suitable for CBDC.

In summary, both "Hybrid" and "Intermediated CBDC" architectures would have better financial resilience than fully backed payment accounts. These options also seem simpler to operate for the central bank than a "Direct CBDC". As the central bank does not directly interact with retail users, it can concentrate on a limited number of core responsibilities, while competing intermediaries handle the operation. Technically, many different infrastructures can

[17] Non-repudiation is a technical property of digital signatures. It means that the validity of a signature can be verified without the cooperation of the signing party.

[18] Such an architecture is referred to as "Decentralized solutions with intermediaries" by Armelius et al (2020). The central bank would not maintain the full record of retail transactions. This could release the central bank from the need to oversee anti-money laundering/combating the financing of terrorism (AML/CFT) risks and also reduce some operational risks for the central bank (such as the attractiveness of the central bank as a target for cyber attacks).

[19] For example, PSPs could commit to encrypted account balances at the central bank and give users payment confirmations with proofs-of-inclusion. This record can be checked against some values published by the central bank. If the balance is not included (and hence accounts at the PSP do not add up), this could be detected and the PSP held accountable.

support this division of responsibilities depending on the desired level of resilience, but central banks need to be aware of the implied technological requirements for their operational setup.

The trade-off between supervisory and operational complexity

The above discussion of the intermediated and hybrid CBDC architectures exemplifies an underlying trade-off for the central bank: it can maintain record-keeping directly, or outsource it to private sector and supervise it. At a technical level, to establish credible direct claims in electronic form, it is sufficient if the central bank has a view (ie read access) to an authoritative source of information. The record-keeping can be delegated to the private sector, who might either be allowed to use proprietary technology or be required to run some open protocols specified by a standardisation body.[20]

Graph 5 illustrates this trade-off conceptually by bringing the high-level options for the information set of the central bank in a complete order (left scale). The right scale (in red) shows how complementary supervisory requirements grow as the integrity of the system, in particular the ability to honour claims, increasingly relies on the availability of consistent and authentic information from private entities.

A new trade-off for the central bank of the future — Graph 5

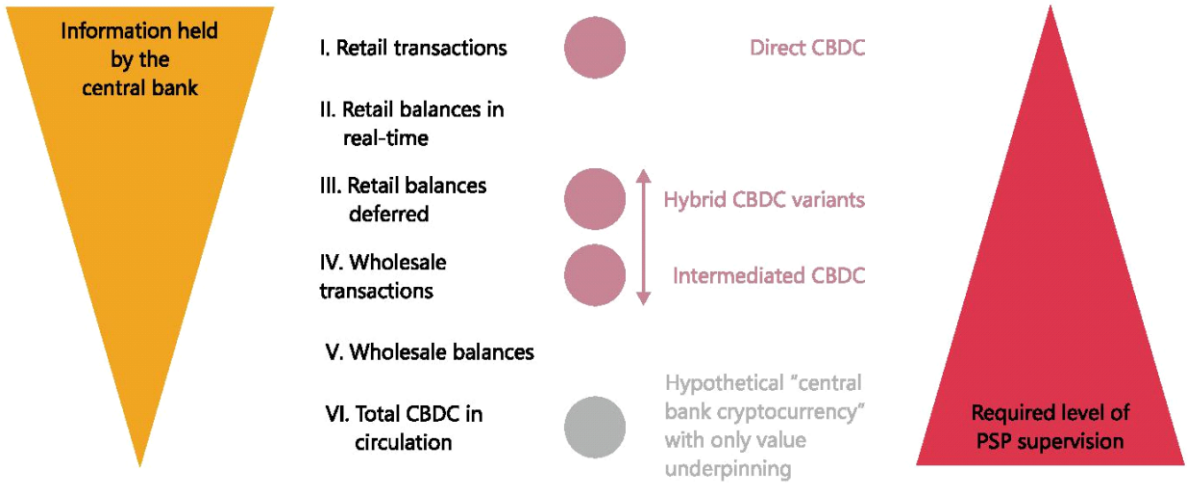

The orange scale on the left visualises the central bank's information set from complete (top) to a single number (bottom). A view on transactions gives the central bank access to the graph of money flows, which are most sensitive in terms of privacy but also most useful for crime prevention. By contrast, the view of balances does not necessarily reveal payment relations. In particular when the balance updates are deferred to the end of a period (eg hour or day), most individual payments are unrecoverable from the mere differences in balances. Balance information is sufficient and necessary to honour claims. The less information is shared with the central bank, supervision is needed to ensure that the relevant information can be retrieved from the private-sector entities if needed. This is visualised by the red scale on the right. The red circles indicate which level of information corresponds to the technical CBDC architectures outlined in the paper (cf. Graph 4). The grey circle illustrates the hypothetical point in design space where the central bank's task is reduced to controlling the monetary supply.

Source: Authors' elaboration.

[20] The decision between proprietary and open standards applies to every architecture. It may affect the competitiveness of the industry, interoperability between CBDCs, and consumers' trust in the CBDC.

Importantly, this supervision is different to conventional banking supervision, which chiefly concerns the integrity of accounting. Recall that the hybrid architecture keeps PSP's balance sheets out of the loop, but the information in their record-keeping systems is crucial. Therefore, the supervision must be focussed on more technical aspects and happen more frequently – perhaps in real time – than currently practiced in banking supervision. It must put aspects like integrity, consistency, information security, and privacy centre stage.[21]

The need for such technical supervision of private sector entities emerges as soon as the central bank is shielded from some retail balances, which a fraudulent or technically compromised PSP could use to appropriate customer funds. On the other end, maximum supervision is required when the central bank has the conceivable minimum information set. This is when its role is reduced to money creation and destruction, hence the only information available to the central bank is the amount of CBDC in circulation. Such a variant would let the central bank focus on its core mandate, monetary policy, and leave all payments operations to the private sector under tight supervision.

Implementing CBDC: cryptocurrencies are not the model

The above discussions of operational architectures and the underlying trade-off between operational and supervisory complexity give rise to the question of adequate technological implementations.

Many approaches proposed by the industry (intentionally not depicted in Graph 2) envision payment systems that feature intermediaries but seek to reduce dependence on them. For example, a number of CBDC prototypes are built on enterprise versions of distributed ledgers, such as Corda, Hyperledger, or Quorum (see Auer et al. 2020a). These versatile software packages were inspired by and borrow concepts from decentralised cryptocurrencies.[22] However, most central bank projects have good reasons for running them in configurations that resemble a redundant but centrally controlled database rather than Bitcoin. The use of ambiguous "blockchain terminology" often obscures this fact.

Some academic proposals,[23] by contrast, are designed to break with conventional databases. They adapt selected principles of decentralised cryptocurrencies to the use case of CBDC. In doing so, they often defend against the wrong threat, namely unaccountable and potentially malicious intermediaries. Those indeed pose challenges to cryptocurrencies running on many computers whose owners are barely identifiable, let alone known and trusted. For CBDC, however, it is unimaginable that a central bank would allow unidentified or unvetted parties to manage critical records. If a CBDC architecture uses designated intermediaries, they would be composed of licensed and supervised banks, established payment service providers, or technology companies if they undergo supervision. Hence, we deem it sufficient for the record-keeping system to ensure that outright malicious actions can be detected and that recovery is smooth. Trying to prevent this failure mode altogether is not an efficient use of resources.

[21] It is conceivable, albeit not near-term, that part of this supervision is carried out by peer control on the basis of public wholesale ledgers and involving digitally signed payments confirmations in the hands of customers (see Auer (2019)).

[22] For academic proposals featuring such permissioned distributed ledger technology (DLT), see Ali and Narula (2020) and Gersbach and Wattenhofer (2020). Auer, Monnet and Shin (2021) examine applications of permissioned DLT in the broader field of finance. They find that that due to the difficulty in ensuring that the underlying economic design sets the right incentives, in many instances centralised designs may be economically more efficient.

[23] See for instance Danezis and Meiklejohn (2016) and Sun et al. (2017).

Another dimension of the implementation concerns the end-user devices enabling access to CBDC. They must also be usable for less technically adept users. Since CBDCs might have advantages for the unbanked and the elderly population, this desideratum cannot be overstated. Many technical proposals adopt Bitcoin's approach of authorising transactions via digital signatures alone. As a result, the security of assets hinges on the secrecy of private keys. And if 20 years of research in usable security teaches us a single lesson, then it is that "Johnny can't encrypt" (Whitten and Tygar, 1999); precisely because end users cannot manage private keys! Given that proficient cryptocurrency users keep losing fortunes due to lost and stolen keys (eg Abramova et al. 2021), there is simply no case for making people's direct claims on the central bank – their money under the mattress – contingent on the use of cryptography without any safety net.[24] Also in this regard, drawing inspiration from Bitcoin, a technology designed to circumvent authority (e.g. Böhme et al. 2015), is clearly not the best blueprint for a public good provided by a central authority.

Stepping up: payments and the central bank of the future

The declining usability of physical cash has led to a growing number of central banks to consider the issuance of a cash-like electronic claim on the central bank that is also available to households, ie a retail CBDC. In this paper, we have argued that replicating all the convenient properties of cash is no easy task. We have attempted to set out the requirements for a CBDC-based payment system, and explored suitable operational architectures and technological implementations.

A CBDC should let central banks provide a universal means of payment for the digital era, while at the same time safeguarding consumer privacy and preserving the private sector's primary role in retail payments and financial intermediation. We have thus argued that it should embody a "minimally invasive" design that achieves the stated goal of offering a cash-like digital payment option without upending the monetary and financial system.

On the technological side, we argue that two-tiered CBDC designs with private sector intermediaries handling retail payments present a viable option. However, a range of different operational arrangements is conceivable. In some, the central bank hosts a database of retail balances (even if anonymised), whereas in others it keeps track only of wholesale balances.

Within this design space, a novel trade-off for central banks emerges: they can operate either a complex technical infrastructure or a complex supervisory regime. There are many ways to pursue either option, but all will require central banks to deepen their technological expertise.

And there are several adjacent discussions that we have not even touched upon. One is the trade-off between accessibility and security from attacks on electronic payment devices. Another is that no CBDC design can be entirely cash-like, as users must rely on a technical infrastructure and may need to check for intermediary failure or at least respond to notifications.

[24] Popular messaging apps, which are celebrated for deploying cryptography at scale, do so by relying on big-tech identity providers for account recovery, and by trading off security for usability. For perspective, the fraction of users understanding and securely using WhatsApps' public key verification is orders below those who grasp basic banknote security features.

Third – and maybe paramount – is the issue of how electronic money could be designed to offer privacy, a property cash offers by default. In any technology for electronic money, privacy is a feature that has to be laboriously engineered rather than being an intrinsic characteristic of the record-keeping system.

All this brings us back to David Chaum, who proposed from a computer science perspective how some of these issues, chiefly anonymity, could be tackled.[25] Thirty-six years later, the time finally seems ripe for computer scientists and central bankers to join forces and bring to light a true electronic complement to cash, as a basic means of payment for the digital era.

Looking ahead, the computer science community will need to accompany the development of CBDC and communicate realistically what current technology can and cannot deliver.[26] As the technology for replicating all properties of cash does not yet exist, the transition to CBDC might transform today's monetary system, which builds on cash as a pivotal element.[27] Given society's high dependence on a functioning and predictable monetary system, we call for a minimally invasive technology: one that offers consumers a digital complement to cash while preserving the two-tier monetary system and the important role of the private sector in it.

References

Abramova, S, Voskobojinikov, A, Beznosov, K and R Böhme (2021): "Bits under the mattress: understanding different risk perceptions and security behaviors of crypto-asset users", forthcoming in *Proc of ACM CHI*.

Adrian, T and T Mancini-Griffoli (2019): "The rise of digital money", IMF Note, no 19/001, July.

Ali, R and N Narula (2020): "Redesigning digital money: what can we learn from a decade of cryptocurrencies?", MIT DCI Working Papers, January.

Agarwal, R and M Kimball (2015): "Breaking through the zero lower bound", IMF Working Papers, no WP/15/224.

Alfonso, V, C Boar, J Frost, L Gambacorta and J Liu, "E-commerce in the pandemic and beyond", BIS Bulletin, no 36, January 2021.

Andolfatto, D "Assessing the impact of central bank digital currency on private banks", The Economic Journal 131 (634), 525-540.

Armelius, H, G Guibourg, S Johansson and J Schmalholz (2020): "E-krona design models: pros, cons and trade-offs", *Sveriges Riksbank Economic Review*, June, pp 80–96.

Arner, D, R Auer and J Frost (2020b): "Stablecoins: risks, potential and regulation", Bank of Spain Financial Stability Review, issue 39.

Auer, R and R Böhme (2020a): "The technology of retail central bank digital currency", *BIS Quarterly Review*, March, pp 85-100.

Auer, R and R Böhme (2020b): "CBDC architectures, the financial system, and the central bank of the future", VoxEU, 29 October.

[25] An implementation of these ideas is discussed in Chaum et al (2021).

[26] Clearly, deriving a workable direct claim from an electronic record requires a legal framework, as would the option to make CBDC legal tender. Hence, engineering and legislation will have to work hand in hand.

[27] Pichler et al. (2020) discuss the specific role of cash for the monetary system. Committee on Payments and Market Infrastructures and Markets Committee (2018) discusses the broader systemic implications of CBDC issuance. These include the potential disintermediation of banks, heightened susceptibly to bank runs during financial turmoil, the speed and extent of monetary policy transmission, and the possibility to implement negative interest rates on CBDC.

Auer, R, G Cornelli and J Frost (2020a): "Rise of the central bank digital currencies: drivers, approaches and technologies", CEPR Discussion Paper 15363, October.

Auer, R, G Cornelli and J Frost (2020b): "Covid-19, cash and the future of payments", *BIS Bulletin*, no 3, April.

Auer, R, J Frost, T Lammer, T Rice and A Wadsworth (2020c): "Inclusive payments for the post-pandemic world", *SUERF Policy Note* no 193.

Auer, R, C Monnet and H S Shin "Permissioned distributed ledgers and the governance of money" BIS Working Papers No 924, 27 January 2021.

Bank for International Settlements (BIS) (2020), "Central banks and payments in the digital era", BIS Annual Economic Report, Ch. III, June.

Bank for International Settlements (2009), "Roles and objectives of modern central banks", Chapter 2 in "Issues in the Governance of Central Banks", 18 May.

Barba Navaretti, G, G Calzolari and A F Pozzolo (2020), "What are the wider supervisory implications of the Wirecard case?", European Parliament Think Tank, PE 651.384

Barontini, C and H Holden (2019): "Proceeding with caution – a survey on central bank digital currency", BIS Papers, no 101, January

Bindseil, U (2020): "Tiered CBDC and the financial system", *ECB Working Paper Series*, no 2351.

Boar, C, H Holden and A Wadsworth (2020): "Impending arrival – a sequel to the survey on central bank digital currency", *BIS Papers*, no 107, January.

Bordo, M. D. and A T Levin (2017) "Central Bank Digital Currency and the Future of Monetary Policy", NBER Working Paper 23711, August.

Borio, C. (2019) "On Money, Debt, Trust and Central Banking"; BIS Working Paper No. 763; February.

Brunnermeier, M and D Niepelt (2019): "On the equivalence of private and public money", *Journal of Monetary Economics*, vol 106, pp 27–41.

Carstens, A (2020): "Shaping the future of payments", *BIS Quarterly Review*, March, pp. 17–20.

Carstens, A (2021a): "Digital currencies and the future of the monetary system", speech at the Hoover Institution, 27 January.

Carstens, A (2021b): "Central bank digital currencies: putting a big idea into practice", speech at the Peterson Institute for International Economics, 31 March.

Central Bank of the Bahamas (CBB) (2019): "Project Sand Dollar: A Bahamas Payments System Modernisation Initiative", 24 December.

Chaum, D (1985): "Security without identification: transaction systems to make big brother obsolete", Communications of the ACM, vol 28, no 10, pp 1030–1044.

Chaum, D C Grothoff and T Moser, 2021. "How to issue a central bank digital currency," Working Papers 2021-03, Swiss National Bank.

Committee on Payments and Market Infrastructures and Markets Committee "Central bank digital currencies" March 2018.

Committee on Payment and Settlement Systems (2003), "The role of central bank money in payment systems" August.

Collins, B. (2020) "Wirecard collapse freezes millions of online bank accounts: will customers ever get their money back?," Forbes.com, Jun 28.

Danezis, G and S Meiklejohn (2016): "Centrally banked cryptocurrencies", proceedings of the 23rd Annual Network and Distributed System Security Symposium, The Internet Society.

Eastern Caribbean Central Bank (2021) "Public Roll-out of the Eastern Caribbean Central Bank's Digital Currency – DCash!", Press Release, March 25.

Friedman, M and A Schwartz (1963): *A Monetary History of the United States, 1867–1960*, Princeton: Princeton University Press.

Frost, J, Shin HS and Wierts P (2020): "An early stablecoin? The Bank of Amsterdam and the governance of money", BIS working paper, forthcoming.

Fernández-Villaverde, J, D Sanches, L Schilling, and H Uhlig (2020): "Central bank digital currency: central banking for all?", University of Chicago, Becker Friedman Institute for Economics Working Paper no 2020-04.

Fung, B. and H. Halaburda (2016): "Central Bank Digital Currencies: A Framework for Assessing Why and How." Bank of Canada Staff Discussion Paper No. 2016-22.

Gersbach H. and R Wattenhofer, 2020. "A Minting Mold for the eFranc: A Policy Paper," CER-ETH Economics working paper series 20/342, Center of Economic Research.

Group of Central Banks (2020): "Central bank digital currencies: foundational principles and core features", Report No 1 in a series of collaborations from a group of central banks, October.

Giannini, C (2011): The age of central banks, Edward Elgar.

Gilbert, S and N Lynch (2002): "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services", ACM SIGACT News, vol 33, no 2, pp 51-59.

Hayek, F (1945): "The use of knowledge in society", American Economic Review, vol 35, no 4, pp 519-30.

Kahn, C (2016): "How are payment accounts special? Payments Innovation Symposium, Federal Reserve Bank of Chicago

Kahn, C and W Roberds (2009): "Why pay? An introduction to payments economics", Journal of Financial Intermediation, vol 18, no 1, pp 1–23.

Kiff, J, J Alwazir, S Davidovic, A Farias, A Khan, T Khiaonarong, M Malaika, H Monroe, N Sugimoto, H Tourpe, and P Zhou (2020): "A survey of research on retail central bank digital currency", IMF Working Paper no 20/104.

Lamport, L. (1978) "Time, Clocks, and the Ordering of Events in a Distributed System". Communications of the ACM, vol 21, no 7, 558–565.

Look C., "ECB's Panetta Floats 3,000-Euro Limit on Digital Cash", Bloomberg, 21 February 2021.

Niepelt, D 2020. "Monetary Policy with Reserves and CBDC: Optimality, Equivalence, and Politics," Working Papers 20.05, Swiss National Bank, Study Center Gerzensee.

Pichler, P, M Summer and B Weber (2020): "Does digitalization require central bank digital currencies for public use?", Monetary Policy & the Economy, Oesterreichische Nationalbank, issue Q4/19, pp 40-56.

Powell, J. (2019) "Letter to Congressman French Hill", November 19.

Rubinstein, A. (1986) "Finite automata play the repeated prisoner's dilemma", Journal of Economic Theory, vol 39, pp 83–96.

Schnabel I and Shin HS (2018): "Money and trust: lessons from the 1620s for money in the digital age" BIS Working Papers, No 698.

Sun, H, H Mao, X Bai, Z Chen, K Hu and W Yu (2017): "Multi-blockchain model for Central Bank Digital Currency," proceedings of the 18th International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT), IEEE, pp 360-367.

Sveriges Riksbank (2019): "The Riksbank proposes a review of the concept of legal tender", press announcement.

Sveriges Riksbank (2020): "The Riksbank's e-krona pilot", February.

Tobin, J (1987): "The case for preserving regulatory distinctions", proceedings of the Federal Reserve Bank of Kansas City Jackson Hole symposium, pp 167–83.

Whitten, A and D Tygar (1999): "Why Johnny can't encrypt: a usability evaluation of PGP 5.0", proceedings of the 8th USENIX Security Symposium.

Previous volumes in this series

947 June 2021	Money, technology and banking: what lessons can China teach the rest of the world?	Michael Kwok Fai Chui
946 June 2021	The pricing of carbon risk in syndicated loans: which risks are priced and why?	Torsten Ehlers, Frank Packer and Kathrin de Greiff
945 May 2021	US monetary policy and the financial channel of the exchange rate: evidence from India	Shesadri Banerjee and M S Mohanty
944 May 2021	Income inequality, financial intermediation, and small firms	Sebastian Doerr, Thomas Drechsel and Donggyu Lee
943 May 2021	Income inequality and the depth of economic downturns	Emanuel Kohlscheen, Marco Lombardi and Egon Zakrajšek
942 May 2021	FX policy when financial markets are imperfect	Matteo Maggiori
941 May 2021	The digitalisation of money	Markus K Brunnermeier, Harold James and Jean-Pierre Landau
940 May 2021	Monetary-fiscal crosswinds in the European Monetary Union	Lucrezia Reichlin, Giovanni Ricco and Matthieu Tarbé
939 May 2021	The constraint on public debt when $r < g$ but $g < m$	Ricardo Reis
938 May 2021	CCP Auction Design	Wenqian Huang and Haoxiang Zhu
937 April 2021	Growth, coal and carbon emissions: economic overheating and climate change	Emanuel Kohlscheen, Richhild Moessner and Előd Takáts
936 April 2021	The anchoring of long-term inflation expectations of consumers: insights from a new survey	Gabriele Galati, Richhild Moessner and Maarten van Rooij
935 March 2021	An empirical foundation for calibrating the G-SIB surcharge	Alexander Jiron, Wayne Passmore and Aurite Werman
934 March 2021	A global database on central banks' monetary responses to Covid-19	Carlos Cantú, Paolo Cavallino, Fiorella De Fiore and James Yetman

All volumes are available on our website www.bis.org.