Raphael Auer
raphael.auer@bis.org

Rainer Böhme
rainer.boehme@uibk.ac.at

The technology of retail central bank digital currency[1]

Central bank digital currencies (CBDCs) promise to provide cash-like safety and convenience for peer-to-peer payments. To do so, they must be resilient and accessible. They should also safeguard the user's privacy, while allowing for effective law enforcement. Different technical designs satisfy these attributes to varying degrees, depending on whether they feature intermediaries, a conventional or distributed infrastructure, account- or token-based access, and retail interlinkages across borders. We set out the underlying trade-offs and the related hierarchy of design choices.

JEL classification: E42, E44, E51, E58, G21, G28.

The question of whether central banks should issue digital currency to the general public has attracted increasing attention. This special feature sketches out some key technological design considerations for a retail CBDC, in the event that a central bank decided to issue one. We do not investigate the case for or against issuance, the systemic implications, or how these might be managed.[2]

We structure our approach around consumer needs and the associated technical design choices. Current electronic retail money represents a claim on an intermediary, rather than functioning as the digital equivalent of cash. CBDCs could potentially provide a cash-like certainty for peer-to-peer payments. At the same time, they should offer convenience, resilience, accessibility, privacy and ease of use in cross-border payments. Different technical designs meet these criteria to varying degrees, with attendant technical trade-offs. We explore these issues. The aim is not to promote or highlight any particular approach, but to lay some groundwork for more systematic discussions.

[1] We thank Morten Bech, Codruta Boar, Claudio Borio, Stijn Claessens, Benoît Cœuré, Jon Frost, Leonardo Gambacorta, Marc Hollanders, Henry Holden, Ross Leckow, Cyril Monet, Hyun Song Shin, Rastko Vrbaski, Amber Wadsworth and Philip Wooldridge for comments, and Haiwei Cao, Giulio Cornelli and Alan Villegas for exceptional research assistance. The views expressed in this article are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

[2] For the systemic implications, see the survey in CPMI-MC (2018). Andolfatto (2018), Kumhof and Noone (2018), and Bindseil (2020) examine how the impact on the central bank's balance sheet can be managed, while Brunnermeier and Niepelt (2019) investigate how financial instability risks can be mitigated.

Key takeaways

- A trusted and widely usable retail CBDC must be secure and accessible, offer cash-like convenience and safeguard privacy.
- Various technical designs satisfy these criteria to different degrees, and the associated trade-offs need to be identified.
- The design of a retail CBDC needs to balance the credibility of direct claims on the central bank with the benefits of using payment intermediaries.

Our approach is graphically represented in the "CBDC pyramid", which maps consumer needs onto the associated design choices for the central bank. This scheme forms a hierarchy in which the lower layers represent design decisions that feed into subsequent, higher-level decisions.

We start by introducing the four main design choices, as represented in the four layers of the CBDC pyramid. We assess the legal structure of claims and the operational roles of the central bank and private institutions in different CBDC architectures. We discuss the choice between distributed ledger technology (DLT) and a centrally controlled infrastructure. We compare token-based systems and account-based systems. Before concluding, we assess how the development of CBDCs might reinforce current efforts to overhaul cross-border payments.

From consumer needs to design choices: the CBDC pyramid

The focus of our approach is the "retail" aspect of CBDC; we ask what consumer needs a CBDC could address.[3] We thus sketch the development of a CBDC through an approach that proceeds from consumer needs to design choices.[4] The left-hand side of the CBDC pyramid (Graph 1) sets out such consumer needs and six associated features that would make a CBDC useful. Starting with cash-like peer-to-peer usability, these features also comprise convenient real-time payments, payments security, privacy, wide accessibility and ease of use in cross-border payments. The pyramid's right-hand side lays out the associated design choices.

The consumer's prime need is that the CBDC embodies a cash-like claim on the central bank, ideally transferable in peer-to-peer settings. Today, even consumers who normally prefer to pay electronically are confident that, if an episode of financial turmoil were to threaten, they could shift their electronic money holdings into cash. This flight to cash has been seen in many crisis episodes, including recent ones. The main concern is that if, in the future, cash were no longer generally

[3] All private sector non-financial users are referred to as "consumers" in what follows. For a discussion of "wholesale" CBDC for use in the financial industry, see CPMI-MC (2018).

[4] The survey in Boar et al (2020) highlights that central banks have advanced other motivations for issuance, including monetary policy implementation and financial stability considerations. These aspects are considered in the CBDC design frameworks of Fung and Halaburda (2016), Bjerg (2017), CPMI-MC (2018), Mancini-Griffoli et al (2018), Wadsworth (2018), Kahn et al (2019) and Adrian (2019). Although it takes a more positive stance towards CBDC, our focus on technical design elements is related to Pichler et al's (2020) analysis of the limits of CBDC when compared with cash.

accepted, a severe financial crisis might create further havoc by disrupting day-to-day business and retail transactions.[5]

At the same time, consumers are unlikely to adopt a CBDC if it is less convenient to use than today's electronic payments. Banks and payment service providers run sophisticated infrastructures that can handle peak demand, such as on Singles Day in China or Black Friday in the United States. And intermediaries help to smooth the flow of payments by taking on risk, for example during connectivity breaks or offline payments.

These two needs – cash-like safety and convenience of use – lead to the foundational design consideration for a CBDC (see lowest layer of pyramid in Graph 1): the choice of the operational architecture, and how it will balance the consumer's demand for a cash-like claim on the central bank with the convenience that intermediaries confer on the payment system. The choice is shaped by two questions. Is the CBDC a direct claim on the central bank or is the claim indirect, via payment intermediaries? What is the operational role of the central bank and of private sector intermediaries in day-to-day payments?

Further, the consumer's need for cash-like payment safety means that a CBDC must be secure not only from the insolvency or technical glitches of intermediaries, but also from outages at the central bank. The choice is whether to base this infrastructure on a conventional centrally controlled database or instead on DLT – technologies that differ in their efficiency and degree of protection from single

The CBDC pyramid

Graph 1

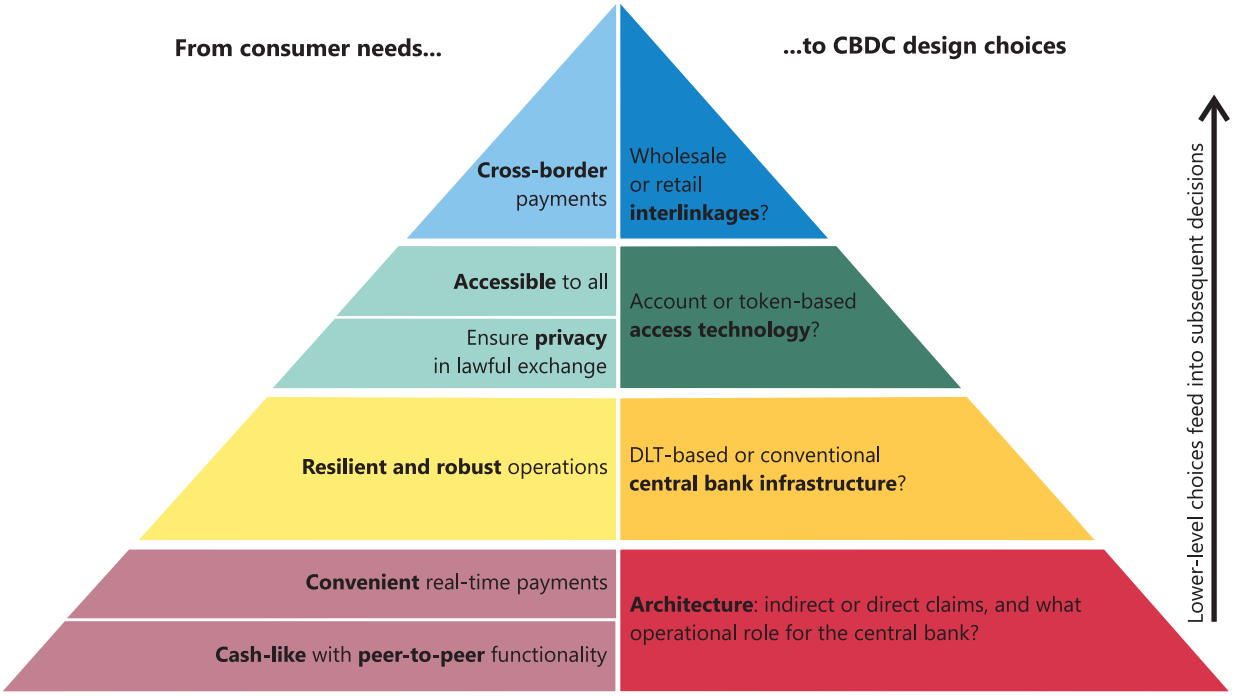

The CBDC pyramid maps consumer needs (left-hand side) onto the associated design choices for the central bank (right-hand side). The four layers of the right-hand side form a hierarchy in which the lower layers represent design choices that feed into subsequent, higher-level decisions.

Source: Authors' elaboration.

[5] In Sweden, where cash use has already declined substantially, considerations along these lines have led the central bank to propose a review of the concept of legal tender (Sveriges Riksbank (2019)).

points of failure. Importantly, this decision can only be made once the architecture has been decided upon, as DLT is only feasible for some operational setups. This is why the choice of infrastructure lies in the pyramid's second layer.

Two further consumer needs are easy, universal access and privacy by default.[6] From a technical perspective, there is an underlying trade-off between privacy and ease of access on the one hand and ease of law enforcement on the other. The associated design choice – the pyramid's third layer – is whether access to the CBDC is tied to an identity system (ie an account-based technology) or instead via cryptographic schemes that do not require identification (ie an access technology based on so-called digital tokens).

The final consumer need we consider is that CBDCs should also enable cross-border payments. At a design level, this could be arranged via technical connections at the wholesale level that are built on today's systems. Alternatively, novel interlinkages could be envisaged at the retail level, ie allowing consumers to hold foreign digital currencies directly. Importantly, the means of implementing the latter option would depend on whether the CBDC was account- or token-based. This is why this design choice belongs in the top layer of the pyramid.

Architecture: indirect or direct claims, and the operational role for the central bank

The CBDC pyramid's bottom layer is the legal structure of claims and the respective operational roles of the central bank and private institutions in payments. Our analysis starts with an overview of possible technical architectures for CBDCs. In all three architectures shown in Graph 2, the central bank is, by definition, the only party issuing and redeeming CBDC. We note that all three architectures could be either account- or token-based, and might run on various infrastructures. These choices are discussed below.

The key differences here are in the structure of legal claims and the record kept by the central bank. In the "indirect CBDC" model (Graph 2, top panel), the consumer has a claim on an intermediary, with the central bank keeping track only of wholesale accounts. In the "direct CBDC" model (centre panel), the CBDC represents a direct claim on the central bank, which keeps a record of all balances and updates it with every transaction. The "hybrid CBDC" model (bottom panel), is an intermediate solution providing for direct claims on the central bank while allowing intermediaries to handle payments.

Consider first the indirect CBDC model (top panel). This term is used by Kumhof and Noone (2018), and is equivalent to the "synthetic CBDC" in Adrian and Mancini-Griffoli (2019). This model is also known as the "two-tier CBDC" for its resemblance to the existing two-tier financial system; a token-based variant is proposed as a "multi-cell CBDC" in Ali (2018). For consumers, this type of CBDC is not a direct claim on the central bank. Instead, the intermediary (labelled "CBDC bank" in Graph 2 for its close resemblance to a narrow payment bank) is mandated to fully back each outstanding indirect CBDC-like liability to the consumer (labelled "ICBDC"

[6] Privacy here means that the consumer's data are used only in steps strictly necessary for the specific purpose of determining whether a transaction is lawful and, if this the case, executing it. "By default" implies that privacy is ensured without requiring any intervention by the user.

in Graph 2) to retail consumers via its holding of actual CBDCs (or other central bank money) deposited at the central bank.[7] Just as in today's system, intermediaries handle all communication with retail clients, net payments and send payment messages to other intermediaries and wholesale payment instructions to the central bank. The latter settles wholesale CBDC accounts with finality.

Besides offering the convenience of today's systems based on intermediaries, the indirect CBDC also relieves the central bank of the responsibility for dispute

An overview of potential retail CBDC architectures — Graph 2

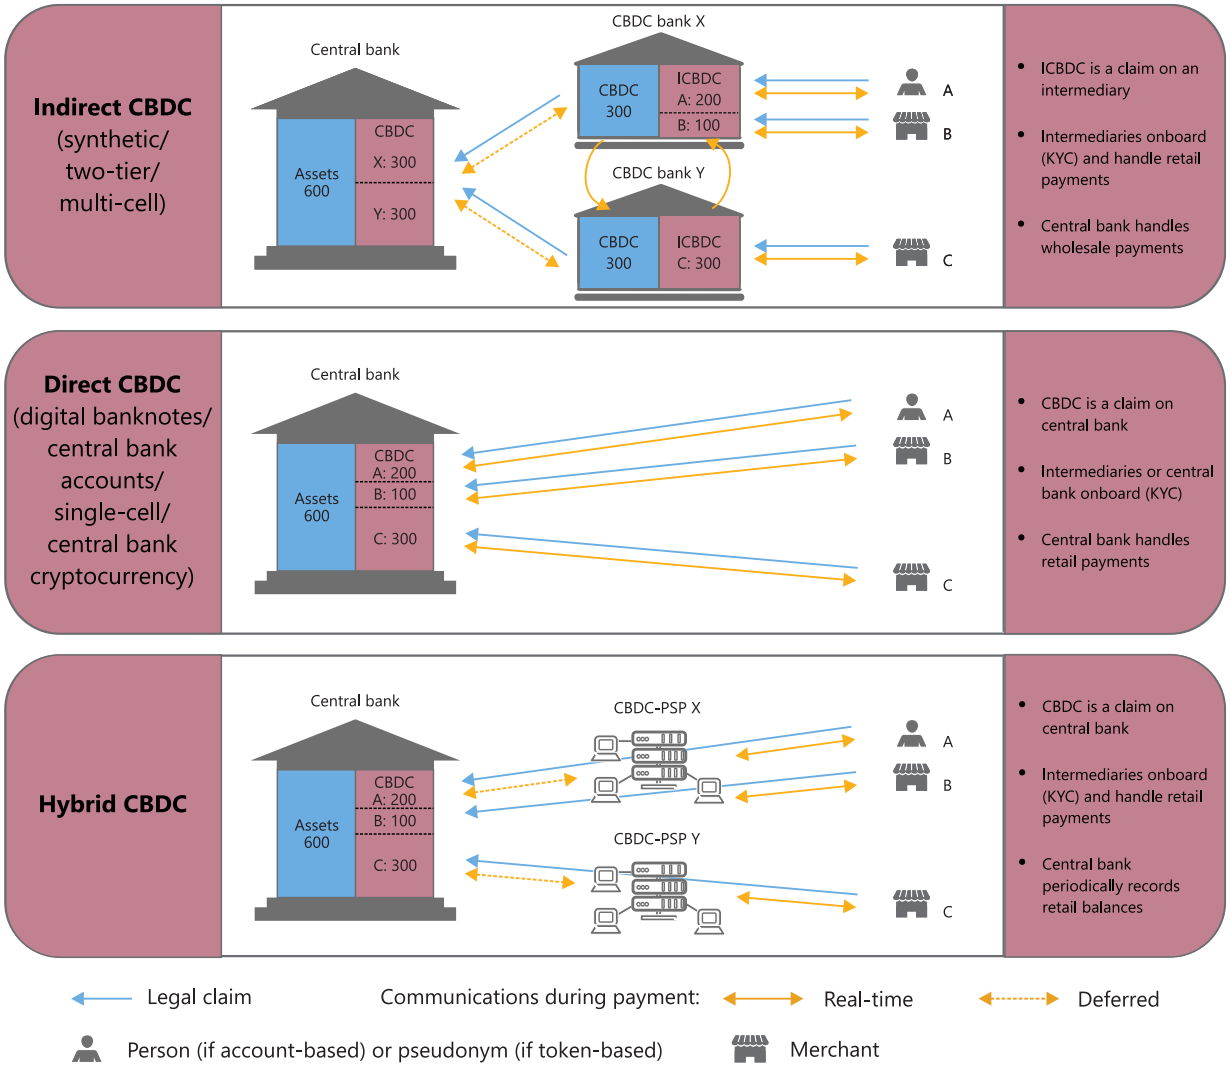

In all three architectures, the CBDC is issued only by the central bank. In the indirect CBDC architecture (top panel), this is done indirectly, and an ICBDC in the hands of consumers represents a claim on an intermediary. In the other two architectures, consumers have a direct claim on the central bank. In the direct CBDC model (centre panel), the central bank handles all payments in real time and thus keeps a record of all retail holdings. The hybrid CBDC model (bottom panel) is an intermediate solution providing for direct claims on the central bank while real-time payments are handled by intermediaries. In this architecture, the central bank retains a copy of all retail CBDC holdings, allowing it to transfer holdings from one payment service provider to another in the event of a technical failure. All three architectures allow for either account- or token-based access.

Source: Authors' elaboration.

[7] Some have argued that this architecture does not warrant the CBDC label. However, the label does apply if one follows CPMI-MC (2018) in defining a retail CBDC as any claim on the central bank that is different from today's wholesale accounts (see also Bech and Garratt (2017)).

resolution, know-your-customer (KYC) and related services. But the downside is that the central bank keeps no record of individual claims (only the intermediaries do, whereas the central bank records only wholesale holdings) nor is there any cash-like direct proof of the claim. Thus, the central bank cannot honour claims from consumers without information from the intermediary.[8] If the intermediary is under stress, determining the legitimate owner might involve a potentially lengthy and costly legal process with an uncertain outcome. This model's regulatory and supervisory issues, as well as those pertaining to deposit insurance, are hence similar to those of today's system.

Consider next a CBDC directly operated by the central bank, the direct CBDC architecture (centre panel). One version would comprise accounts managed by the central bank. Several private sector companies are developing token-based variants, or "digital banknotes".[9] In this architecture, KYC and customer due diligence could be handled by the private sector or the central bank or another public sector institution. The central bank, however, would be the only institution handling payment services.

The direct CBDC is attractive for its simplicity, as it eliminates dependence on intermediaries by doing away with them. However, this entails compromises in terms of the payment system's reliability, speed and efficiency. One aspect is that building and operating technical capacity on this scale is often viewed as being better undertaken by the private sector, as seen in today's credit card networks. Second, even if a central bank were to build the necessary technological capability, the resulting CBDC might be less attractive to consumers than today's retail payment systems. Electronic payments must deal with connectivity outages or offline payments, which involves risk-taking by intermediaries. Importantly, it is the customer relationship – based on KYC – that allows the intermediary to accept such risks. Unless a central bank were to take on responsibility for KYC and customer due diligence – which would require a massive expansion of operations, well beyond existing mandates – it would find it difficult to provide this service.[10]

In addition to these two pure options, one can also envisage novel future solutions that merge elements of both the indirect and the direct CBDC.[11] We label this third type of architecture the hybrid CBDC (bottom panel). In this model, a direct claim on the central bank is combined with a private sector messaging layer. Again, variations on this theme might include both token- and account-based ones.

One key element of the hybrid CBDC architecture is the legal framework that underpins claims, keeps them segregated from the balance sheets of the payments service providers (PSPs), and allows for portability. If a PSP fails, holdings of the

[8] A further difficulty is that it is unclear what the holder of an ICBDC would actually be entitled to, as, by definition, retail investors are prohibited from holding the actual CBDCs issued by the central bank.

[9] These token-based versions are termed "single-cell" CBDC structures in Ali (2018) and "central bank cryptocurrencies" in Berentsen and Schär (2018).

[10] The respective advantages and disadvantages of direct and indirect CBDC architectures mirror those of the direct and indirect security holding systems that are discussed in the context of the future of settlement in Bech, Hancock, Rice and Wadsworth (2020, in this issue).

[11] Although these authors do not spell out the underlying structure of legal claims, several ways to distribute payment functions and communications over multiple parties have been studied in the field of computer science. One example is the proposal of Danezis and Meiklejohn (2016), which shifts real-time communications from the central bank to dynamically appointed intermediaries.

CBDC are not considered part of the PSP's estate available to creditors. The legal framework should also allow for portability in bulk, ie give the central bank the power to switch retail customer relationships from a failing PSP to a fully functional one.[12] The second key element is the technical capability to enable the portability of holdings. Since the requirement is to sustain payments when one intermediary is under technical stress, the central bank must have the technical capability to restore retail balances. It thus retains a copy of all retail CBDC holdings, allowing it to transfer retail CBDC holdings from one PSP to another in the event of a technical failure.[13]

The hybrid CBDC would have both advantages and disadvantages vis-à-vis the indirect or direct CBDC architectures. As an intermediate solution, it might offer better resilience than the indirect CBDC, but at the cost of a more complex to operate infrastructure for the central bank. On the other hand, the hybrid CBDC is still simpler to operate than a direct CBDC. As the central bank does not directly interact with retail users, it can concentrate on a limited number of core processes, while intermediaries handle other services including instant payment confirmation.

Conventional or DLT-based central bank infrastructure?

What infrastructure might the different CBDC architectures require for the central bank, and how could they be implemented in the most resilient way? This choice, represented as the second tier of the CBDC pyramid, follows immediately after the decision on architecture because the infrastructure requirements for the central bank differ substantially across the three architectures shown in Graph 2.

For the central bank, the indirect CBDC implies loads similar to those of today's system. By contrast, the direct CBDC would require massive technological capabilities, as the central bank processes all transactions by itself, handling a volume of payments traffic comparable with that of today's credit or debit card operators. The hybrid CBDC architecture is more complex to operate than the indirect model, as the central bank does maintain retail balances. Nevertheless, it could be implemented at scale using today's technology and with a relatively modest infrastructure even in the world's largest currency areas.[14]

The infrastructure could be based on a conventional centrally controlled database, or on a novel distributed ledger. Graph 3 shows how elements of DLT could play a role in CBDC. The first DLT-related design choice hinges on whether

[12] While functionally similar, such segregation differs from deposit insurance in terms of legal procedures and associated delays. Today's deposits are often insured but, in the case of a bank failure, the funds can only be retrieved through a reimbursement process. Further, deposit insurance may be limited in amount and ultimately depends on the strength of the deposit insurer (see Baudino et al (2019) for an overview).

[13] Note that a variant of this CBDC architecture could allow users to retain cryptographic proofs of their CBDC balances, rather than oblige the central bank to hold them. These proofs could be used to retrieve balances in case of a technical failure. The advantage would be to circumvent potential privacy and legal issues connected with the central bank storing retail account balances. The disadvantage would be that entrusting users with cryptographic proofs may open the door to loss and theft of funds.

[14] For example, even for a payment area with a billion users, it would be feasible to verify each digital signature (computationally, the most costly operation of a transaction) for all accounts on an hourly basis with a two-digit number of standard servers.

the authority to update the database is centralised or delegated to a network of identified and vetted validators.[15]

Conventional and DLT-based infrastructures often store data multiple times and in physically separate locations. The main difference between them lies in how data are updated. In conventional databases, resilience is typically achieved by storing data over multiple physical nodes, which are controlled by one authoritative entity – the top node of a hierarchy. By contrast, in many DLT-based systems, the ledger is jointly managed by different entities in a decentralised manner and without such a top node. Consequently, each update of the ledger has to be harmonised between the nodes of all entities (often using algorithms known as "consensus mechanisms"). This typically involves broadcasting and awaiting replies on multiple messages before a transaction can be added to the ledger with finality.

The overhead needed to operate a consensus mechanism is the main reason why DLTs have lower transaction throughput than conventional architectures. Specifically, these limits imply that current DLT could not be used for the direct CBDC except in very small jurisdictions, given the probable volume of data throughput. However, DLT could be used for the indirect CBDC architecture, as the number of transactions in many wholesale payment systems is comparable with that handled by existing blockchain platforms, as also demonstrated in several wholesale

Elements of decentralisation: DLT and token-based access — Graph 3

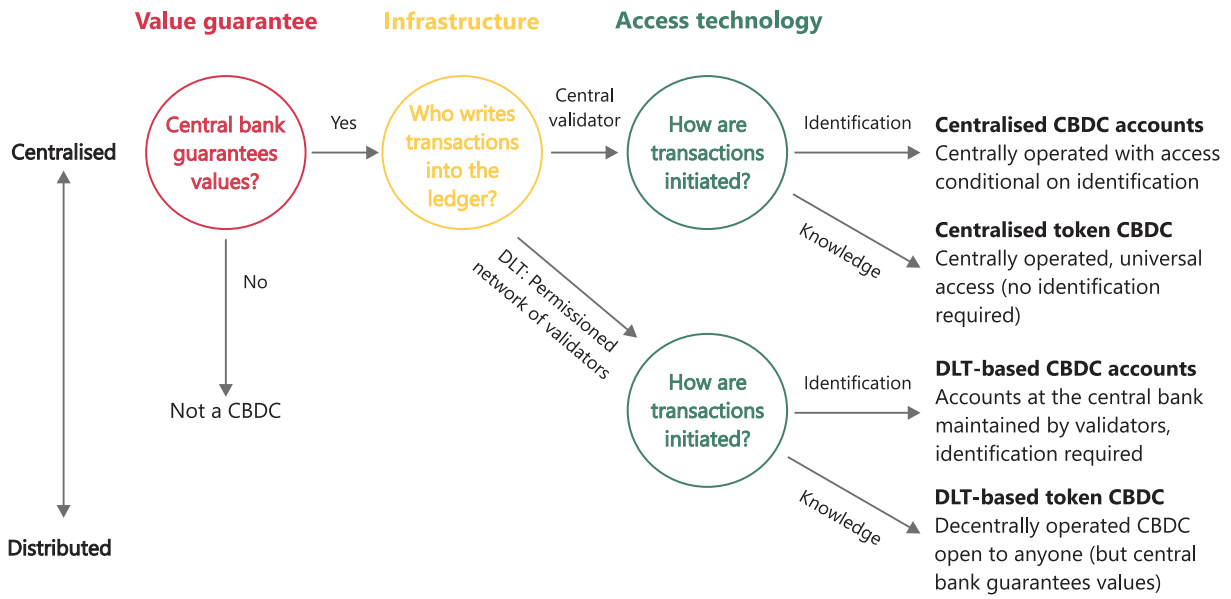

This graph maps out the four possible combinations of whether a CBDC infrastructure is distributed or centralised and whether access is based on identification (accounts) or cryptographic knowledge (digital tokens). All four combinations are possible for any CBDC architecture (indirect, direct or hybrid), but in the different architectures, the central bank and the private sector operate different parts of the respective infrastructure.

Source: Authors' elaboration.

[15] Most likely, central banks would consider only "permissioned" DLT, in which a network of pre-selected entities performs the updating. While it is technically possible to use "permissionless" technology, in which unknown validators perform the updating, the economic cost of this process is very high (see Böhme et al (2015) for an introduction for the case of Bitcoin, Auer (2019a) for a discussion of the underlying economics and Ali and Narula (2020) for a specific analysis in the context of CBDCs).

CBDC experiments conducted by central banks (Bech, Hancock, Rice and Wadsworth (2020, in this issue)). Enterprise versions of DLT might also be feasible for the hybrid CBDC architecture.

When it comes to achieving resilience, neither a DLT-based system nor a conventional one has a clear-cut advantage. The vulnerabilities are simply different. The key vulnerability of a conventional architecture is the failure of the top node, for example via a targeted hacking attack. The key vulnerability of DLT is the consensus mechanism, which may be put under pressure, for example, by a denial-of-service type of attack.

Overall, one needs to weigh carefully the costs and benefits of using DLT. This technology essentially outsources to external validators the authority to adjust claims on the central bank balance sheet,[16] which is advantageous only if one trusts this network to operate more reliably than the central bank. Ongoing assessments of DLT-based proofs-of-concept tend to be negative (see box for a brief overview). Among the DLT-based projects that are still ongoing, it remains to be seen whether scalable implementations will actually rely on the technology.[17]

That said, even if one decides against using DLT as the backbone infrastructure of a CBDC, one closely related technology might still be useful. Whether or not the infrastructure is based on DLT, access can still be based on cryptography rather than identification – Graph 3 outlines the possible combinations, and the box shows which combinations are being investigated by central banks.

Token- or account-based access, and how to safeguard privacy?

Once the CBDC's architecture and infrastructure have been chosen, the question arises of how and to whom one should give access. This is the third layer of the CBDC pyramid.

A first option is to follow the conventional account model and tie ownership to an identity (Graph 4, left-hand side). Claims are represented in a database that records the value along with a reference to the identity, just as in a bank account. This has drawbacks in the case of CBDCs. In particular, it depends on "strong" identities for all account holders – schemes that map each individual to one and only one identifier across the entire payment system. Such schemes can present a challenge in some jurisdictions, thus impairing universal access.[18]

The second option is for the central bank to honour claims solely when the CBDC user demonstrates knowledge of an encrypted value – an option sometimes

[16] In the indirect CBDC architecture, validators of the central bank ledger update the wholesale accounts, while in the hybrid and direct architectures they update the retail accounts.

[17] Experiments are based on enterprise versions of distributed ledgers, which allow for decentralisation but, in practice, are often run under centralised control. Ali and Narula (2020, p 6) note that the platforms typically used "are useful for experimentation and prototyping because of their flexibility and features […]. However, what is helpful for prototyping might not be good for practice; these complex platforms make trade-offs when it comes to security, stability, and scale."

[18] There are broader benefits to universal digital identity frameworks, such as the scope for supporting open banking and enabling the distribution of other financial services. D'Silva et al (2019) discuss the Indian experience.

referred to as digital tokens (Graph 4, right-hand side). One example is when the secret part of a public-private key pair is used to sign a message, a technology outlined by Auer, Böhme and Wadsworth (2020, in this issue).

A token-based system would ensure universal access – as anybody can obtain a digital signature – and it would offer good privacy by default. It would also allow the CBDC to interface with communication protocols, ie be the basis for micropayments in the internet of things. But the downsides are severe. One is the high risk of losing funds if end users fail to keep their private key secret. Moreover, challenges would arise in designing an effective AML/CFT framework for such a system. Law enforcement authorities would run into difficulties when seeking to identify claim owners or follow money flows, just as with cash or bearer securities. Retail CBDCs would thus need additional safeguards if they followed this route.[19]

We emphasise that the privacy dimension goes far beyond the question of whether the system is based on accounts or digital tokens. Transaction-level financial data reveal sensitive personal data. Hence, two aspects of privacy by default are crucial for the design of a CBDC. First is the amount of personal information transaction partners learn about each other when the system is operating normally.[20] Second is the risk of large-scale breaches of data held by the system operator or intermediaries.

Crucially, a CBDC that lets merchants collect and link payment data to customer profiles transforms the very nature of payments, from a simple exchange of value to the exchange of value for a bundle of data. Hence, a CBDC should preserve its users' privacy vis-à-vis their transaction partners, ie by default, transaction partners

Account-based access compared with token-based access — Graph 4

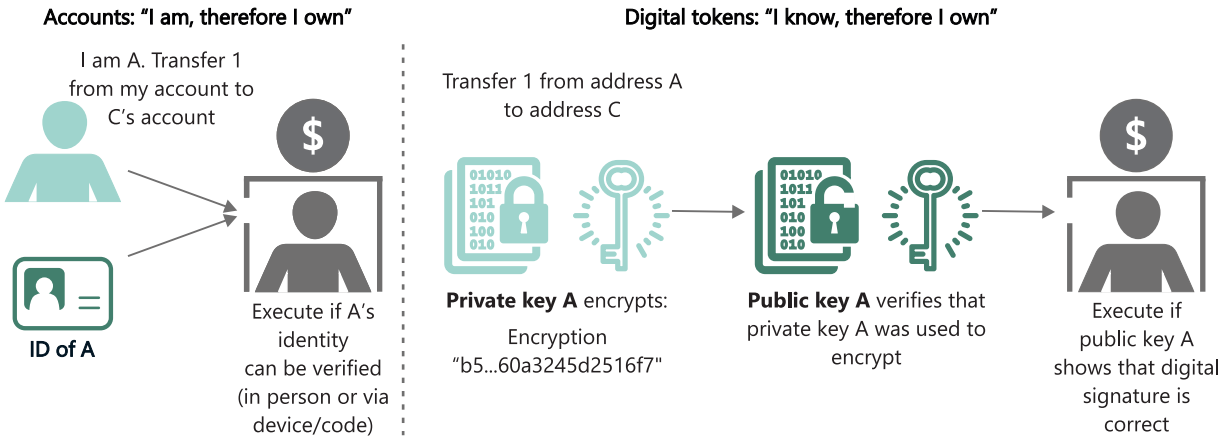

In an account-based CBDC (left-hand side), ownership is tied to an identity, and transactions are authorised via identification. In a CBDC based on digital tokens (right-hand side), claims are honoured based solely on demonstrated knowledge, such as a digital signature.

Source: Authors' elaboration.

[19] The legal framework would need to allow claims to be put "on hold" until the legitimacy of a transaction history has been demonstrated (Böhme et al (2015)). This could be part of a broader regulatory framework allowing for "embedded supervision", ie an approach in which supervisory and other public authorities automatically monitor market ledgers to check for compliance with regulatory goals (Auer (2019b)).

[20] See ECB (2019) for a practical proposal and Frost et al (2019) for a broader discussion of the use of data in finance.

would interact via "unlikable pseudonyms", as envisaged in Chaum's (1985) pioneering work on electronic money. In such a system, a merchant is presented with a proof that the payment for a specific invoice has been made, but no information about the payee is revealed.

Depending on the involvement of intermediaries and the information they receive, technical safeguards for data protection need to be complemented by a legal framework restricting data collection by front-end applications, for example the smartphone payment app. Data loss is a further threat, given that payment systems are a prime target for cyber attacks. In this context, it must be noted that not all privacy-enhancing technologies are mature. For example, some so-called zero-knowledge proofs have already been shown to be vulnerable (Ruffing et al (2018)). The only sure-fire way to avoid losing much data is not to store it or to irrevocably delete old transactions as soon as possible. This principle of data minimisation is embodied in many data protection laws. Where this is not an option, aggregation and anonymisation must be relied on. A last resort is storage in physically separated (and offline) places guarded by legal access procedures.

Cross-border payments: wholesale or retail linkages?

Once a CBDC's configuration is clear, as well as how resident consumers can access it, the question arises whether it can be used only domestically or also elsewhere. This is the topmost layer of the CBDC pyramid.

The demand for seamless and inexpensive cross-border payments has grown in parallel with growth in international e-commerce, remittances and tourism. A CBDC might come with the same wholesale interlinkage options explored in the current system (Bech, Faruqui and Shirakami (2020, in this issue)).

Here, one noteworthy aspect is that a coordinated CBDC design effort could take a clean-slate perspective and incorporate these interlinkage options right from the start. This would represent a unique opportunity to facilitate easier cross-border payments (eg Carney (2019) and Cœuré (2019)), reducing inefficiencies and rents by shortening the payment value chain.

CBDCs would also permit novel retail interlinkages if they were to allow consumers to hold multiple currencies. In today's account-based system, a cross-border transaction is inseparably linked to a foreign exchange transaction. The intermediary processing the transaction can apply extra fees and unfavourable exchange rates. In contrast, if consumers were given the option of buying foreign currency in advance, before spending it abroad, just as they can with cash, this would separate the payment from the foreign exchange transaction. In turn, this would open up the possibility of interfacing retail wallets directly with competitive foreign exchange markets.

Importantly, the scope for such retail interlinkages and their design would depend on the national access framework. If a national system is based on digital tokens, it will by default be accessible to foreign residents. If it is account-based, interoperability would be a design choice, one that could also be coordinated internationally.

Conclusion

As central banks play a key role in payment systems, both the declining use of cash and related developments in the private sector may require them to "step up" and take a more active role (Carstens (2019 and 2020, in this issue)). Should they wish to do so, many ways are open to them.

This feature has gone down a hypothetical road by investigating the choices that might be encountered during the design stage of a CBDC, and how the related decision-making process could be structured. On the way, we have highlighted how consumer needs might translate into technical trade-offs. Some design-related considerations emerge from our analysis, for example, regarding the feasibility of DLT-based vis-à-vis that of more conventional technical infrastructures, but other choices remain less clear-cut.

With a framework for decision-making in mind, more hands-on experience with specific design choices could be helpful. The box surveys ongoing technical design efforts by central banks along the technical dimensions identified in this feature. As most projects are still in their early stages, the most important takeaway is that central banks around the world are investigating a rich set of prototypes, spanning almost the full range of potential designs encompassed in the CBDC pyramid. If the results of these experiments are shared internationally, a clearer picture will emerge of which technological choices are generally suited for CBDCs, and how the optimal design might depend on the specific circumstances of each jurisdiction. This, in turn, could help to inform the debate on whether and how CBDCs should actually be issued.

Taking stock: ongoing retail CBDC projects

Raphael Auer, Giulio Cornelli and Jon Frost

Among the many central banks that are exploring the possibility of a retail CBDC (Boar et al (2020)), several have published research or statements on the related motivations, architectures, risks and benefits. The table below shows 17 selected projects or reports published before 19 February 2020. It does not cover wholesale CBDCs or cross-border payment projects that do not involve a CBDC. When it comes to the four main design choices (Graph 1 in the main text), many central banks are still considering multiple options, and it is not always possible to classify them. Regarding their architecture (Graph 2 in the main text), five projects focus on a direct CBDC, two on an indirect CBDC, and 10 investigate several designs or do not specify the architecture.

As for infrastructure (Graph 3 in the main text), only one project focuses on a conventional technology, whereas five focus on DLT. However, experience with the latter technology has not always been encouraging. Sveriges Riksbank (2018) notes that DLT still suffers from inadequate performance and scalability. The National Bank of Ukraine (2019) concludes that DLT may offer no fundamental advantages in a centralised issuance system. More generally, ECCB (2020) notes that DLT could not ensure cash-like resilience in the case of prolonged electricity outages.

On the access technology (Graph 4 in the main text), three projects provide for access based on digital tokens, whereas three focus on account-based access.

Regarding the focus on cross-border interlinkages, no CBDC project has an explicit focus on payments beyond the central bank's jurisdiction. It is noteworthy that several central banks are working on cross-border payment trials with a consumer focus in parallel to their CBDC efforts. Moreover, wholesale initiatives such as Project Jasper (Bank of Canada), Project Ubin (Monetary Authority of Singapore), Project Stella (ECB and Bank of Japan) and Project Lion Rock-Inthanon (Hong Kong Monetary Authority and Bank of Thailand) might potentially help support more efficient retail transactions through the banking system.

Only very few projects have already been completed, with considerable variation in the results. A few jurisdictions, including Denmark and Switzerland, have determined that, currently, the costs of a retail CBDC would outweigh the benefits. A larger number continues to actively develop retail CBDCs; Boar et al (2020) find that over a third of all surveyed central banks say that issuing a retail CBDC is a medium-term possibility. Looking ahead, the overall conclusion from a technological perspective is that a rich set of technical designs are currently under consideration. This underscores the need for international coordination to share experience.

Selected retail CBDC projects

Table A

Design choices				Project/country	Notes on status, motivation and conclusion
Architecture[1]	Infrastructure[2]	Access[3]	International[4]		
D	U	A	N	Rafkróna Iceland	Research; aims to address "steadily diminishing use of banknotes and coin"; "many issues have yet to be clarified, and they must be dealt with appropriately before a position can be taken".
D	U	A	N	Sand Dollar The Bahamas	Pilot; improve "financial inclusion …, [reduce] the size of legitimate but unrecorded economic activities, [strengthen] national defences against money laundering and other illicit ends [and]… deliver government services through digital channels, thereby improving tax administration and increasing the efficiency of spending".

D	U	U	N	E-krona* Denmark	Research; "the potential benefits of introducing CBDC [are not assessed to] match the considerable challenges that the introduction would present".
D	U	U	N	E-krona* Norway	Working group; focus on "independent back-up solution, credit risk-free alternative to bank deposits, competition, legal tender"; "more information is required before a conclusion can be reached".
D	U	U	N	E-krona Sweden	Ongoing work; "within a few years, if the current trend continues, we will find ourselves in a situation where cash is no longer generally accepted as a means of payment"; "an account-based e-krona could rationalise payments from agencies and make them less dependent on commercial agents".
I	D	T	N	Digital fiat currency Brazil	Research; "Improve the efficiency of the monetary function, … payment processes and systems, …. financial inclusion and … user experience".
I	D	U	I	E-euro* ECB	Research; "CBDC with the status of legal tender could guarantee that all users have, in principle, access to a cheap and easy means of payment"; "proof of concept also highlights a number of areas where there is room for improvement".
U	C	A	N	Dinero Electrónico Ecuador	Pilot; "means of payment available to absolutely all Ecuadorians". Operated 2014–16; discontinued.
U	D	T	I	DXCD Eastern Caribbean	Pilot; aims to address the "high cost of current payment instruments and banking services", needs of customers and inefficient cheque settlement.
U	D	U	N	Bakong Cambodia	Pilot; aims to "increase access to quality formal financial services"; "decrease demand for… cash".
U	D	U	N	E-hryvnia Ukraine	Pilot; test DLT "as a technological framework for e-hryvnia issuance and circulation"; no fundamental advantage in using DLT in a centralised model.
U	U	T	N	Electronic legal tender South Africa	Expression of interest; "The scope of this project is specific to the use of a CBDC as electronic legal tender (ELT), similar to the characteristics of, and complementary to, cash."
U	U	U	N	Billete Digital Uruguay	Pilot; "Digital bills that aim to have same functions and uses as physical bills"; ongoing evaluation.
U	U	U	N	DC/EP (Digital Currency/Electronic Payments) China	Ongoing work; aims to create digital alternative to cash and coins for retail use.
U	U	U	N	E-shekel Israel	Research; "help in the struggle against … unreported transactions"; "contribute to the high-tech sector (fintech)"; Conclusion that "the team does not recommend that the Bank of Israel issue digital currency (e-shekel) in the near future".
U	U	U	U	E-euro* France	Research; "account-based model would offer better results for a retail CBDC. However, it might also lead to a greater loss of resources for banks".
U	U	U	U	E-franc Switzerland	Research; "Examine the opportunities and risks of introducing a cryptofranc (e-franc)"; "additional benefits currently low and outweighed by risks".

[1] D = direct; I = indirect; U = unspecified or multiple options under consideration. [2] C = conventional; D = DLT; U = unspecified or multiple options under consideration. [3] A = account-based; T = token-based; U = unspecified or multiple options under consideration. [4] I = international; N = national; U = unspecified or multiple options under consideration. * Not an official designation.

Sources: Central bank websites; www.unescap.org; www.efd.admin.ch; www.cf40.org.cn.

References

Adrian, T (2019): "Stablecoins, central bank digital currencies, and cross-border payments", lecture at International Monetary Fund-Swiss National Bank Conference, Zurich, 14 May.

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

Ali, R (2018): "Cellular structure for a digital fiat currency", paper presented at the P2P financial system international workshop, Federal Reserve Bank of Cleveland, 27 July.

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

Andolfatto, D (2018): "Assessing the impact of central bank digital currency on private banks", Federal Reserve Bank of St Louis, *Working Papers*, no 2018-25.

Auer, R (2019a): "Beyond the doomsday economics of 'proof-of-work' in cryptocurrencies", *BIS Working Papers*, no 765.

——— (2019b): "Embedded supervision: how to build regulation into blockchain finance", *BIS Working Papers*, no 811.

Auer, R, R Böhme and A Wadsworth (2020): "An introduction to public-private key cryptography in digital tokens", *BIS Quarterly Review*, March, p 73.

Baudino, P, R Defina, J Fernández Real, K Hajra and R Walters (2019): "Bank failure management – the role of deposit insurance", *FSI Insights on policy implementation*, no 17, August.

Bech, M and R Garratt (2017): "Central bank cryptocurrencies", *BIS Quarterly Review*, September, pp 55–70.

Bech, M, U Faruqui and T Shirakami (2020): "Payments without borders", *BIS Quarterly Review*, March, pp 53–65.

Bech, M, J Hancock, T Rice and A Wadsworth (2020): "On the future of securities settlement", *BIS Quarterly Review*, March, pp 67–83.

Berentsen, A and F Schär (2018): "The case for central bank electronic money and the non-case for central bank cryptocurrencies", *Federal Reserve Bank of St Louis Review*, vol 100, no 2, pp 97–106.

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

Bjerg, O (2017): "Designing new money – The policy trilemma of central bank digital currency", *Copenhagen Business School Working Papers*.

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.

Böhme, R, N Christin, B Edelman and T Moore (2015): "Bitcoin: economics, technology, and governance", *Journal of Economic Perspectives*, vol 29, no 2, pp 213–38.

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 (2019): "The future of money and the payment system: what role for central banks?", lecture at Princeton University, 5 December.

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

Chaum, D (1985): "Security without identification: transaction systems to make big brother obsolete", *Communications of the ACM*, vol 28, no 10, pp 1030–44.

Cœuré, B (2019): "Digital challenges to the international monetary and financial system", conference on "The future of the international monetary system", Luxembourg, 17 September.

Committee on Payments and Market Infrastructures and Markets Committee (2018): *Central bank digital currencies*, March.

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

D'Silva, D, Z Filková, F Packer and S Tiwari (2019): "The design of digital financial infrastructure: lessons from India", *BIS Papers*, no 106, December.

Eastern Caribbean Central Bank (2020): "ECCB digital EC currency pilot: what you should know", accessed 27 January.

European Central Bank (2019): "Exploring anonymity in central bank digital currencies", *In Focus*, no 4, December.

Frost, J, L Gambacorta, Y Huang, H S Shin and P Zbinden (2019): "BigTech and the changing structure of financial intermediation", *BIS Working Papers*, no 779, April.

Fung, B and H Halaburda (2016): "Central bank digital currencies: a framework for assessing why and how", Bank of Canada, *Staff Discussion Papers*, no 22.

Kahn, C, F Rivadeneyra and T-N Wong (2019): "Should the central bank issue e-money?", *Federal Reserve Bank of St Louis Working Papers*, no 3.

Kumhof, M and C Noone (2018): "Central bank digital currencies – design principles and balance sheet implications", *Bank of England Working Papers*, no 725.

Mancini-Griffoli, T, M Peria, I Agur, A Ari, J Kiff, A Popescu and C Rochon (2018): *Casting light on central bank digital currencies*, International Monetary Fund, November.

National Bank of Ukraine (2019): *Analytical Report on the E-hryvnia Pilot Project*, February.

Pichler, P, M Summer and B Weber (2020): "Does digitalization require central bank digital currencies for public use?", *Monetary Policy and the Economy*, forthcoming.

Ruffing, T, S Thyagarajan, V Ronge and D Schroder (2018): "Burning Zerocoins for fun and for profit – a cryptographic denial-of-spending attack on the Zerocoin protocol", proceedings of the Crypto Valley conference on Blockchain Technology, Institute of Electrical and Electronics Engineers, pp 116–9.

Sveriges Riksbank (2018): "The Riksbank's e-krona project: Report 2", October.

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

Wadsworth, A (2018): "The pros and cons of issuing a central bank digital currency", *Reserve Bank of New Zealand Bulletin*, vol 81, no 7, June.