Study Guide

CBDC Papers Lecture Series

CBDCTHINKTANK

Baruch COLLEGE ZICKLIN SCHOOL OF BUSINESS | Frankfurt School Blockchain Center | DEA | ZIBS

Contents

INTRODUCTION

1. About Us

The CBDC TT is a group of bankers, researchers, academics, members of international financial institutions (IFIs) and non-government organizations (NGOs) interested in (and, dare we say, passionate about) the future of digital currencies and payments. https://cbdctt.com"

2. Welcome: Leadership

Prof. Jamiel Sheikh

Founder, CBDC Think Tank

CBDCs continue to evolve and today are part of the conversation on innovation. Despite having a late start, central banks, academics and the private sector have stepped up to innovate in the space in surprising and interesting ways. We look forward to you joining us in exploring the most current innovations and understanding what the future of CBDCs may look like as these innovations become more and more nuanced, complex and intricate. The papers presented in this series should be a shot in the arm for those looking to accelerate with the current CBDC thinking from some of the best thinkers in the space.

Bruno Silvestre

Chief Liaison Office, CBDC Think Tank

Most professionals passionate about digital currencies tend to agree. The question isn't whether CBDCs will see the light but when. Here at the CBDC Think Tank, we ask ourselves: Is that so? And why? The Think Tank represents the best and most appropriate venue to try and answer these essential questions. Educating is part of our DNA, but so is researching, analyzing, and discussing CBDC-related issues. Our upcoming Papers Lecture Series will be a major step in doing just that and we have no doubt that all the presentations will provide significant, noteworthy and substantial food for thoughts for anyone interested in this new challenge for Central banks.

John Kiff

Managing Director, Education

Welcome to the CBDC Think Tank Workshop's CBDC Papers Lecture Series. The papers presented here represent the leading edge of the "academic" yet practical work on CBDC, and the authors come from some of the most prestigious institutions working in this space. The papers have been carefully selected to cover all of the key facets of the CBDC thought process, from a deep dive into which central banks are doing what, on to how to design CBDCs and their ecosystems for success, and ending with some deep dives into the potential impacts of CBDC on the banking sector. And all the sessions are designed to encourage interaction, so please do not be shy to challenge the authors with comments and questions!

3. Welcome: Advisors

Prof. Dr. Jiaying Jiang

Advisor, CBDC Think Tank

Assistant Professor of Law at UF Law and Hauser Global Fellow at NYU

While 105 countries, representing over 95 percent of global GDP, are exploring CBDCs, CBDC Think Tank has taken the lead in educating business leaders and the general public on CBDCs. If adopted, CBDCs would have enormous impacts on businesses and households. Learning about CBDCs has been necessary and urgent for everyone.

Jonas Gross

Advisor, CBDC Think Tank

Chairman Digital Euro Association

We are thrilled that you are interested in doing a deep dive into the world of central bank digital currencies (CBDCs). To properly understand the implications of CBDCs and to evaluate potential CBDC designs, it is necessary not only to learn about CBDC projects worldwide, but also to study academic CBDC papers. For this reason, the CBDC Thank Tank, created this study guide to give you some first papers on hand to study CBDCs from the academic perspective. The Digital Euro Association is collaborating closely with the CBDC Think Tank to educate around digital money, build a community, and provide valuable impulses to policy-makers that are currently evaluating a CBDC issuance.

Prof. Joseph Onochie

Advisor, CBDC Think Tank

Associate Professor of Finance, Zicklin School of Business at Baruch College

As the world moves, inexorably, down the path of coexistence with the impact of technology on business models everywhere, business educators continue to grapple with the most effective ways to prepare our students for the ever-evolving environment. Everywhere, there is a need to strike a balance between fundamental concepts and cutting-edge concepts. Invariably, educators must choose among competing topics to focus on and then hope that, as practitioners, our students will continue to educate themselves as required. A study guide on the Central Bank Digital Currency (CBDC) will be an invaluable source of specific, focused and readily available information on this fast changing field.

BEHIND THE SCENES OF CENTRAL BANK DIGITAL CURRENCY

FINTECH

NOTES

Behind the Scenes of Central Bank Digital Currency

Emerging Trends, Insights, and Policy Lessons

Gabriel Soderberg

In collaboration with Marianne Bechara, Wouter Bossu, Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, Tommaso Mancini-Griffoli, Tao Sun, and Akihiro Yoshinaga

NOTE/2022/004

FINTECH NOTE

Behind the Scenes of Central Bank Digital Currency

Emerging Trends, Insights, and Policy Lessons

Prepared by Gabriel Soderberg in collaboration with Marianne Bechara, Wouter Bossu, Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, Tommaso Mancini-Griffoli, Tao Sun, and Akihiro Yoshinaga

February 2022

Behind the Scenes of Central Bank Digital Currency
Emerging Trends, Insights, and Policy Lessons

Note 2022/004
Prepared by Gabriel Soderberg
In collaboration with Marianne Bechara, Wouter Bossu, Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, Tommaso Mancini-Griffoli, Tao Sun, and Akihiro Yoshinaga

Names: Soderberg, Gabriel, author. | Bechara, Marianne, author. | Bossu, Wouter, author. | Che, Natasha Xingyuan, author. | Davidovic, Sonja, author. | Kiff, John, author. | Lukonga, Inutu, author. | Mancini-Griffoli, Tommaso, author. | Sun, Tao, 1970-, author. | Yoshinaga, Akihiro, author. | International Monetary Fund, publisher.
Title: Central bank digital currency behind the scenes : emerging trends, insights, and policy lessons / prepared by Gabriel Soderberg in collaboration with Marianne Bechara, Wouter Bossu, Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, Tommaso Mancini-Griffoli, Tao Sun, and Akihiro Yoshinaga.
Other titles: Emerging trends, insights, and policy lessons. | Fintech Notes (International Monetary Fund).
Description: Washington, DC : International Monetary Fund, 2022. | February 2022. | Note 2022/004. | Fintech notes. | Includes bibliographical references.
Identifiers: ISBN 9798400201219 (paper)
ISBN 9798400200373 (ePub)
ISBN 9798400200403 (WebPDF)
Subjects: LCSH: Digital currency. | Banksand banking, Central. | International finance.
Classification: LCC HG1710.S63 2022

The paper was written under the supervision of Tommaso Mancini-Griffoli, and drew on work by Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, and Tao Sun. Marianne Bechara, Wouter Bossu, and Akihiro Yoshinaga researched and wrote the section on legal foundations of CBDC. While the authors are responsible for any mistakes, the paper greatly benefitted from interactions with central bank representatives: Cleopatra Davis and Kimwood Mott (Central Bank of the Bahamas), Scott Henry, Francisco Rivadeneyra, and Dinesh Shah (Bank of Canada), Changchun Mu, Naji Liu, and Yuan Lyu (People's Bank of China), Sharmyn Powell (Eastern Caribbean Central Bank), Gabriela Guibourg, Johan Schmalholz, and Mithra Sundberg (Sveriges Riksbank), and Adolfo Sarmiento (Banco Central de Uruguay). The paper benefitted from comments by IMF staff. Carlos Padilla-Chavez provided outstanding research assistance, and Erica Sandoval provided precious editorial assistance.

Publication orders may be placed online or through the mail:
International Monetary Fund, Publication Services
P.O. Box 92780, Washington, DC 20090, U.S.A.
T. +(1) 202.623.7430
publications@IMF.org
IMFbookstore.org
elibrary.IMF.org

Contents

1. Introduction

Central banks are increasingly pondering whether to issue their own digital currencies to the general public, so-called retail central bank digital currency (CBDC).[1] The majority of IMF member countries are actively evaluating CBDCs, with only a few having issued CBDCs or undertaken extensive pilots or tests.[2]

This paper shines the spotlight on the handful of countries at the frontier in the hope of identifying and sharing insights, lessons, and open questions for the benefit of the many countries following in their footsteps. Clearly, what can be gleaned from these experiences does not necessarily apply elsewhere. The sample of countries remains small and country circumstances differ widely. However, the insights in this paper may inspire further investigation and allow countries to gain time by building on the experience of others. Importantly, the purpose of this paper is not to evaluate the courses taken by different jurisdictions, but to study and discuss their key experiences and lessons.

The paper studies six advanced CBDC projects, drawing on collaboration and exchanges with the respective central banks to get insights beyond what has previously been published. Unless a specific published source is cited, all information stems from interviews and workshops with members of CBDC project teams in each jurisdiction.[3]

The chosen CBDC projects fulfill at least one of the following criteria:

a. *A CBDC is already issued*. Selected project: Central Bank of The Bahamas (CBOB).

b. *A pilot CBDC has been or is being tested involving actual households and firms*. Selected projects: People's Bank of China (PBOC), Eastern Caribbean Central Bank (ECCB),[4] and Banco Central de Uruguay (BCDU).

c. *A CBDC project has been brought onto the country's political agenda and is being analyzed by government or parliamentary bodies outside of the central bank*. Selected project: Sveriges Riksbank.

d. *The central bank has carried out a CBDC project and decided against issuing a CBDC for the time being*. Selected project: Bank of Canada (BOC).

Importantly, these countries have different national contexts and their CBDC projects are at different stages of development (see Box 1 for a quick overview). Thus, the information that central banks can provide differs. Whether or when these projects, except for that of the Bahamas, eventually evolve into an officially launched CBDC offered to the general public remains to be seen.

The structure of this paper is based on the primary considerations for a CBDC project and is summarized graphically in Figure 1. Importantly, all these considerations should be viewed as being carried out with sound processes for risk identification and mitigation.[5]

This paper first explores the policy goals of the different jurisdictions. It then reviews the operational models for CBDC, that is, who issues and distributes CBDC and the respective roles of the central bank and the private sector. The paper then turns to the design features of CBDC, which range from ways to

[1] CBDC is digital money issued by a central bank and is conceivable in both retail and wholesale form. Retail CBDC, or sometimes general purpose CBDC, refers to CBDC that can be held and used by individuals, whereas wholesale CBDCs are available only to a selected set of financial institutions. For more on these different types, see BIS (2018).

[2] For a recent survey of CBDC projects around the world, see Boar and Wehrli (2021). For online resources that are updated continuously, see Atlantic Council (2021) and Kiff (2021).

[3] These central banks have also been given the opportunity to read and comment on the text before publication. Any errors remain the responsibility of the author.

[4] The ECCB is the monetary authority in the Eastern Caribbean Currency Union (ECCU), which is a monetary union consisting of Anguilla, Antigua and Barbuda, Dominica, Grenada, Montserrat, St. Kitts and Nevis, Saint Lucia, and St. Vincent and the Grenadines.

[5] For a discussion of different risks that a central bank needs to consider in a CBDC project, see Kiff and others (2020) and Khan and Malaika (2021).

BOX 1. Current Status of CBDC Projects

- *CBOB, Sand Dollar:* The Sand Dollar was officially launched in October 2020. In late 2021, there were around 20,000 active Sand Dollar wallets in a population of about 400,000, and functions are continuously being developed.
- *BOC:* The BOC has not found a pressing case for a digital currency given the present state of the Canadian payments system. However, it continues to build the technical capacity to issue a CBDC, and monitor developments that could increase its urgency.
- *PBOC, e-CNY:* No formal decision has been taken to launch the e-CNY. The PBOC runs a pilot in parallel in different regions. By October 2021, there were over 123 million e-CNY wallets registered with individuals and about 9.2 million wallets held by firms—a rapid increase from approximately six million active e-CNY wallets in April 2021. In a population of nearly one and a half billion, the share of e-CNY users is now approaching 10 percent.
- *ECCB, DCash:* No decision has been made to formally issue DCash. In March 2021, the ECCB launched a pilot program to successively extend DCash throughout the countries of the Eastern Caribbean Currency Union (ECCU) and run the program for 12 months. Given its rapid adoption, ECCB is now considering transitioning to an official CBDC launch.
- *Sveriges Riksbank, e-krona:* No decision has been made to issue the e-krona. The Riksbank has developed a proof of concept and is exploring technological and policy angles of CBDC. A government inquiry is investigating the role of the state in the digital payments system, including the potential role of a CBDC.
- *BCDU, e-peso:* After ending a pilot in 2018, the BCDU has changed leadership and has opted to not pursue a second pilot due to other priorities and a lack of resources. Potentially, a second pilot will be launched in the future.

Figure 1. The Main Choices and Considerations for a Central Bank Digital Currency Project

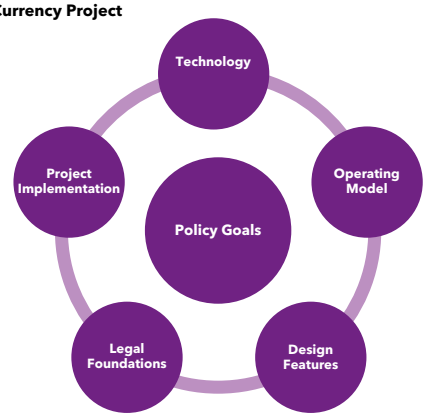

Source: IMF staff.

mitigate risks to uses in cross-border payments. Next , the paper considers options available to jurisdictions on specific technologies and moves to the legal foundations of CBDC. The last section examines the process of exploring and testing CBDC, such as organizational choices and staffing. It also includes insights identified as particularly important by the jurisdictions themselves on the way forward.

2. Policy Goals of CBDC Projects

Policy goals for CBDC naturally guide the ensuing exploration and work. These goals also help establish guidelines to make design and technology choices.

The goals differ across jurisdictions, reflecting factors like the characteristics of the payment systems and various perceived domestic challenges. Mandates may also be a consideration. Central bank laws often establish the function of promoting efficient, safe, and secure payment systems, or set efficient and effective monetary policy, both of which may be relevant to CBDC.

However, the themes of modernizing and/or future-proofing countries' payment systems ran across the various goals stated by the central banks reviewed in this paper. Modernizing is about improving the payment system through increasing digitalization. And future-proofing refers to updating a payment system that is already extensively digitalized to counter potential future risks associated with continuous innovation.

This section discusses the different policy goals that each jurisdiction identified as crucial. Other goals may also exist and be important, but of lesser priority.

A. Financial Inclusion

Financial inclusion is a common policy goal for CBDC projects. Financial inclusion entails access to appropriate and affordable financial services and is associated with poverty reduction worldwide.[6] But despite significant progress, large parts of the world's population remain financially underserved. Increasing financial inclusion has many challenges, including access to digital technology. CBDC could potentially facilitate financial inclusion by increasing access to digital payments and thus serving as a gateway to wider access to financial services.

Most of the six jurisdictions in this survey identify financial inclusion as a top policy goal. In the Bahamas, pockets of the population are excluded from financial services because they live in regions where it is not profitable for commercial actors to operate. Approximately 20 percent of the adult population is estimated to have no bank account.[7] Geography exacerbates the problem since the Bahamas consists of many islands, which are costly to serve.

Likewise, the ECCU consists of island nations where it has been difficult for financial institutions to develop economies of scale and find profitable channels of expansion. Foreign banks have increasingly withdrawn from the region, citing low profitability. The result is lower financial inclusion.

Uruguay has also seen a sluggish development of financial services for a significant part of the population. The government has actively sought to stimulate its development, including by making a digital option mandatory for essential payments.[8]

While China has made rapid progress in financial inclusion and digitalization, the population in remote regions remain underbanked and underserviced by mobile payment operators. The PBOC has sought to promote digital payments and financial inclusion for two decades, but estimates that around 10 percent of the Chinese population still lack access to basic financial services.[9] Meanwhile, some financial institutions that focus on local business have difficulties in digitalizing due to their technological capabilities. PBOC now sees extending financial inclusion to this part of the population as a key policy goal for the e-CNY.

6 For an overview of financial inclusion, see World Bank (2021), Ozili (2020), and Dev (2006).

7 IMF (2019), p.13.

8 The Financial Inclusion Law was enacted in 2014.

9 For more on financial inclusion in China, see World Bank and PBOC (2018).

B. Access to Payments

Helping facilitate payments among the population is an important objective for central banks in most countries.[10] Access to payments is associated with, but not identical to, financial inclusion. Even countries with high levels of financial inclusion, such as Sweden, can still face access to payments challenges. Some central banks are concerned that private payment service providers might not find extending services to all parts of the population sufficiently profitable, and that a declining use of cash will exacerbate the problem. Some jurisdictions are therefore exploring if a CBDC could help achieve or safeguard universal access to payments.

Access to payments may encounter multiple hurdles, including shortage of cash, firms' refusal to accept cash, and lack of or recurring disturbances of digital infrastructure. In the Bahamas, for example, the island geography creates difficulties in both distributing cash and extending digital infrastructure. This is why the CBOB has listed access to payments—regardless of age, social status, or location—as one of its most important goals.[11]

In countries in which cash usage is dwindling, access to payment is also a key concern. Some segments of the population still rely on, or prefer, making cash payments, but may run into limitations. One of the Riksbank's top priority goals for the e-krona project is to ensure broad access to payments in the years ahead.[12] In particular, the Riksbank has identified the elderly and groups with certain disabilities as potentially adversely affected in a cashless society. While the Riksbank is committed to ensuring that cash will still be available and possible to use in the future,[13] it is also exploring how CBDC could facilitate the creation of digital payments especially suitable for these groups as a complement to cash.

The BOC also emphasizes access to payments as a key policy goal despite near-universal financial inclusion. If cash availability falls beneath a certain level, some groups might experience difficulties in making payments. These groups include individuals in remote areas where private firms find it unprofitable to operate, with low income, and with different forms of impairments.[14] A potential CBDC could hence be designed with universal access in mind.[15]

C. Making Payments More Efficient

In countries where cash and check use is high, operational costs are elevated. And in some countries, existing digital payments are also relatively expensive. CBDC is therefore a potential policy tool to offer digital forms of payments that are cheaper to operate. The non-profit nature of central banks means that they could potentially offer low-cost payments as a public good, potentially subject to the need to eventually recover costs.

The Bahamas and the ECCU are high-cost jurisdictions for both physical and existing digital payments. In the Bahamas, an important additional consideration has been the high cost for government agencies to make cash-based payments to citizens who lack bank accounts. There are plans to integrate government agencies in the Sand Dollar network to support digital government payments to individuals to lower this cost.[16]

10 For a discussion on the general role of central banks in payments, see BIS (2003).

11 CBOB (2019).

12 Sveriges Riksbank (2018).

13 The Riksbank is also analyzing legal forms to strengthen cash. See Sveriges Riksbank (2021a).

14 BOC (2020), p.7.

15 Miedema and others (2020).

16 CBOB (2019).

While the Chinese payments market in urban areas is already highly digitalized, the PBOC has expressed a desire to improve its payment services. It sees this as part of an ongoing international effort by central banks to improve their services to the public, comparable to the roll-out of instant payments platforms.

D. Ensuring Resilience of Payments

Ensuring the ability to pay and extending government transfers to individuals under severe circumstances is important for all jurisdictions, but the urgency of this policy goal is especially high in disaster-prone nations. For the Bahamas and the ECCU, resilience is thus considered a key policy goal. Both consist of islands in a region where natural disasters are frequent. Destruction of physical, financial infrastructure and impediments to shipping cash are immediate concerns. In the Bahamas, a hurricane in 2019 precipitated the start of the Sand Dollar pilot in the same year to facilitate assistance payments to and within afflicted areas.

Likewise, the ECCB accelerated the expansion of its DCash pilot to areas affected by a volcano eruption in St. Vincent and the Grenadines in 2021.

Countries with a highly digitalized payment sector are concerned about disruption to digital services and concentration risks where there are only a few large operators. In China, for example, the mobile payment market is dominated by two firms, AliPay and TenPay/WeChat Pay. The PBOC has expressed concern that the failure of such firms could have serious consequences to the Chinese payments system. One of the crucial policy goals expressed by the PBOC is for the e-CNY to function as a backup to existing digital payment solutions.

Similarly, the Riksbank has identified single points of failure among a few dominant actors as a potential risk that would be exacerbated in a society in which cash is no longer available as a backup or "redundancy" system. The resilience of payments has also become an important part of the country's ongoing modernization of civil defense.[17] While the Riksbank advocates the continued existence of cash, the e-krona could potentially serve as an additional backup to existing forms of digital payments.

The BOC has also noted that cash can function as a backup when digital payments are unfunctional, and that falling cash usage might thus mean impaired payments resilience. CBDC could therefore potentially play a role as an additional backup.[18]

E. Reducing Illicit Use of Money

Some features of cash, including anonymity and the lack of an audit trail,[19] make it attractive for illicit transactions (for example, tax evasion, money laundering, and terrorist financing). CBDC could potentially reduce this problem.

At this point, however, only the Bahamas has reduction of the illicit use of money as a top policy objective for its CBDC. The background to this objective is an ongoing campaign to strengthen the Anti-Money Laundering / Combating Financing of Terrorism (AML/CTF). The Bahamas was put on the Financial Action Task Force (FATF) *grey list* in 2018 due to strategic deficiencies in its AML/CFT framework, which resulted in increased monitoring. The Bahamian authorities subsequently implemented an action plan aimed at addressing the identified deficiencies, and as a result, the Bahamas was de-listed in December 2020.[20]

17 Utredningen om civilt försvar (2021).

18 BOC (2020), Miedema and others (2020).

19 FATF (2015).

20 FATF (2020).

F. Monetary Sovereignty

While currency substitution has long been a risk facing countries, it is possible that new forms of digital currency might have a competitive advantage relative to older forms of currencies. If a sufficiently large portion of a country's population adopts a foreign digital currency or a global stablecoin, the ability of the country to carry out several crucial central bank functions might be impaired, such as monetary policy and lender of last resort.[21]

The BOC has stated that serious consideration of a CBDC might be triggered if monetary sovereignty were to become an issue—say if Canadians began adopting a non-Canadian digital currency or stablecoin.[22] Likewise, the PBOC has said that one motivation for investigating CBDC was to secure monetary sovereignty in a digital future.[23]

G. Competition

CBDC could potentially increase competition in a country's payments sector in two ways: directly, by competing with existing forms of payments; and indirectly, should the CBDC be designed as a platform open to private payment service providers (see the following section, Operating Model). The latter would ensure low barriers of entry for new firms seeking to offer new payment services.

The Riksbank, in particular, sees competition as a potentially important contribution by the e-krona. The payments market, according to Riksbank analysis, displays clear network effects that tend to favor the concentration of a few large actors. This may lead to high fees or stagnating innovation in the future. The e-krona could be a way to ensure more competition and enhance market efficiency.[24]

The BOC has also said that the high concentration of service providers in the Canadian financial system may be contributing to the high costs of payments. If cash were to decline significantly, competition in the Canadian payments market would decline even more. This is among the reasons why the BOC is monitoring cash usage and building capacity to launch a potential CBDC.

H. Summary of Policy Goals

The different policy goals of the jurisdictions in this survey are summarized in Table 1.

Table 1. Jurisdictions' Stated Policy Goals of Central Bank Digital Currency

Country	Financial Inclusion	Access	Efficiency	Illicit Use of Money	Resilience	Sovereignty	Competition
Bahamas	✓	✓	✓	✓	✓		
Canada		✓			✓	✓	✓
China	✓	✓	✓		✓	✓	✓
ECCU	✓		✓		✓		
Sweden		✓	✓		✓		✓
Uruguay	✓		✓				

Sources: Central banks.
Note: ECCU = Eastern Caribbean Currency Union.

21 For example, IMF (2020a) and BIS and others (2021).

22 BOC (2020).

23 For example, Mu (2021).

24 Sveriges Riksbank (2017, 2018); Soderberg (2019).

3. Operating Model

A crucial choice is how CBDC will be issued and circulated, and what the role of the central bank and the private sector will be. We refer to this overarching structure as the operating model.[25] Different names and classifications are used in other literature and there is no established standard for the typology of different operating models.[26]

In the first model, which we call *unilateral CBDC*, the central bank carries out all functions in the payments system, from issuing the CBDC to distributing it, and interacting with end-users.

The second model entails issuance by the central bank, but includes a role for private sector firms to interact with the end-user. We refer to these agents as *intermediaries* and the model in which they operate *intermediated CBDC*. The intermediary role can be filled by financial firms, but also other types of companies such as payment service providers and mobile phone operators. Most would likely be privately-owned and for-profit firms, but state-owned intermediaries and cooperatives may also be involved.[27] This second model would require the central bank to regulate and/or oversee other actors, which adds an extra layer of legal and operational complexity.[28]

In the third model, digital currency is issued not by the central bank but by private firms that back the issuance by holding central bank liabilties. Hence, the third model is not a CBDC, but rather a stablecoin, or a special type of e-money, as it is not issued by a central bank and may be referred to as *synthetic CBDC* or *sCBDC*. But as it is backed one-to-one by central bank-issued assets, it may be considered by some central banks as an alternative to CBDC, and is therefore included in this paper.

These conceptual models should not be seen as mutually exclusive. Some central banks are considering the intermediated model as their main operating model, but also offering basic payment services through a unilateral model to ensure universal access and resilience. Likewise, an sCBDC is not necessarily a replacement for CBDC and could, for instance, be issued by private firms alongside, or even backed by, CBDC.[29]

These three conceptual models are depicted in Figure 2.

These conceptual operating models are useful starting points for discussions on CBDC design. So far, there is a convergence on the intermediated model. No central bank in this survey has explored the unilateral or synthetic CBDC models and the rest of this section will focus on the intermediated model.[30]

25 This follows the concept used in Kiff and others (2020).

26 For example, see Armelius and others (2020), BOE (2020), and Auer and Böhme (2020).

27 For a discussion on how state-owned enterprises have developed in the recent decades, see Bruton and others (2014).

28 For a discussion, see Kiff and others (2020).

29 See Adrian and Mancini-Griffoli (2019, 2021), Auer and Böhme (2020), and Auer and others (2021).

30 Recently, however, the Hong Kong Monetary Authority issued a whitepaper outlining an approach similar to sCBDC. Cash in Hong Kong SAR is currently also mainly issued by private institutes rather than the monetary authority. For more on this, see HKMA (2021).

Figure 2. Three Conceptual CBDC Operating Models

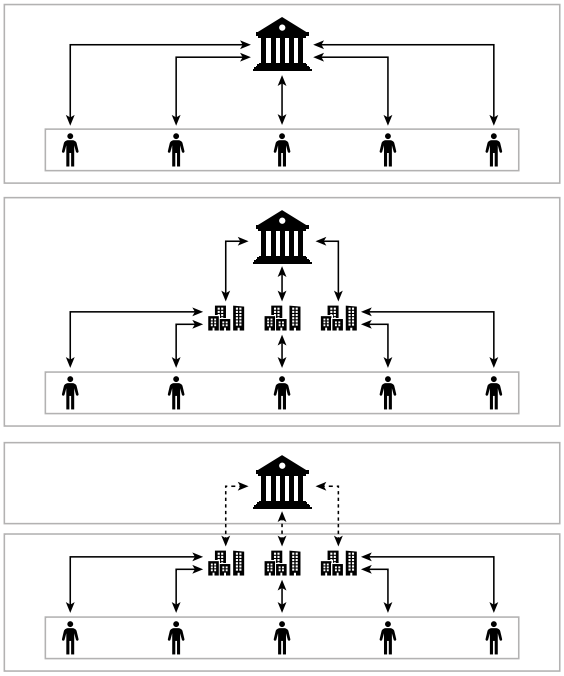

Unilateral CBDC

Central bank issues money and performs all functions, including direct interaction with end users

Intermediated CBDC

Central bank issues money, but delegates functions to non-central bank intermediaries who interact with end users

Synthetic CBDC

Non-central bank actors issue money that is backed by central bank assets that they acquire from the central bank (dashed line)

Source: IMF staff.
Note: CBDC = central bank digital currency.

A. Central Bank and Private Sector Functions

The intermediated model can take different forms depending on how functions are distributed between the central bank and private intermediaries, as illustrated in Figure 3.

When discussing the distribution of functions between actors, it is useful to distinguish between the owner of the technical system necessary to carry out a specific function and the executor of the function itself. These are not always the same; indeed the system might be owned by the central bank while the

Figure 3. Functions to Be Carried Out in a CBDC Environment

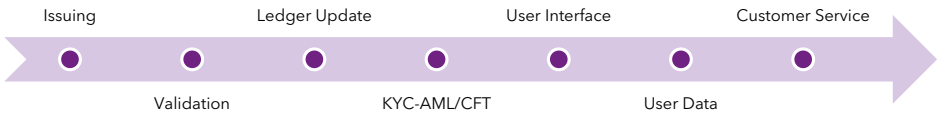

Source: IMF staff.
Note: AML/CFT = anti–money laundering/combating the financing of terrorism; KYC = know your customer.

function is carried out by a private company.[31] For instance, the IT system that intermediaries use to monitor users of the Sand Dollar relating to AML/CFT is owned by the CBOB, which ensures a standardized approach. As mentioned above, delegating functions to private actors still requires regulation and monitoring.

Issuing is obviously a crucial function of all types of money. As discussed above, all central banks in the study currently explore models in which the CBDC are their own liability, just like cash. It is possible, though, for central banks to let a private company own the technical systems that enable CBDC issuance. For instance, in the Uruguay e-peso pilot, a private vendor owned and operated a technical system that converted pesos created by the central bank into e-pesos, effectively making the issuance of e-peso into a two-stage process.[32]

Validation refers to validating a transaction. The concept is often associated with the validation that takes place in a distributed ledger technology (DLT) network but can also refer to more traditional processes including checking the user's identity, the authenticity of money, and the availability of funds.[33] In some cases, these functions are divided between the central bank and private entities. For instance, in the e-krona proof of concept, the central bank owns and operates a notary node that ensures money has not been spent before, while private intermediaries carry out remaining validations, such as checking the authenticity of e-kronas.

Each transaction entails a *ledger update* when users transfer holdings of CBDC between each other. A ledger is a database of records of monetary holdings but can be either centralized or distributed across a network. Updating the ledger means updating the records of CBDC balances after payments have been made. The centralized ledger, owned and updated by a single entity, is still the standard approach among central banks, whereas DLT is a potential new approach.

In the case of DLT, at least three alternatives exist. First, the central bank owns the infrastructure of the entire ledger and updates it (for example, the Bahamas Sand Dollar). Second, the central bank owns the ledger, but private intermediaries update it. And third, a private intermediary owns part of the ledger and updates that same part of the ledger, conditional on the central bank's approval (in the Swedish e-krona proof of concept, the intermediary can update the ledger after the central bank's notary node has checked that no double spending has taken place).

The remaining functions of Figure 3 concern the interaction with the end-users. KYC-AML/CFT refers to the process of implementing Know Your Customer (KYC) and AML/CFT requirements (for example customer due diligence measures) aimed at combating illicit flows. *User interface* denotes the means through which users can interact with and/or pay with their CBDC holdings, such as through applications on mobile phones. *User data* refers to the function of handling the personal data of users, and *customer service* captures the process of helping users connect to the CBDC, handling errors, and solving other issues.

Table 2 summarizes how functions are allocated between central banks and private sector actors in the six CBDC projects.[34] In some cases the distribution of functions may change when the project progresses from pilot to formal launch. For instance, the ECCB currently offers private intermediaries a ready-made application for users to interact with DCash wallets, but states that after a formal launch, intermediaries would develop their own user interfaces.

[31] Importantly, systems ownership does not preclude outsourcing of certain aspects but it entails responsibility for the overall development, maintenance, and functionality of the system, including outsourced services. For a definition of system owner, see NIST (2021).

[32] For more on potential operational risks of outsourcing central bank activities, see Kiff and others (2020)

[33] For an overview of DLT in payments, see BIS (2017) and Shabsigh and others (2020). DLT can be permissionless, meaning that anyone can join the network and partake in performing crucial functions, or permissioned, meaning there are strict requirements to joining the network. All central banks that are currently exploring DLT are focusing on the permissioned variant. For more on this, see Natarajan and others (2017).

[34] Private sector here includes state-owned enterprises and cooperatives. See Bruton and others (2014) for a discussion.

Table 2. The Distribution of CBDC Functions Between the Central Bank and the Private Sector

	Issuing		Validation		Ledger Update		KYC-AML/CFT		User Interface		User Data[1]		Customer Service	
	Owner	Executor	Owner	Executor	Owner	Executor	Owner	Executor	Owner	Executor	Owner	Executor	Owner	Executor
Bahamas														
Canada														
China														
ECCU														
Sweden														
Uruguay														

Color scheme: Central Bank Both Private Still Exploring
Source: Central banks and IMF staff.
Note: AML/CFT = anti–money laundering/combating the financing of terrorism; KYC = know your customer.
[1]One option is to grant central banks access to data stored by intermediaries. This is the practice in China.

B. The Business Model of CBDC

The business model of CBDC is a key concern for both private firms and the central bank. If private firms are expected to carry out a function in the CBDC ecosystem, they will have to make a profit at least in the medium term. Similarly, though central banks are not-for-profit organizations, they will need to decide whether to seek cost-recovery for their expenditures of building the CBDC system. Central banks may also decide to subsidize the use of CBDC to increase adoption if supported by a particular policy goal.

Among the CBDC projects reviewed in this paper, there is almost universal consensus that the main business model for private intermediaries is fees on payments. The role of the central bank is seen as providing a free or low-cost platform on which private intermediaries can operate. None of the central banks favor allowing private intermediaries to gather payments data, which could be used for commercial purposes. The PBOC notes it does not charge intermediaries or users, and intermediaries cannot charge individual users in the e-CNY project. However, intermediaries have the choice of charging merchants. The PBOC views this as a substantial incentive for firms to enter the market, and keeping fees in check.

The BOC states the choice of business model is complex. One possibility is for the central bank itself to provide a basic CBDC payment function to the public, possibly but not necessarily charging a fee for using it. The Riksbank is also considering this approach.

The question of whether a central bank should charge intermediaries for using the CBDC system is also connected to the question of whether it anticipates recovering its development costs. There is a risk that if central banks collect fees, intermediaries will in turn pass the cost downstream and raise the price of payments, which may counter initial policy goals. The question of whether and how to cover costs remains an open question, and the BOC states this as one of its most important areas of research.

Staff at the Riksbank also state that charging intermediaries fees is difficult because of the current regulatory framework. Another issue is that charging fees would possibly contradict its commitment to offering payments as a public good. Revenues for the central bank would likely solely be in the form of seignorage.

While subsidizing the adoption of CBDC is currently not seen as a viable path, the Riksbank is discussing if it might subsidize the cost of developing certain functions that the private sector would not find profitable. Examples of this include increasing payments resilience and developing payment solutions for minorities.

4. Design Features

CBDCs can be designed in different ways with different characteristics and functions. We refer to these characteristics and functions as *design features*. Design features are more specific than the operating model, and CBDCs with the same operating model can still differ in their design features. CBDC designs are generally intended to support policy goals or mitigate risks that could arise from issuing CBDC. This section is divided by topic and ends by summing up design features of individual projects in Table 3.

A. Restrictions Aimed at Ensuring Financial Stability

Central banks engaged in CBDC projects have committed to not jeopardizing financial stability and avoiding any sudden shifts to the structure of the financial system.[35] The literature discusses the potential risk that the introduction of CBDC could create, including crowding out banks and facilitating bank runs.[36] In addition, the literature discusses different ways to mitigate these risks by either restricting CBDC balances or taxing the use or balances of CBDC above a threshold.[37]

All CBDCs that are currently circulating, either as official currency or through a pilot, are designed with restrictions that limit the competitiveness of CBDC versus bank deposits. At the time of writing, however, only three of the six central banks in this study—the central banks of the Bahamas, China, and the Eastern Caribbean—have circulating CBDC. Other central banks are still analyzing these questions conceptually.

Limits on CBDC fall under two main categories: restrictions on remuneration of CBDC and quantitative restrictions on holdings and transactions of CBDC.

Restrictions on the Remuneration of CBDC

The Bahamas, China, and the ECCU currently do not pay interest on CBDC holdings. In all three cases, the reason is to limit CBDC competition with bank deposits. If there is no interest, CBDC can still be attractive as a means of payment, even while its attractiveness as a store of value (savings instrument) diminishes.[38]

There is a potential policy trade-off between limiting competition with bank deposits and ensuring an effective transmission mechanism of monetary policy. Staff at the Riksbank, reflecting the past Swedish experience of low interest rates, point out that a zero percent interest rate on CBDC could limit the ability to carry out a negative interest rate monetary policy.[39] Also, the attractiveness of bank deposits versus CBDC would shrink with lower policy rates. A possible solution is a CBDC with an interest rate that is consistently lower than the policy rate. The Riksbank is investigating the legal issues related to paying interest rates (whether positive or negative) on CBDC.[40]

An alternative discussed in the literature is to impose fees on transactions above a certain threshold, but so far none of the CBDC projects have tried this.

Quantitative Restrictions

All three active CBDC projects were designed with quantitative restrictions. The goal is to explicitly limit competition with bank deposits but also to foster financial inclusion. To lower the threshold to onboard new

35 Central Banks and BIS (2020), G7 (2021).

36 For example, see Kumhof and Noone (2018), Juks (2018), and Bindseil (2019).

37 In particular, see Bindseil (2020).

38 For more on this, see Agur and others (2022).

39 For example, see Armelius and others (2018) and Armelius and others (2020).

40 For a legal discussion on interest rates and CBDC, see Bossu and others (2020).

users, small CBDC holdings are allowed without the need for identification or other KYC procedures (see more in the section on Anonymity).

Without special arrangements, it is not possible to send money to a wallet that has reached its specified limit, and the sender will typically receive an error message when trying to do so. CBDC holdings may also be connected to a bank account to which excess holdings of CBDC may automatically be transferred. Such a function is currently under development in the Bahamas Sand Dollar.

Limits to CBDC balances may also help limit pilot programs. Uruguay limited both the number of users and the amount of total e-pesos each user could hold. This made the pilot more manageable, but also lowered risks of disruptions and to the reputation of the central bank (were something to go wrong).

The ECCB DCash pilot also plans to limit the total amount of DCash that can be created. For now, however, DCash is offered to meet demand. No decision has been made on when to set the limit or how high it should be, in part to meet unexpectedly high demand due to the COVID-19 pandemic and support the economic recovery.

The Riksbank is exploring different technological options that could allow quantitative restrictions, and has tested a payments card which carries a limited amount of e-krona.

B. Anonymity

Anonymity is one of the key traits of cash, and the rise of digital payments threatens the lawful or legitimate preference for anonymity by certain segments of the public or for certain purposes—such as buying a present for one's spouse. Anonymity is also connected to financial inclusion: non-anonymous payment services often require forms of identifications that can be difficult or costly to obtain.

However, anonymity can also be used for illicit purposes and can undermine AML/CFT measures. Anonymity, therefore, poses a policy trade-off—the more anonymity, the larger the risk for illicit use.

All three active CBDC projects have chosen the same way to handle the policy trade-off between anonymity/financial inclusion and AML/CFT compliance. Their approach has been to provide a tiered selection of wallets with different levels of thresholds. Those with lower thresholds allow for greater anonymity. As a result, CBDC can more easily be rolled out into rural or disadvantaged areas where virtual identification can be difficult.

The use of tiered CBDC wallets thus gives rise to "policy synergies" between anonymity, risk-reduction (of bank runs), and financial inclusion.

C. Off-Line Capacity

The ability to pay when not connected to main telecommunication systems is important to increase resilience in crisis situations, such as during natural disasters and armed conflicts. Off-line capacity is hence linked to the policy goal of resilience and is especially important in disaster-prone or geopolitically tense areas. The PBOC also emphasizes that off-line functionality is important in areas with patchy telecom access, which also often correspond to areas of low financial inclusion.

In practice, off-line functionality has turned out to be complicated technologically.[41] Further, the exact definition of off-line transactions differs. Often the term refers to being off-line from the Internet but still reliant on a local network, such as Bluetooth. But some events such as prolonged black-outs or electromagnetic disturbances might affect local networks as well.

The Bahamas considers off-line functionality to be vitally important but has encountered difficulties in achieving it. The pilot revealed that the planned solution of local off-line networks—built on introducing local redundancies to the main telecommunication system—did not fully achieve the policy goal. The

[41] For a discussion, see Chohan (2021), Armelius and others (2021), and ECB (2020).

telecommunication towers required in the solution are vulnerable to the same weather conditions as the main telecommunication system. Also, the geographical reach of the local networks is limited, which makes it difficult to make payments between islands. Presently, the CBOB is working with its main contractor to identify alternative solutions. Staff at the CBOB state that the decision to explore alternative off-line solutions was the most significant change that resulted from the practical experience of the pilot.[42]

The PBOC has tested different solutions, and reports that sufficiently safe and efficient off-line payments are now in place. These include hardware-based e-CNY wallets placed inside mobile phones, or held as cards that can make payments to another mobile phone wallet in physical proximity without Internet access. To reduce the effects of illicit tampering with the devices, which could lead to double spending and counterfeiting, each user can only perform a limited number of off-line payments before needing to go back online to access the main ledger. In addition, offline payments in e-CNY involve a variety of technologies, including digital signature and encrypted storage to further reduce risks.

The team working on the Swedish e-krona proof of concept has identified a number of potential solutions to establish offline functionality and is proceeding to test these. Participants have identified several challenges, such as how to prevent double spending and ensure the authenticity of e-kronas while off-line.[43]

D. Cross-Border Payments Using CBDC

Central banks and international organizations are increasingly evaluating the use of CBDCs to enhance the efficiency of cross-border payments, which is generally considered costly and inefficient.[44] The G20 has instigated an ongoing collaboration between international organizations and national central banks to explore ways to further this goal, including through CBDCs.[45]

Retail CBDC projects are carried out primarily with domestic purposes in mind, at least so far. Nonetheless, discussions on how CBDC could potentially be used in cross-border payments are ongoing. And technical experiments on wholesale CBDC for cross-border usage have been conducted for several years.[46] Adverse macroeconomic implications, such as increased currency substitution and vulnerability to financial shocks, are possible as retail CBDC become available across borders.[47] These potential risks, and means of mitigating them, are being discussed in the context of the G20 roadmap to enhance cross-border payments.

The six jurisdictions in this study are exploring cross-border issues carefully but largely on the side of their domestic considerations. Canada, China, and Sweden are represented in the Future of Payments Working Group, which stems from the G20 roadmap. In addition, the PBOC is also exploring how a retail CBDC, such as the e-CNY, can be used for cross-border payments, and has partnered with the BIS Innovation Hub and other national central banks in the multi-CBDC Bridge project—an experimental CBDC arrangement leveraging DLT to facilitate cross-border payments.[48]

The PBOC states three principles for their ongoing work on cross-border payments for CBDC. The first is the principle of "no disruption," which in practice means to avoid negative spillovers on the Chinese economy and that of other nations, such as significant currency substitution. The second rule is that any CBDC cross-border payments system must be compliant with the rules and regulations of all connected countries, including capital flow management measures. In addition, according to the PBOC, information

[42] The tested solution relied on smaller local networks that were meant to act as backups to the main telecommunication system rather than independently of telecommunications.

[43] See also Sveriges Riksbank (2021b).

[44] BIS and others (2021).

[45] FSB (2020).

[46] For example, see BOC and others (2018) and BOT and HKMA (2020).

[47] For more on potential macrofinancial implications of cross-border use of CBDC, see IMF (2020a) and BIS and others (2021).

[48] BISIH and others (2021).

flows between countries should be improved to help authorities counter illicit use of money, including tax evasion. According to the third rule, cross-border payments should involve interoperability across domestic CBDCs or between domestic CBCDs and incumbent payment systems rather than a single CBDC used for transactions on both sides of the border. The PBOC thus prefers a system where domestic CBDCs are converted to other currencies as payments cross borders.

The Bahamas does not currently allow the Sand Dollar to be used outside its borders. It has stated that the Sand Dollar is exclusively intended for domestic purposes and that cross-border payments must take place through commercial banks in traditional non-CBDC Bahamian dollars. However, foreigners can own and pay with Sand Dollars when visiting the Bahamas after registering for an account with a low limit on both balances and monthly transactions.[49] Nevertheless, the central bank is planning to explore cross-border functionality for the Sand Dollar within the next three years.

Staff at the ECCB look favorably to using CBDC for cross-border payments, given the importance of trade and overseas remittances for the countries in the ECCU. The ECCB has begun discussions with other regional central banks regarding the interoperability with legacy payment systems and platforms to enable remittances and trade in the region. At present, however, the main priority is ensuring that DCash works for domestic purposes. That being said, as the ECCU consists of eight different nations, the DCash technically represents the first trial of a single CBDC used for cross-border payments, although within a monetary union and with the same central bank.

Staff at the six central banks have raised the following main hurdles to using CBDC across borders:

- *Technical interoperability:* The lack of coordination on technology and messaging standards in initial stages of development could imply that retrofitting CBDC for cross-border use will be costly and complex. Collaboration on the G20 roadmap may help, while decentralized forms of compatibility between different DLT systems may also be promising.[50]
- *Legal and regulatory harmonization:* At present, all the jurisdictions have carried out their legal investigations based on their domestic legal systems. However, some harmonization may be needed regarding the treatment of data and privacy, tax and payments laws, and capital flow management measures.

E. Summary of Design Features

Table 3 summarizes the design features under consideration or deployed by the six central banks.

Table 3. Design Features of CBDC Projects

	Carry Interest or Not	Quantitative Restrictions	Anonymity	Offline	Cross-Border Payments
Bahamas	No	Yes	For lower tier	Yes/exploring	Future project
Canada	Undecided	Undecided	Undecided	Exploring	International collaboration
China	No	Yes	For lower tier	Yes	Experimenting/international collaboration
ECCU	No	Yes	For lower tier	No	Future project
Sweden	Undecided	Exploring	Undecided	Exploring	International collaboration
Uruguay	No	Yes	Yes, but traceble	No	Possible future project

Source: Central bank staff and published sources.

[49] It is possible, however, to integrate a Sand Dollar account with a bank account so that CBDC is exchanged into commercial bank money before making the cross-border payment.

[50] BIS and others (2021), Herlihy (2018).

5. Technology

CBDCs rely on technology, which must be appropriately selected to operationalize the policy goals discussed earlier. Even in an intermediated CBDC model, the central bank must build a core system for issuing CBDC and processing transactions. One of the great difficulties is making decisions while much of the technology is still developing and remains relatively untested. Central banks must decide where to acquire technology, if they do not build it in-house, and which technology best suits their purposes.

A. Technology Suppliers

A central bank typically needs to acquire technology from or partner with external vendors to develop proprietary solutions. So far, there are two main approaches to the technology supply question. The first is to choose a main contractor that supplies the technology and collaborates with the central bank to develop the CBDC. This "CBDC package solution" was chosen by the Bahamas (NZIA), the ECCB (Bitt), Sweden (Accenture), and Uruguay (Roberto Giori). In the case of Sweden and Uruguay, however, contractors were used to deliver a specific test solution and thus would not necessarily be relied upon to further develop, and potentially launch, CBDC.

In the second approach, the central bank relies to a greater extent on internal resources and has different contractors for different areas as necessary. This approach tends to be more onerous for the central bank in terms of internal capacity and resources, but also offers more control over the development process. Canada and China have chosen this path. The BCDU indicates that a second Uruguay CBDC pilot may likely be based on this approach to avoid relying on a single vendor.

Intermediaries can also be selected as development partners. The PBOC, for instance, has partnered with specific e-CNY intermediaries to develop payments solutions and functions that have been added to the e-CNY ecosystem.

B. Distributed Ledger Technology vs Centralized Technology

Distributed ledger technology (DLT), the best known of which is blockchain, has in recent years emerged as a promising alternative to technologies that are based on centralized ledgers. Central bankers are therefore faced with another technology choice.[51] The choice is particularly difficult as DLT is still developing, and its capacity and suitability are being explored. Some pilots and proofs of concept are therefore testing DLT without necessarily expecting to select it for further development.

The experiences so far suggest that there is no universal case for DLT as the primary engine of CBDC, and jurisdictions have different views on the potential merits of the technology. The Bahamas and the ECCB have DLT-based systems, and staff from both central banks cite the security of the technology as valuable for their needs.

The PBOC, on the other hand, has tested DLT during its pilots and decided that its capacity to process transactions and store data does not meet its requirements. It is particularly concerned about e-CNY's ability to handle days with very high levels of transactions, such as the "Singles Day" (November 11, China's equivalent to Black Friday in the United States).

However, the PBOC has committed to what it refers to as a "hybrid architecture." Thus, DLT is being used in the e-CNY system but only in limited areas where it is deemed to have an advantage over other technologies. Intermediaries can also base their activities on any technology, including DLT, and still function in the e-CNY ecosystem.This openness to different technologies is part of what the PBOC calls a "Long Term

[51] See Kiff and others (2020), Auer and Böhme (2020).

Evolution System," through which new features of technology can continue to be added to the e-CNY even though its core is a centralized ledger.

The e-peso did not rely on DLT, but BCDU staff acknowledge that a potential second e-peso pilot might test the appropriateness of DLT, or a hybrid system that incorporates DLT for particular purposes.

The BOC has not decided on technology but is carrying out multiple technological workstreams, including DLT. Its staff has expressed some skepticism about the suitability of DLT for central bank purposes but acknowledges that DLT can support some important functions. One possibility would be to combine different technologies to achieve different purposes.[52]

The Riksbank is currently exploring a DLT-based proof of concept, but its staff stress that a potential future e-krona does not necessarily have to be built on DLT. A second e-krona proof of concept or pilot could thus be based on a different technology.

C. Summary of Technology Choices

Table 4 sums up the technology choices described above.

Table 4. Summary of Technology Choices

	Main Tech Contractor	DLT
Bahamas	NZIA	Yes
Canada	–	–
China	–	Hybrid
ECCU	Bitt	Yes
Sweden	Accenture	Testing
Uruguay	Roberto Giori	No

Source: Central banks and published sources.
Note: DLT = distributed ledger technology; ECCU = Eastern Caribbean Currency Union.

[52] BOC (2020).

6. Legal Foundations for CBDC

CBDC[53] requires a legal framework that clarifies whether the central bank has the mandate to issue CBDC and what status it would have legally.[54] Existing legal frameworks were typically enacted in a pre-digital age, and investigating CBDC therefore also entails ascertaining whether law reform is necessary to ensure that a CBDC can be issued by the central bank.

The status of the six surveyed jurisdictions is summarized in Figure 4.

Figure 4. Status of Law Reforms in the Six Jurisdictions

Sources: Central banks and IMF staff.
Note: ECCU = Eastern Caribbean Currency Union.

To issue the Sand Dollar, The Bahamas enacted a revised legal framework, the Central Bank of Bahamas Act, in 2020. The currency issuance function is broadly worded, and the definition of "currency" explicitly includes not only banknotes and coins but also "electronic money" issued by the Central Bank.[55] Moreover, the Act specifically grants the Central Bank the power to issue currency in the form of "electronic money."[56] To support this, the Act also grants the Central Bank regulatory powers to prescribe "the framework under which electronic money issued by the Central Bank…may be held or used by the public."[57]

Among the countries that have not yet formally issued a CBDC law, reform is still being investigated and prepared. For example, China is preparing for a general revision on People's Bank of China Law (draft), which suggests that Chinese currency includes both physical and digital forms (e-CNY) and thus confirm the legal tender status of e-CNY.[58] The draft law provides the central bank with the broad power to plan, organize, and supervise the payment system and financial infrastructures. The Central Bank will have responsibility to coordinate the work on national financial security, with the goal of developing a cyber-resilient CBDC. In addition, the draft law explicitly prohibits and imposes fines on the production, sale, and circulation of "illegal CBDC."

The ECCB has prepared a draft amendment to its central bank act. The draft amendment would establish the legal foundation of CBDC by extending the definition of "currency" to "digital currency."[59] Further, it

53 This section was written by Marianne Bechara, Wouter Bossu, and Akihiro Yoshinaga.
54 On the analytical model for assessing those questions, see Bossu and others (2020).
55 Central Bank Act of The Bahamas (2020). Sections 5(1) and 8(1).
56 Central Bank Act of The Bahamas (2020). Section 12(7).
57 Central Bank Act of The Bahamas (2020). Section 15.
58 PBOC (2020). Articles 18 and 19.
59 Eastern Caribbean Central Bank Agreement (Amendment) Order (2020), Article 2.

explicitly attributes legal tender status to digital currency and clarifies the central bank's sole right to issue digital currency.[60]

In Sweden, the legal questions are currently being investigated in a government inquiry launched after a petition was sent to Parliament by the Riksbank in 2019.[61] In parallel, the central bank is actively analyzing whether existing means of payment and legal mechanisms in Sweden would be fit for e-krona operations or whether new types of assets or legal mechanisms should be created by law.

Since it has decided not to issue a CBDC at this time, Canada is not currently looking into law reform. When Uruguay completed its six months e-Peso pilot in 2018, the legal framework was considered sufficient at the time for the central bank to carry out the testing without the need for legal amendments. Such amendments, however, would be necessary for an official roll-out of the e-peso, according to the central bank.

Surveyed central banks flagged several legal challenges (in addition to the challenge of legal harmonization mentioned in the section on cross-border payments) to issuing CBDC, as well as lessons that could be drawn from that process.

"Law follows technology": In many cases, the operating and even legal design of CBDC was initially driven by technological developments, often under the advice of consulting firms. Central banks are therefore recommended to initiate legal reflections very early in the process. This should go hand in hand with building sufficient internal legal capacity in central banks' legal departments.

Understanding the legal nature of CBDC: In many countries, this new form of money poses significant legal challenges under public and private law. In some countries, some fundamental issues still need to be decided, such as the legal nature and ramifications of issuing digital currency (for instance, rights of holders subsequent to the insolvency of authorized providers). Given the many legal complexities, several central banks relied on external counsel to develop the legal-regulatory framework for CBDC. Central banks should consider combining the abovementioned internal capacity building with a needs assessment for external counsel. This could go hand in hand with a close dialogue with financial intermediaries, in particular, to gain insights into for instance how they see CBDC impacting their business models.

Legal tender status: While central banks acknowledge that the technical means to receive CBDC in payment (such as devices or internet access) is not universal in their countries, most of the surveyed central banks nevertheless advocate granting legal tender status to CBDC. This approach could be possible under a fairly "relaxed" legal conception of legal tender status, with ample space for contractual derogations. That said, it is also acknowledged that without wide acceptance and circulation of CBDC, the reputation of the issuing central bank would be at risk. Against this background, a few jurisdictions have started a fundamental debate on the role of legal tender currency.

Flexibility and law reform in preparatory phase vs. final phase: Several central banks indicated that they saw no need for law reform in the pilot phase, but that law reform would be necessary for the final phase (roll-out). Maintaining this type of flexibility during the pilot phase may be useful for other central banks, in particular at a stage where CBDC is not yet issued as an actual liability of the central bank, and central banks may alter fundamental design features subsequent to the pilot.

Specific vs. general law reform: Modifying the central bank law and other laws only to strengthen the legal basis for CBDC issuance may be the fastest route. However, a few central banks chose to anchor these amendments into a broader reform of the central bank's charter to address other legal issues. This was the case in the Bahamas. Whilst such an approach may somewhat slow down the law reform process, it yields the benefit that other aspects of the central bank's legal framework can be strengthened in conjunction. Going forward, countries should assess whether CBDC-related law reform could be an opportunity to introduce other legal amendments.

60 Eastern Caribbean Central Bank Agreement (Amendment) Order (2020), Articles 18(1) and 18(3).

61 SOU (2021).

7. Project Implementation

CBDC projects are generally large undertakings for central banks and need to be organized, staffed, and financed. The availability of resources differs across central bank, as does the importance of CBDC projects relative to other undertakings. Carrying out a pilot further requires planning and execution, supportive institutional structures, and investment in staff education, skills, and retention. A key part of a CBDC project is also to ensure that there is enough staff to identify and monitor operational risks.[62]

This section investigates the different organizational paths taken by this paper's six central banks, the learning curve of staff engaged in CBDC, and the main challenges they have faced.

A. Organizational Changes at the Central Bank

Central banks investigating CBDC must decide whether to make formal organizational changes, or work with existing structures. Some have created new committees, divisions, or research centers.[63] Others have instead reprioritized the work of staff in existing divisions.

At one end of the spectrum, and reflecting the size of its undertaking, the PBOC first set up a specialized work team in 2014, but two years later created a new specialized institute, the Digital Currency Institute of the People's Bank of China (PBCDCI). The PBCDCI has set up subsidiaries across geographical areas to help organize the e-CNY pilots.

Similarly, the CBOB created a new unit devoted to developing the Sand Dollar but under the supervision of a policy steering committee made up of representatives from the different departments of the bank.

The decision to make formal organizational changes can also arise as work on CBDC advances. The BOC carried out a substantial part of its CBDC analysis by drawing on the resources of two departments—coordinated by a fintech senior officer—neither of which were exclusively devoted to CBDC. Then, in 2020, after presenting its official position on CBDC (see Box 1), the BOC formed a research team to investigate technology that would help build capacity for a successful CBDC launch; the team was also tasked with monitoring conditions that could trigger the need to proceed.

Likewise, the Riksbank started its e-krona project with a project team consisting of members from different departments. However, a new division was created in 2019 devoted specifically to developing its e-krona proof of concept. CBDC policy analysis, however, remains part of the payments department's general policy work. The two divisions work closely together.

In contrast, the ECCB has not initiated any changes in its organizational structure and instead draws personnel from across different departments at the central bank to form an internal working group. Similarly, the BCDU did not initiate any organizational changes while conducting the e-peso pilot.

B. Internal Staffing

The number of staff at central banks involved in CBDC projects varies mainly with the degree of outsourcing to private vendors. Another important factor is the size of the pilots undertaken by the central banks. For instance, the staff working on the Chinese e-CNY project grew from around 40 to around 300. Importantly, this number does not include private-sector employees that have been working in collaboration with the PBOC.

For the CBOB and the ECCB, both operating in smaller countries and teaming up with a main contractor, the numbers involved are considerably smaller. At its peak during the launch, the Sand Dollar employed 35

62 Khan and Malaika (2021).

63 For a discussion on CBDC projects and central bank governance, see Bechara and others (2021).

people at varying levels of time commitment. Currently, 15 people work full-time on the Sand Dollar. The ECCB is currently managing its DCash project with 12 people, all of whom in addition have other duties. This has been possible thanks to considerable technical expertise from outside. The Uruguay e-peso pilot similarly employed five full-time and five part-time employees. Again, these numbers refer only to central bank staff, and the full amount of personnel involved on the private sector side is likely considerably larger.

Table 5. Number of Central Bank Staff Engaged in CBDC Projects in Late 2021

Central Bank	Number of Staff
CBOB, Sand Dollar	15
BOC	50
PBOC, e-CNY	300
ECCB, DCash	12
Riksbank, e-krona	20
BCDU, ePeso	0 (10 during pilot)

Source: Central banks.
Note: This table does not include private sector personnel. Further, it does not distinguish between those working full time or part time on the CBDC project. The reason is the difficulty in comparing the time spent by part-time employees who, in some phases of the project, may work more than full time. Part-time employed, therefore, often means that they have other tasks besides CBDC. BCDU = Banco Central de Uruguay; BOC = Bank of Canada; CBDC = central bank digital currency; CBOB = Central Bank of Bahamas; ECCB = Eastern Caribbean Central Bank; PBOC = People's Bank of China.

C. Organization and Design of Pilots

This section focuses on the organization and execution of pilots. Canada and Sweden have not launched pilots, so are not discussed in this section. The three main aspects of a pilot are its general organization, how users are recruited and what results they yield, and how those results are incorporated.

General Organization of Pilots

The first main pilot design factor is how limited it will be. Pilots can be limited in time by having a clear termination date communicated in advance. But they can also be limited in scope in terms of how many users can participate or how much money will be issued.

The second main pilot design factor is its goals and the ability to revisit these as the pilot progresses. For instance, central banks may choose to develop and test new functions after the initial launch. Pilots are also sometimes directly used to further specific policy goals, for instance, by being extended to certain areas to support economic development or recovery after a natural disaster.

The four pilots studied in this paper are described below in the order they were first launched.

Planning for the Uruguay e-peso began in 2016 and the pilot ran from November 2017 to April 2018. Compared to the other pilots in this study, this effort was more contained in time and scope. It was clear from the outset that the pilot would end after six months, and all e-pesos owned by test users at the end of the pilot would be cashed in and destroyed. Total issuance was set at 20 million e-pesos, and no more than 10,000 end-users could take part by downloading an app on their cell phones. The pilot was also contained in terms of functional involvement of the private sector: from the outset, the different functions of the pilot were distributed among, as well as funded and developed by, a group of firms that were primarily interested

in testing aspects of their respective technologies. Thus, the e-peso was not an open platform, and commercial banks were not involved.

The CBOB launched the Sand Dollar pilot in December 2019 after more than three years of research into CBDC and planning the pilot. When considering suitable test areas, the bank opted to first roll out the pilot in the Exuma District in the South East Bahamas. Usage of mobile phones is high in Exuma, and testing the pilot there was ideal to ensure as many test users as possible. To create a baseline for measuring progress, the CBOB conducted a survey of the level of financial inclusion and willingness of the population to adopt digital payments.[64]

The pilot was rolled out to a second test area, the Abaco Islands, in February 2020. The Abacos infrastructure was severely damaged by a hurricane in September 2019, and the area was still recovering economically. The Sand Dollar pilot in this area, therefore, served a double purpose—to test off-line payments solutions as well as a means to support relief efforts and economic recovery.

In total, the Sand Dollar pilot included around 2,000 wallets, and around 35 persons at the central bank were involved in its launch. The COVID-19 pandemic made the execution of the Abacos tests more difficult but did not change the pilot plans.

The scale of the Chinese e-CNY pilot is unique. By October 8, 2021, over 123 million e-CNY wallets were held by individuals and around 9.2 million wallets were held by firms. To help organize such a large trial in different areas in China, the Digital Currency Institute of the People's Bank of China (PBCDCI) created several subsidiaries in Shenzhen, Suzhou, and Shanghai, and is considering creating more in other areas.

The PBCDCI has the main responsibility for planning and executing the e-CNY trials, but collaborates with local authorities, private intermediaries, and technology firms. So far, trials have been conducted in more than 10 cities and regions. The scale of the trials has allowed the PBCDCI to test both core technologies for raw payment processing but also for ancillary and add-on features such as identification, off-line payments, and programmability. Trials have increasingly been conducted in rural areas, following regional economic development goals.

The ECCB rolled out the DCash pilot to four countries in the ECCU in March 2021: Antigua and Barbuda, Grenada, Saint Kitts and Nevis and Saint Lucia. The pilot is scheduled to run for 12 months, after which all DCash are to be cashed in and destroyed. DCash has been issued on demand as the number of users grew, but the central bank announced that there would be a total limit on how much would be issued. The central bank initially stated that the pilot will be successful when DCash reaches 4,000 end-users and 35 merchants per country in the ECCU. But with experience, the ECCB has adjusted these goals to reflect differences between countries.

The ECCB altered the plan for its pilot because of two external events. First, the COVID-19 pandemic led to an unanticipated increase in demand for both DCash and online shopping. In response, the ECCB decided to expand the pilot to include online purchases using a web browser. Second, a volcano eruption in St. Vincent and the Grenadines also prompted the ECCB to accelerate the pilot in the affected area to help it recover by increasing access to payments.

In sum, the pilots of several jurisdictions were modified to address external events, thus highlighting the importance of a flexible approach.

Recruitment of Users for the Pilots

A CBDC pilot requires users willing to learn to use a new payments solution and trust it with their money. All central banks with pilots stressed the importance of information campaigns to recruit test users. In addition, financial incentives can be used. In the e-peso pilot, the BCDU's technology partner Roberto Giori paid for the information campaign to recruit test users and funded financial incentives: the first 1,000 users

[64] For a summary of the Exuma survey, see CBOB (2019).

received 1,000 e-pesos (approximately $23) for free, and 20 awards of 1,000 e-pesos were granted to the most active users for each month of the pilot.

The PBOC similarly launched a series of lotteries in collaboration with local authorities offering free e-CNY, which could be spent at merchants also joining the pilot. The local authorities provided the funding for these lotteries.

In the Bahamas and the ECCU, central banks have relied to a great extent on public information campaigns that stress the convenience and safety of paying with CBDC compared to physical means of payment. Recently, however, the ECCB added the incentive to get a percentage of expenses rebated in DCash at the end of the day for payments made in DCash at registered merchants.

Results of Pilots

Pilots can identify which areas need more testing, potentially through new pilots or extensions of pilots. The BCDU concluded that a potential second Uruguay CBDC pilot would need to be based on different principles compared to the first, including multiple vendors, and the participation of commercial banks.

Results of a pilot can also be used to improve the pilot or an officially launched CBDC. Staff at the CBOB stress that the most important gain from the Sand Dollar pilot was to better understand the motivations for potential users to adopt CBDC and for firms to join the CBDC network as intermediaries. This motivated the central bank to step up its communication and educational efforts on the local level. Further, the pilot showed the importance of increasing interoperability with the retail banking system to make it easier for users to convert bank deposits to Sand Dollars. As mentioned earlier, the pilot revealed that the planned off-line payment solution did not work as intended, and was the one major revision of the pilot in the officially launched Sand Dollar.

The ECCB's pilot is still in its early stages, so results are considered preliminary. However, demand has been sufficiently strong that ECCB now believes that ending the pilot after the planned 12 months might impair the payments system. Therefore, it is considering formally launching the CBDC and extending access to all countries in the ECCU rather than ending the pilot.

The PBOC reports that it is so far very pleased with the results of the e-CNY pilot. It has enabled testing a wide variety of different technological solutions for various features, including off-line capacity (see the sections above on design features and technology), payments methods using facial recognition, and tap-and-go. Surveys among test users, and the public, on the progress of the e-CNY, are also reported as being very favorable.

D. Stakeholders and Public CBDC Communication

Potential stakeholders in a CBDC project include the potential users, but also private intermediaries, incumbents in the payments and financial markets, as well as government agencies, representative political bodies, and governments. Some government agencies with a need to facilitate payments to invidividuals, such as tax agencies, social welfare agencies, or in some cases ministries of finance, might in particular have a stake in improving payments methods through CBDC.[65]

The introduction of CBDC requires approval that goes beyond the central bank. For instance, legal changes are often needed that are typically enacted by politicians in legislative bodies. Further, getting people to test, understand, and trust a CBDC pilot does not come automatically, and so without the buy-in of the public, there will never be a meaningful level of adoption. Therefore, communication with stakeholders is a key part of CBDC projects.

The CBOB, as mentioned above, stressed the importance of reaching out to potential user communities. The bank partnered with communication experts and marketing agencies, and the pilot and official launch

[65] For example, see IMF (2020b).

were accompanied by surveys and market research. The bank invited representatives of different industries to discussion forums to discuss what the Sand Dollar could mean for them. Specifically, the bank took time to promote to commercial banks what benefits the Sand Dollar could bring in terms of lower costs of handling cash and a potentially larger customer base.

The Bahamas government, and other government agencies, were supportive from the outset of the Sand Dollar project. A key potential benefit for government agencies is lower costs of handling public transfer payments to individuals. However, staff at the CBOB states that it would have been beneficial to have had more engagements with these stakeholders to ensure that digitalization efforts were more synchronized.

The PBOC and the ECCB also stress the importance of organized public information campaigns. In China, public information is extended by the PBOC but also by the private intermediaries which carry out face-to-face interactions with end-users. The ECCB partnered with market research agencies to better understand the public's needs and to receive real-time feedback as the pilot progressed. The e-peso was also accompanied by an educational campaign.

The BOC and the Riksbank, however, have not engaged in organized information campaigns. The Riksbank stresses that its communication strategy has been openness about its project rather than education or promotion. It has published regularly on its CBDC work, as well as participated in both national and international conferences, forums, and bilateral meetings with representatives of different stakeholders. The second e-krona report, published in 2018, was accompanied by a call for comments from stakeholders, which were published on the bank's website.

In 2019, the Riksbank sent in a petition to the Swedish Parliament to create a government inquiry into the future role of the state in the digital payments market – including assessing the pros and cons of a Swedish CBDC. Parliament approved the petition, and the inquiry was launched in 2021. This is an example of how central banks can directly solicit key stakeholders and elicit a policy response, in this case starting the process of potentially changing legislation to allow for the creation of a CBDC.

E. Major Challenges and Hindrances

Investigating, testing, and even launching CBDC comes with its challenges. The central banks studied in this paper raise several common themes.

Lack of precedents: Several central banks pointed out the difficulty of designing a project where there is little or no experience, nor established standards. However, prior research, even if conceptual, was of value to guide choices along the way. The central banks all emphasize the need to continue learning and experimenting.

Lack of resources: As demonstrated in this paper, CBDC projects are resource-intensive and become even more so as their scale increases. Thus, the PBOC raises resources as a constraint. Likewise, resource constraints are one of the key reasons why Uruguay has not yet launched a second e-peso pilot. Staff at the ECCB also stated that the financial cost of the DCash project has been one of the major obstacles to overcome.

Unwillingness to adopt digital payments among the population: Some jurisdictions mentioned that part of the population is suspicious about CBDC and digital payments in general. The CBOB has pointed out that part of the population still does not trust that their money is safe if converted to Sand Dollar and that they are concerned about privacy issues.

Legal issues: The need to make amendments or change laws and regulations is mentioned as one of the key obstacles by several jurisdictions.

Cyber security: The PBOC said that the risks from cyberattacks are substantial if the e-CNY becomes a crucial payments system. Creating an acceptable level of cyber security is one of the main challenges it sees.

Technological uncertainty: As technology is still developing, choosing the best technology is deemed a challenge. For instance, ECCB staff were uncertain whether DCash's DLT technology was sufficiently scalable to meet the demands of large-scale adoption. It is therefore open to considering another model.

F. Key Insights

Central bank staff gain experience and insights from running CBDC pilots and interacting with intermediaries and users. This section summarizes key insights raised by the staff at the six jurisdictions. Thus, these insights reflect the experiences of staff at individual central banks and are not necessarily immediately applicable in other contexts.

The importance of market research: Based on its experiences with the pilot and official launch of the Sand Dollar, the CBOB stresses the need to perform extensive market research to understand the needs of potential users.

Collaboration with participating private intermediaries: The CBOB underscores the need for the central bank to have strong collaboration and open communication with private firms that have face-to-face contact with the end-users. This point is also emphasized by the PBOC.

Technology neutrality: The PBOC is a strong proponent of neutrality. The e-CNY is designed as a hybrid system which, though its core is based on centralized technology, is fully compatible with DLT or other technologies that intermediaries choose to use. This reflects the PBOC's key recommendation that no technology is perfect and that being open to using different technologies is key. Similarly, the BCDU said that the simplest and most appropriate technology for the purposes of the CBDC should be favored, a principle that it followed when setting up the e-peso pilot.

Importance of cross-border payments: The PBOC stresses the importance of exploring cross-border payments with CBDC and adhering to the principles of "no disruption, compliance, and interoperability."

The anonymity/privacy trade-off: The PBOC emphasizes the need to manage the tension between anonymity and privacy, but that full anonymity for all transactions cannot be considered.

Allowing the public access to information on CBDC: The Riksbank highlights the importance for a central bank to be open about its work on CBDC. The first reason is that issuing CBDC is fundamentally about how to organize a society's payments system and therefore concerns everyone. The second reason is that understanding CBDC can take a long time and the process of communicating with the public (and decision-makers) should begin early in the process.

The importance of non-technical aspects: The BCDU stresses that a CBDC is not only a technical process but also a cultural one. The introduction of CBDC will have to be guided by careful knowledge about the cultural aspects of users and preferences for the characteristics of money.

8. Conclusions

This paper discussed six CBDC pioneer projects. It illustrated the importance of individual country context and policy goals for the design and implications of CBDC. Just as there is no universal case for CBDC, there is no universal design or recipe to implement CBDC.

CBDC is still in its infancy, and there are still open issues as well as commonly identified obstacles. Open issues referred to by several jurisdictions include the nature of sustainable business models that will ensure cost recovery and provide sufficient incentives for private sector participation. Other issues have to do with pushing the boundaries of innovation to allow for important features such as off-line capacity. The choice of technology is also frequently highlighted, including the use cases and limits of DLT. Key difficulties going forward include making choices in a very new and rapidly evolving field, as well as costs associated with the development process.

A new trend among some of the jurisdictions in this study, spearheaded by the PBOC, is a pragmatic view of technology. The choice between centralized and distributed technology does not need to be either-or. And central banks could adopt CBDCs that utilize different technologies for different ends.

While individual country contexts remain important, there are also areas of convergence. All central banks have explored the intermediated operational model. Countries are seeking a balance between preserving key aspects of the traditional monetary and financial system while at the same time updating the role of central banks in the digital era. Relatedly, all CBDCs currently in circulation have design characteristics that limit competition with bank deposits.

Examples of policy trade-offs were evident during the discussions, but policy synergies were also identified. The relationship between anonymity and illicit use of money, for example, presents a policy trade-off, but there are policy synergies between anonymity, risk reduction, and financial inclusion. Managing policy trade-offs and leveraging policy synergies could be a potential area of increased central bank attention in the future.

Pilot designs differ among the jurisdictions from strictly limited in time, scope, and goals to more open-ended. Pilots are also used as policy tools. The exact dividing line between an open-ended pilot and an officially launched CBDC is therefore not always clear-cut, especially since an officially launched CBDC can continue to be upgraded and developed after launch. To some degree, a pilot could therefore lead to a "soft launch."

CBDC exploration is still in an early stage, and not all country experiences can be easily ported abroad. There are still open questions, and CBDC remains an uncharted territory, raising challenges as well as opportunities. Increased international information-sharing of insights learned from individual CBDC projects and cooperation on policy and design issues will be important going forward. This paper represents an early contribution to this ongoing process.

References

Adrian, Tobias, and Tommaso Mancini-Griffoli. 2019. "The Rise of Digital Money." IMF Fintech Note 19/01, International Monetary Fund, Washington, DC.

Adrian, Tobias, and Tommaso Mancini-Griffoli. 2021. "Public and Private Money Can Coexist in the Digital Age." IMFBlog (blog), Februay 18. https://blogs.imf.org/2021/02/18/public-and-private-money-can-coexist-in-the-digital-age.

Agur, Itai, Anil Ari, and Giovanni Dell'Ariccia. 2022. "Designing Central Bank Digital Currencies." *Journal of Monetary Economics*, forthcoming.

Atlantic Council. 2021. "Central Bank Digital Currency Tracker." Accessed on December 17, 2021. https://www.atlanticcouncil.org/cbdctracker.

Armelius, Hanna, Paola Boel, Carl Andreas Claussen, and Marianne Nessén. 2018. "The E-krona and the Macroeconomy." *Sveriges Riksbank Economic Review* 2018 (3): 43-65.

Armelius, Hanna, Gabriela Guibourg, Stig Johansson, and Johan Schmalholz. 2020. "E-krona Design Models: Pros, Cons and Trade-offs." Sveriges Riksbank Economic Review 2020 (2): 80-96.

Armelius, Hanna, Carl Andreas Clausen, and Isaiah Hull. 2021. "On the Possibility of a Cash-Like CBDC." Staff Memo, Sveriges Riksbank, Stockholm.

Auer, Raphael, and Rainer Böhme. 2020. "The Technology of Retail Central Bank Digital Currency." BIS Quarterly Review March: 85-100.

Auer, Raphael, and Rainer Böhme. 2021. "Central Bank Digital Currency: The Quest for Minimally Invasive Technology." BIS Working Paper 948, Bank for International Settlements, Basel, Switzerland.

Auer, Raphael, Jon Frost, Leonardo Gambacorta, Cyril Monnet, Tara Rice, and Hyun Song Shin. 2021. "Central Bank Digital Currencies: Motives, Economic Implicatios, and the Research Frontier." BIS Working Paper 976, Bank for International Settlements, Basel, Switzerland.

Bechara, Marianne, Wouter Bossu, Yan Liu, and Arthur Rossi. 2021."The Impact of Fintech on Central Bank Governance." IMF Fintech Note 21/01, International Monetary Fund, Washington, DC.

Bindseil, Ulrich. 2019. "Central Bank Digital Currency: Financial System Implications and Control." International Journal of Political Economy 48 (4): 303-335.

Bindseil, Ulrich. 2020. "Tiered CBDC and the Financial System." ECB Working Paper 2351, European Central Bank, Frankfurt.

BIS (Bank for International Settlements). 2003. "The Role of Central Bank in Payment Systems." Committee on Payment and Settlement Systems, Basel, Switzerland.

BIS (Bank for International Settlements). 2017. "Distributed Ledger Technology in Payment, Clearing and Settlement: An Analytical Framework." CPMI Paper 157, Committee on Payments and Market Infrastructures, Basel, Switzerland.

BIS (Bank for International Settlements). 2018. "Central Bank Digital Currencies." Committee on Payments and Market Infrastructure, Basel, Switzerland.

BIS (Bank for International Settlements), BIS Innovation Hub, the International Monetary Fund, and the World Bank. 2021. "Central Bank Digital Currencies for Cross-border Payments." Joint report to the G20.

BISIH (Bank for International Settlements Innovation Hub), Bank of Thailand, Central Bank of the United Arab Emirates, Hong Kong Monetary Authority, and People's Bank of China. 2021. "Building a Multi CBDC Platform for International Payments." Joint report.

Boar, Condruta, and Andreas Wehrli. 2021. "Ready, Steady, Go? Results of the Third BIS Survey on Central Bank Digital Currency." BIS Paper 114, Bank for International Settlements, Basel, Switzerland.

BOC (Bank of Canada), Bank of England, and Monetary Authority of Singapore. 2018. "Cross-Border Interbank Payments and Settlements. Emerging Opportunities for Digital Transformation." Joint report.

BOC (Bank of Canada). 2020. "Contingency Planning for CBDC." Background Note. Bank of Canada, Ottawa.

BOE (Bank of England). 2020. "Central Bank Digital Currency: Opportunities, Challenges, and Design." Future of Money Discussion Paper, Bank of England, London.

Bossu, Wouter, Masaru Itatani, Catalina Margulis, Arthur Rossi, Hans Weenink, and Akihiro Yoshinaga. 2020. "Legal Aspects of Central Bank Digital Currency : Central Bank and Monetary Law Considerations." IMF Working Paper 20/254, International Monetary Fund, Washington, DC.

BOT (Bank of Thailand) and HKMA (Hong Kong Monetary Authority). 2020. "Inthanon-LionRock. Leveraging Distributed Ledger Technology to Increase Efficiency in Cross-Border Payments." Joint report.

Bruton, Garry, Mike Peng, David Ahlstrom, Stan Ciprian, and Kehan Xu. 2014. "State-Owned Enterprises Around the World as Hybrid Organization." *The Academy of Management Perspectives* 29 (1) : 92-114.

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

Central Banks and BIS. 2020. "Central Bank Digital Currencies: Foundational Principles and Core Features". Joint report by Bank of Canada, European Central Bank, Bank of Japan, Sveriges Riksbank, Swiss National Bank, Bank of England, Board of Governors Federal Reserve System and Bank for International Settlements.

Central Bank Act of The Bahamas. 2020. Nassau.

Chohan, Usman. 2021. "The Double Spending Problem and Cryptocurrencies." *SSRN Electronic Journal* Doi: 10.2139/ssrn.3090174.

Dev, Mahendra. 2006. "Financial Inclusion: Issues and Challenges." *Economic and Political Weekly* 41 (41): 4310-4314.

ECB (European Central Bank). 2020. "Report on a Digital Euro." European Central Bank, Frankfurt.

FATF (Financial Action Task Force). 2015. "Money Laundering Through the Physical Transportation of Cash." FATF, Paris, and MENAFATF (Middle East & North Africa Financial Action Task Force), Manama, Bahrain.

FATF (Financial Action Task Force). 2020. "FATF Removes The Bahamas from the List of Jurisdictions under Increased Monitoring." Press release, December 18, 2020. https://www.fatf-gafi.org/publications/high-risk-and-other-monitored-jurisdictions/documents/bahamas-delisting-2020.html.

FSB (Financial Stability Board). 2020. "Enhancing Cross-Border Payments. Stage 3." Financial Stability Board, Basel.

G7 (Group of Seven). 2021. "Public Policy Principles for Retail Central Bank Digital Currencies." October.

IMF (International Monetary Fund). 2019. "The Bahamas. Financial Sector Assessment Program. Technical Note on Financial Inclusion, Retail Payments, and SME Finance." IMF Country Report 19/201, International Monetary Fund, Washington, DC.

IMF (International Monetary Fund). 2020a. "Digital Money Across Borders: Macro-financial Implications." IMF Policy Paper, International Monetary Fund, Washington, DC.

IMF (International Monetary Fund). 2020b. "Digital Solutions for Direct Cash Transfers in Emergencies." Special Series on Fiscal Policies to Respond to COVID-19, International Monetary Fund, Washington, DC.

Herlihy, Maurice. 2018. "Atomic Cross-Chain Swaps." In Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing. Association for Computing Machinery, New York, 245-254.

HKMA (Hong Kong Monetary Authority). 2021. "e-HKD: A technical perspective." Request for Comments, Hong Kong Monetary Authority.

Juks, Reimo. 2018. "When a central bank digital currency meets private money: effects of an e-krona on banks." *Sveriges Riksbank Economic Review* 2018 (3): 79-99.

Khan, Ashraf, and Majid Malaika. 2021. "Central Bank Risk Management, Fintech and Cybersecurity." IMF Working Paper 21/105, International Monetary Fund, Washington, DC.

Kiff, John. 2021. "Jurisdictions Where Retail CBDC Is Being Explored." *Kiffmeister Chronicles*. http://kiff-meister.blogspot.com/2019/12/countries-where-retail-cbdc-is-being.html.

Kiff, John, Jihad Alwazir, Sonja Davidovic, Aquiles Farias, Ashraf Khan, Tanai Khiaonarong, Majid Malaika. 2020. "A Survey of Research on Retail Central Bank Digital Currency." IMF Working Paper 20/104, International Monetary Fund, Washington, DC.

Kumhof, Michael, and Clare Noone. 2018. "Central Bank Digital Currencies – Design Principles and Balance Sheet Implications." Staff Working Paper 725, Bank of England, London.
Natarajan, Harish, Solvej Krause, and Helen Gradstein. 2017. "Distributed Ledger Technology and Blockchain." FinTech Note No. 1. World Bank, Washington, DC.
NIST (National Institute of Standards and Technology). 2021. "System Owner." Computer Security Resource Center, Glossary. https://csrc.nist.gov/glossary/term/system_owner.
Miedema, John, Cyrus Minwalla, Martine Warren, Dinesh Sha. 2020. "Designing a CBDC for Universal Access." Staff Analytical Note 2020-10, Bank of Canada, Ottawa.
Mu, Changchun. 2021. "Navigating the Uncharted Waters of Retail CBDC." Presentation at the BIS Innovation Summit (virtual conference), March 25. https://www.bis.org/events/bis_innovation_summit_2021/agenda.htm.
Ozili, Peterson. 2020. "Financial Inclusion Research Around the World: A Review." *Forum for Social Economics* 50 (4): 457-479.
PBC (People's Bank of China). 2020. Revised Law for People's Bank of China. Draft for consultation.
Shabsigh, Ghiath, Tanai Khiaonarang, and Harry Leinonen. 2020. "Distributed Ledger Technology Experiments in Payments and Settlements." IMF Fintech Note 20/01, International Monetary Fund, Washington, DC.
Soderberg, Gabriel. 2019. "The E-krona – Now and for the Future." Economic Commentary No. 8, Sveriges Riksbank, Stockholm, Sweden.
SOU (Statens Offentliga Utredningar). 2021. "Committee Terms of Reference. Role of the State in the Payment Market." Statens Offentliga Utredningar, Stockholm.
Sveriges Riksbank. 2017. "The Riksbank's e-krona Project. Report 1." Sveriges Riksbank, Stockholm, Sweden.
Sveriges Riksbank. 2018. "The Riksbank's e-krona Project, Report 2." Sveriges Riksbank, Stockholm, Sweden.
Sveriges Riksbank. 2021a. "The Position of Cash as Legal Tender Needs Strenghtening." In *Payments Report 2021*. Sveriges Riksbank, Stockholm, Sweden.
Sveriges Riksbank. 2021b."E-krona Pilot. Phase 1." Sveriges Riksbank, Stockholm.
Utredningen om civilt försvar. 2021. "Struktur för ökad motståndskraft." Statens Offentliga Utredningar 2021:25. Regeringskansliet, Stockholm.
World Bank. 2021. *The State of Economic Inclusion Report 2021: The Potential to Scale*. Washington, DC: World Bank.
World Bank and PBOC (People's Bank of China). 2018. "Toward Universal Financial Inclusion in China: Models, Challenges, and Global Lessons." Joint Report. World Bank, Washington, DC.

FINTECH

NOTES

Behind the Scenes of Central Bank Digital Currency

Emerging Trends, Insights, and Policy Lessons

Gabriel Soderberg

In collaboration with Marianne Bechara, Wouter Bossu, Natasha Che, Sonja Davidovic, John Kiff, Inutu Lukonga, Tommaso Mancini-Griffoli, Tao Sun, and Akihiro Yoshinaga

NOTE/2022/004

Adoption of Central Bank Digital Currency: Lessons from E-Money Schemes in Asia

To Be Published Shortly.

Designing a Central Bank Digital Currency with Support for Cash-Like Privacy

Designing a Central Bank Digital Currency with Support for Cash-like Privacy

Jonas Gross[a], Johannes Sedlmeir[a,b], Matthias Babel[a], Alexander Bechtel[c], Benjamin Schellinger[a,d]

[a]University of Bayreuth; [b]Project Group Business & Information Systems Engineering of the Fraunhofer FIT; [c]University of St. Gallen; [d]FIM Research Center, University of Bayreuth

ARTICLE VERSION

January 14, 2022

ABSTRACT

Most central banks in advanced economies consider issuing central bank digital currencies (CBDCs) to address the declining use of cash as a means of payment and to position themselves against increased competition from Big Tech companies, cryptocurrencies, and stablecoins. One crucial design dimension of a CBDC is the degree of transaction privacy. Existing solutions are either prone to security concerns or do not provide full (cash-like) privacy. Moreover, it is often argued that a fully private payment system and, in particular, anonymous transactions cannot comply with anti-money laundering (AML) and countering the financing of terrorism (CFT) regulation. In this paper, we follow a design science research approach (DSR) to develop and evaluate a holistic software-based CBDC system that supports fully private transactions and addresses regulatory constraints. To this end, we employ zero-knowledge proofs (ZKPs) to enforce limits on fully private payments. Thereby, we are able to address regulatory constraints without disclosing any transaction details to third parties. We evaluate our artifact through interviews with leading economic, legal, and technical experts and find that a regulatorily compliant CBDC system based on ZKPs that supports full (cash-like) privacy is feasible.

KEYWORDS

Anonymity, CBDC, Compliance, Design Science, Digital Identity, Digital Wallet, Electronic Cash, Payment System, Privacy by Design, Regulation, Self-Sovereign Identity, Zero-Knowledge Proof

Corresponding author: Jonas Gross, University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany, jonas.gross@uni-bayreuth.de.

Excutive Summary

We propose a two-tiered (retail) central bank digital currency (CBDC) system consisting of transparent and private CBDC accounts. The main focus of this paper lies on the private accounts provided by the *privacy pool*. Our general architecture is illustrated in Figure 1. The central bank issues CBDC into transparent accounts. These accounts are maintained either by the central bank or by payment service providers (PSPs), such as banks, on behalf of the central bank. CBDC account holders can deposit CBDC into the privacy pool either via transparent CBDC accounts or deposit cash via special automated teller machines (ATMs). End-users can conduct three different types of payments differing in the degree of privacy, namely fully private, semi-private, and fully transparent transfers.

- Fully private (cash-like) transfers take place inside the privacy pool. In fully private transfers, the identities of both involved parties, as well as the transaction amount, remain hidden to third parties.
- Semi-private transfers take place between the privacy pool and the transparent CBDC accounts. They reveal the amount of the transfer to the PSP and the central bank, and only the identity of the holder of the transparent CBDC account to the involved PSPs and potentially, depending on the design, also to the central bank.
- Fully transparent transfers take place between transparent CBDC accounts. In fully transparent transfers, the sender's and receiver's PSPs and, depending on the design, potentially also the central bank, know the identities of the involved parties and the transaction amount, similar to current electronic payments via commercial banks.

For semi-private and fully private transfers, the privacy guarantees are provided *by design*, i.e., no third parties and, in particular, neither PSPs nor the central bank nor regulatory authorities, will learn about the transaction amount and the involved parties even if they collude or, in the most extreme case, behave maliciously.

Addressing regulatory constraints related to anti-money laundering (AML) and combating the financing of terrorism (CFT) regulation while allowing for fully private payments, we impose balance, transfer, and turnover limits. We enforce these limits without sharing transfer details beyond the information that we discussed previously with any third parties by using general-purpose zero-knowledge proofs (ZKPs)

in the form of zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs).

We develop an *unspent account state (UAS)* system to obtain the privacy and, in particular, unlinkability guarantees of Zcash with a similar technical construction based on commitments, nullifiers, and ZKPs. More precisely, we do not use a token-based model (corresponding to *coins* or *digital banknotes*), as it is unclear how turnover and balance limits could be implemented in such an approach. Instead, we use commitments and nullifiers to obfuscate transfers in an account-based model (including an *identity* and a *balance*) and make them unlinkable. This approach facilitates the enforcement of turnover and balance limits for fully private payment, as it funnels all associated transfers conducted by one particular user into one unique privacy pool account that they control. In our design, all end-users privately register and maintain their own private CBDC account and only send cryptographic proofs of the correct local accounting and compliance with the imposed limits to the ledger that is maintained by the central bank. A transfer proposal is only accepted by the central bank if the ZKP rightfully proves the compliance with the limits imposed. To ensure that every user only has access to one account, and hence make turnover and balance limits effective, our construction relies on the availability of a *unique* digital identity for end-users. For this purpose, we suggest using government-issued, certificate-based, digital IDs, as currently explored in various countries. Alternatively, PSPs such as banks could issue these identities, which would, however, require coordination among these intermediaries to prevent multiple registrations.

From the perspective of a user, an anonymous transfer corresponds to invalidating their previous account state and creating a new account state, where the difference between the previous balance and the new balance equals the amount of the transfer. Due to the fully-fledged user-sided accounting, our CBDC system would also allow integrating remuneration in private CBDCs accounts that directly affects the difference between the previous and the new balance. Third parties cannot trace back a transaction history to one specific user, and transfers by the same user are unlinkable. The only entities in the system that have access to their account details, including their ID information, balance, and turnover, are the users themselves. Users keep their CBDC in a digital wallet, e.g., in a dedicated app on a mobile device. The digital wallet stores the users' digital ID, cryptographic keys, and account details for their private CBDC account as well as access credentials for their transparent CBDC account. The wallet also supports users in maintaining their private and transparent accounts. For example, if one of the limits in the privacy pool prevents a user from conducting an

anonymous payment, the wallet may automatically suggest conducting the payment on the transparent side.

The implementation of our system can already be regarded as feasible. From a user's perspective, the processing time for a transfer would be comparable to private transfers in Zcash. Hence, payments would be processed within a few seconds or even less. Given the rapid improvements on the performance of generating ZKPs observed over the last few years, it seems likely that processing times will see further significant improvements in the near future. Additionally, the operational burden for the central bank is relatively low because the succinct proofs in zk-SNARKs are small and offer fast verification in few milliseconds. The main challenge of our approach is the prevention of a black market for private CBDC accounts: Getting control over a large number of private CBDC accounts by buying them from users that are not interested in using the privacy pool or by blackmailing would allow illicit actors to circumvent the limits and use the privacy pool for a *money mule* business. However, this threat can be effectively mitigated by binding users' digital identities to secure hardware, e.g., their mobile phones, and by requiring a ZKP of ownership of the private key stored in this secure hardware and associated with a non-expired, non-revoked unique digital identity in each transaction. This approach provides a high degree of assurance that users own the digital ID associated with the privacy pool account for every private transfer without the need to reveal correlatable identity-related information to third parties. Users that seek to pass on their accounts would need to pass on their unlocked mobile phone, on which their digital ID is stored, including all associated rights and permissions. Our discussions with stakeholders indicate that this risk is comparatively low and thus acceptable.

Figure 1.: High level architecture of the proposed CBDC design.

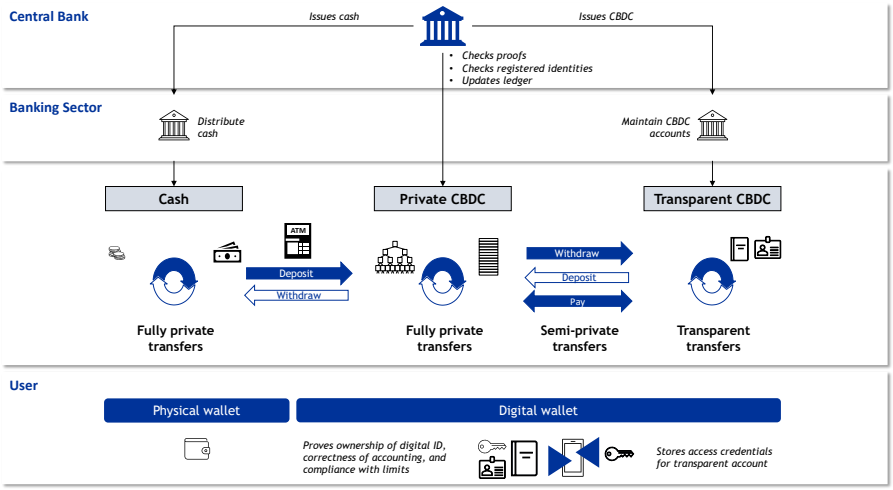

1. Introduction

The monetary system is changing. In many advanced economies, the use of cash as a means of payment has declined steadily over the last decade and in an accelerated way during the COVID-19 pandemic (European Central Bank, 2020b). Moreover, public money faces increasing competition from novel, private sector-issued forms of money, such as cryptocurrencies and stablecoins, and from Big Tech payment systems (European Central Bank, 2020a). Consequently, central banks take actions to preserve their monetary sovereignty. Today, 86 % of central banks around the world consider issuing their own digital currencies, known as central bank digital currencies (CBDCs) (Boar & Wehrli, 2021). While the Bahamas has already launched a CBDC and some other countries have introduced CBDC pilots (e.g., China), most jurisdictions are still debating and analyzing design options. In this context, the appropriate degree of transaction privacy receives great attention from central bankers, policy-makers, and academics. For instance, in its announcement to start a project on the digital euro, the European Central Bank (ECB) stressed that the two-year investigation phase aims to identify "the design options to ensure privacy and avoid risks for euro area citizens, intermediaries and the overall economy" (European Central Bank, 2021c).

Keeping transaction data private is crucial, among other reasons, to avoid identity theft, threats to personal security, data exploitation, and harassment based on potentially embarrassing but legal purchases (e.g., Chaum, Grothoff, & Moser, 2021; Choi, Henry, Lehar, Reardon, & Safavi-Naini, 2021; Kahn, 2018; Kahn, McAndrews, & Roberds, 2005). Privacy is also considered essential from an economic perspective, as it can help to avoid price discrimination (Acquisti, Taylor, & Wagman, 2016; Odlyzko, 2004), making privacy a public good (Garratt & Van Oordt, 2021). Moreover, privacy constitutes a fundamental civil right enshrined in Article 12 of the United Nations Declaration of Human Rights (UDHR) (United Nations, 1948), in Article 8 of the Convention for the Protection of Human Rights and Fundamental Freedoms (European Court of Human Rights, 1950) as well as in Article 7 and 8 of the Charter of Fundamental Rights of the European Union (European Convention, 2000). In the context of CBDC, a consultation of European citizens demonstrated that they see privacy as the *most* important requirement for a CBDC (European Central Bank, 2021b).

A CBDC that stores transaction details in a centralized database operated by the central bank (or payment service providers (PSPs) on behalf of the central bank) bears the risk of losing trust and causing security incidents such as data breaches, e.g., due to human misbehavior or cyber attacks. Incidents like the hack of New Zealand's

central bank in 2021 demonstrates that cyber risks are indeed a threat that should be taken seriously (The Guardian, 2021). Furthermore, if sensitive data are stored centrally, end-users have to *trust* the operator that the privacy promises will not be compromised in the future. However, in such a setup, operators could potentially change their mind or secretly analyze (historic) transaction data and share it with further parties, thereby potentially undermining privacy and trust. Trust in a payment system that inevitably processes sensitive data can be increased by following a *privacy-by-design* approach in which customers do not need to trust the operator for privacy protection and where large-scale data breaches are naturally excluded. In this case, private data would only be stored with the end-user involved in a transaction and not aggregated in a centralized system, thereby providing *trustless privacy*.

Today, cash is the only regulatorily compliant form of money that provides full privacy by design. If payments are conducted digitally, e.g., through mobile payments, bank transfers, or credit cards, the transaction data is stored with the involved PSP. Contrary to public perception, cryptocurrencies such as Bitcoin and Ether do not ensure a high degree of privacy either, as transaction details are stored on a public ledger, and the pseudonymous addresses that send and receive cryptocurrencies can often be traced back to the users that control them (Biryukov, Khovratovich, & Pustogarov, 2014) through analyzing metadata (such as IP addresses) and information from exchanges that need to conduct know-your-customer (KYC) measures (Silfversten et al., 2020). Against this background, privacy-oriented cryptocurrencies such as Zcash and Monero have been developed. They use cryptographic techniques such as zero-knowledge proofs (ZKPs) to enable fully private payments (Fauzi, Meiklejohn, Mercer, & Orlandi, 2019). However, these cryptocurrencies do not conform with prevailing regulations, as unlimited anonymous payments open the door for illicit activities, such as money laundering and terrorist financing (Silfversten et al., 2020).

To secure access to a fully private, regulatorily compliant form of money in an increasingly digital environment, a CBDC should provide a high degree of transaction privacy and offer (at least) the same privacy-preserving features as cash. Multiple central banks have already indicated their willingness to consider privacy-enhancing features for their CBDCs (Bank of Canada, 2020; European Central Bank, 2021a; Lane, 2020), and first CBDC solutions have been proposed by both central bankers and academic researchers that provide some degree of transaction privacy. Naturally, these suggestions also consider regulatory constraints. However, software-based CBDC designs proposed by central banks, e.g., the ECB's anonymity voucher proposal (European Central Bank, 2019), or by academic researchers (Chaum et al., 2021; Dold, 2019;

Tinn & Dubach, 2021), do not support fully private transactions. Besides software-based designs, CBDC solutions can use hardware elements, e.g., built in computers, mobile phones, or smart cards, as gateways to access the CBDC infrastructure (European Central Bank, 2020a) and could, therefore, technically replicate the trustless privacy guarantees of cash. As an example, a hardware-based smart card system is currently being tested by the Central Bank of the Bahamas (Mastercard, 2021). However, such hardware-based solutions still exhibit considerable security challenges (Chaum et al., 2021; European Central Bank, 2021c), and mitigating the risks of sophisticated attacks would likely compromise privacy guarantees (Chaum et al., 2021).[1]

Fortunately, the maturity of privacy-enhancing cryptographic techniques, and particularly of ZKPs, has grown considerably in recent years, offering new opportunities for enhanced privacy. ZKPs have already seen considerable adoption in the context of privacy-oriented cryptocurrencies, where they are used to ensure the integrity of payment systems, e.g., to prevent double-spending, while maintaining a high degree of privacy for the user. However, ZKPs can be more broadly employed, with particular emphasis on enforcing further monetary or regulatory rules in a privacy-oriented payment system. Literature on cryptography already acknowledges the suitability of ZKPs for reconciling privacy and integrity or compliance requirements for electronic payments, e.g., through imposing turnover or per-transaction limits (e.g. Garman, Green, & Miers, 2016). Bontekoe (2020) specifically proposed an extension of Zcash in which third parties escrow users' digital identities and ZKPs allow the enforcement of turnover limits.

Still, regulators and central bankers have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Armelius, Claussen, & Hull, 2021; Auer & Boehme, 2021). This statement clearly indicates a lack of communication between the different research streams. Moreover, neither CBDC nor cryptographic literature has so far provided a rigorous, holistic evaluation of a payment system design that addresses regulatory requirements while supporting fully private payments and that evaluates the design with key stakeholders, such as central bankers and regulators. To address this research gap, we follow a rigorous design science research (DSR) approach to design and evaluate a holistic, software-based CBDC system that is based on ZKPs and supports fully private payments. We first consolidate proposals from the cryptographic and CBDC literature to develop an account-based CBDC payment sys-

[1] We recommend to refer to the comprehensive discussion in Chaum et al. (2021) for a more detailed overview of the challenges associated with hardware-based solutions, particularly if deployed on a larger scale, and to Grothoff and Dold (2021) for additional arguments why a software-based CBDC design may be beneficial.

tem that is fully private by design while addressing regulatory requirements by using per-transaction, turnover, and balance limits. We also instantiate our design through an implementation of the core transaction types using ZKPs. We then evaluate and refine our IT artifact in four evaluation cycles consisting of a total of 22 interviews with 44 experts in the area of regulation, cryptography, central banking, identity, and payments. We find using ZKPs for CBDCs can replicate cash-like privacy in the digital realm and ensure adherence to regulatory constraints. Against this background, ZKPs enable strict privacy protection by design, storing personal transaction data only on the end-users' devices (*trustless privacy*).

The theoretical contribution of our paper is twofold: First, our innovative software-based CBDC payment system combines elements from different strands of the literature, including cryptography, privacy-by-design concepts, and CBDCs. Second, the rigorous evaluation of our CBDC system by key stakeholders in the cadre of DSR allows us to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs and also to discuss risks and potential mitigation measures.

Our paper is structured as follows: In Section 2, we introduce essential background knowledge on CBDCs, different notions of transaction privacy, regulatory aspects of CBDCs, and ZKPs. We then present our DSR approach in Section 3. Subsequently, we discuss related work and describe the design of our IT artifact in Section 4, followed by the presentation of our evaluation cycles in Section 5. Section 6 summarizes our main findings, describes limitations, and gives an outlook on potential future applications of our design and related research opportunities.

2. Theoretical Foundations

2.1. Central Bank Digital Currencies

In general, there are two forms of CBDCs, wholesale and retail CBDCs (Bech & Garratt, 2017). A wholesale CBDC is a digital form of central bank money accessible for financial institutions to optimize the settlement of wholesale payments and tokenized financial assets. A retail CBDC, in contrast, constitutes a novel form of central bank money available to the general public. In this paper, we solely refer to a retail CBDC, as we focus on end-user payments. A retail CBDC unites features of today's predominant forms of money: cash and bank deposits (Bech & Garratt, 2017). While cash is issued by central banks in physical form, bank deposits are issued by commercial banks in digital form. As central bank money, CBDCs bear no counterparty risk

because the central bank issuing the CBDC cannot – by definition and in contrast to commercial banks – become bankrupt.[2] CBDCs would hence provide a safer and practically riskless form of money for end-users.

CBDCs can be designed and implemented in different ways (Allen et al., 2020; Auer & Böhme, 2020; European Central Bank, 2020b; Kiff et al., 2020). Auer and Böhme (2020) identify *architecture*, *access*, and *technology* as the three main design considerations for a CBDC.[3] The architecture specifies the role of the central bank and other market participants in the CBDC ecosystem. The account management, onboarding processes, and distribution of a CBDC might be conducted directly by the central bank (direct model) or by private sector PSPs (intermediated model). The access model defines how CBDC transaction data is stored and how access is managed. In an account-based model, the CBDC is stored in accounts, and hence the ownership of a CBDC is tied to an identity. In a token-based model, the central bank issues digital bearer instruments and ties the CBDC ownership to the (proof of) ownership of the CBDC units itself, similar to cash today. Regarding technology, a CBDC can be issued either via a centralized or a distributed ledger. If a centralized ledger is used, the central bank manages and controls the CBDC system. In the case of a distributed ledger, data processing, storage, and governance can be distributed across additional private or public sector institutions.

2.2. Privacy and Regulatory Compliance of Payments

In this paper, we distinguish between private, anonymous, and fully (cash-like) private transactions. In a *private* transaction, the transaction amount remains unknown, but the sender and receiver, i.e., the transaction parties, might be known to third parties (e.g., PSPs, the central bank, or regulatory authorities). In an *anonymous* transaction, the identities of the sender and receiver remain hidden, but the transaction amount might be known. *Fully private* transactions are private and anonymous; neither the transaction amount nor the sender or receiver are revealed to third parties. Therefore, our definition of full privacy is similar to the concept of secrecy as the concealment of

[2]Today, deposit insurance schemes are established to address the risk of commercial bank money and avoid that, in the case of bankruptcy of a commercial bank, costumers face substantial financial losses. However, deposit insurance schemes are not available in all countries equally, and commercial bank money is only secured until a specific threshold, e.g., in the euro area up to 100,000 Euro per client per financial institution.

[3]The fourth dimension refers to retail and wholesale interlinkages. Such interlinkages are especially relevant for cross-border CBDC payments. As we abstract from cross-border use in this paper, we do not consider this dimension.

information (Bok, 1989; Tefft, 1980). Full privacy or secrecy describes the attempt of consumers to avoid sharing information in order to prevent third parties from creating a digital representation of the real self (Dinev, Xu, Smith, & Hart, 2013; Zwick & Dholakia, 2004).

Today, fully private and regulatorily compliant transactions are only possible with physical cash and, to a certain extent, with e-money. All other payment methods are either not fully private (e.g., credit cards, bank transfers, Bitcoin), or they are not regulatorily compliant (e.g., privacy-oriented cryptocurrencies such as Monero and Zcash). In order to restrict the large-scale financing of illicit activities, regulators usually enforce per-transaction, turnover, and/or balance limits for anonymous payments. For instance, there are *per-transaction limits* for fully private cash payments in many euro area countries such as Greece (500 EUR), France and Portugal (1,000 EUR), Italy (2,000 EUR), Spain (2,500 EUR), Belgium (3,000 EUR), and Slovakia (15,000 EUR) (Pocher & Veneris, 2021). For anonymous e-money transactions, the 5th Anti-Money Laundering Directive specifies a monthly *turnover* and *balance limit* of 150 EUR (European Union, 2018). As, to date, regulatory frameworks do not capture CBDCs, there are no such regulatory limits for CBDCs yet. However, it seems reasonable to expect that similar limits would need to be introduced for anonymous CBDC payments, similar to today's restrictions for anonymous cash and anonymous e-money transactions.

2.3. Zero-Knowledge Proofs

The notion of ZKPs was first introduced in the 1980s, describing "proofs that convey no additional knowledge other than the correctness of the proposition in question" (Goldwasser, Micali, & Rackoff, 1989, p. 186). ZKPs refer to cryptographic protocols in which a *prover* can convince a *verifier* about a mathematical statement, for example, that the prover knows a piece of data that has specific properties. This statement may refer to the knowledge of a pre-image of a publicly known value under a hash function or about properties of the result of a publicly known algorithm that was executed on public or private data. In this setting, with a ZKP, the prover can convince the verifier without disclosing *any* information beyond the statement under consideration (Ben-Sasson, Bentov, Horesh, & Riabzev, 2018; Ben-Sasson, Chiesa, Genkin, Tromer, & Virza, 2013). If the statement refers to the output of an algorithm that was applied to data that is public or that was publicly committed to, a ZKP can enforce *computational integrity* without the need for the verifier to replicate the computation. Besides

providing *confidentiality* for data and intermediate steps in a computation, an appealing property of many ZKPs is that they are *succinct*, i.e., the size of proofs and the computational complexity of proof verification is significantly smaller than applying the algorithm. In the case of the zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs) that we use in our CBDC system, both proof size and verification complexity are even *independent* from the complexity of the computation that is to be verified (Ben-Sasson et al., 2013). However, general-purpose ZKPs that can cover a large class of statements come at a high computational overhead for the prover.

In the 25 years after their discovery, researchers have leveraged special types of ZKPs in some contexts, such as enforcing *correct behaviour* in multiparty computations or *selective disclosure* in digital identity management schemes with *anonymous credentials*. The latter describes digital certificates that a trusted organization signed digitally and that their owner can use to prove claims about parts of the content of these certificates without revealing all of the contained information. In particular, when verifiably presenting attributes attested in the anonymous credential, strongly correlating contents such as the value of the digital signature itself do not need to be revealed (Ben-Sasson et al., 2013). Lately, these anonymous credentials have seen first adoption in so-called decentralized or self-sovereign identity projects as explored by the public and private sector in Canada and Germany, among others (Kubach & Sellung, 2021). However, practical applications remained rare, as for general-purpose ZKPs beyond these very specific cases, the computational complexity for the prover was prohibitive. Also, seemingly, there was not a considerable need for deploying ZKPs because information systems (IS) were generally designed with a service provider that was trusted with respect to both integrity and confidentiality. However, this paradigm started to shift with the advent of Bitcoin and the decentralization as facilitated by blockchain technology. Building on blockchain technology, a new type of IS, decentralized applications, emerged that do not involve a third party that is trusted with respect to integrity by performing computations redundantly (Rossi, Mueller-Bloch, Thatcher, & Beck, 2019). However, the replicated execution of operations on blockchains immediately leads to considerable challenges from a scalability and confidentiality perspective (Ben-Sasson et al., 2018; Kannengießer, Lins, Dehling, & Sunyaev, 2020).

In this context, general-purpose ZKPs started to find applications in privacy-oriented cryptocurrencies such as Zcash or applications on Ethereum such as Tornado-Cash, building on prior academic work (Ben-Sasson et al., 2014) to provide a Bitcoin-like payment system with fully private transactions. In the last few years, addi-

tionally, the succinctness of proofs has been leveraged by various projects on the Ethereum blockchain and novel cryptocurrencies. In zk-rollups, an untrusted third party batches many operations and proves the correctness of the resulting state transition with a ZKP. Through their ability to solve privacy challenges in cryptocurrency and blockchain projects (Partala, Nguyen, & Pirttikangas, 2020), ZKPs have received increased attention in academia and business, and have hence considerably matured in terms of performance and applicability. Consequently, IS building on blockchains and general-purpose ZKPs have already seen first adoption in industry consortia that leverage blockchain technology, e.g., in the context of medical supply chains (Mattke, Hund, Maier, & Weitzel, 2019).

3. Method

This paper follows a rigorous DSR approach (Hevner, March, Park, & Ram, 2004; March & Smith, 1995; Peffers, Tuunanen, Rothenberger, & Chatterjee, 2007) to design, develop, and evaluate a CBDC system that provides full privacy while addressing regulary requirements related to anti-money laundering (AML) and combating the financing of terrorism (CFT). We structure our paper as proposed by Gregor and Hevner (2013). To ensure methodological rigor, we use the widely accepted DSR methodology proposed by Peffers et al. (2007). Thus, we apply the following six steps procedure to derive our IT artifact: (1) Problem identification, (2) objectives definition, (3) design and development, (4) demonstration, (5) evaluation, and (6) communication (see Figure 2).

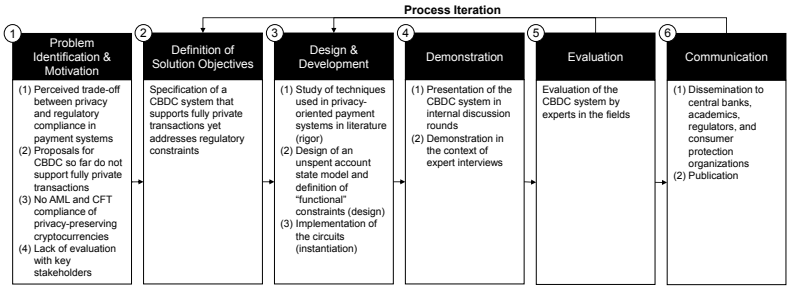

Figure 2.: DSR approach to design our account-based CBDC system according to Peffers et al. (2007).

DSR was established to enable IS practitioners to find solutions to previously unsolved problems through a continuous build-and-evaluate process (Hevner et al., 2004; March & Smith, 1995). For an IT artifact to make a valuable contribution to IS research, it must address both a relevant business need (Hevner et al., 2004) and a general problem (Iivari, 2015). First, we identified the underlying problems and derived design requirements for our CBDC system. We screened the most relevant primary literature on CBDCs, namely ECB (2020a) (European Central Bank, 2020a), BoC (2020) (Bank of Canada, 2020), Fed (2021) (Cheng, Lawson, & Won, 2021), BoE (2020) (Bank of England, 2020), and BIS et al. (2020) (Bank for International Settlements et al., 2020), and identified both users' and central banks' requirements for a CBDC (see Table 1).

We derived the following key requirements for end-users: privacy protection, high security and transaction speed, fast settlement, low costs, high usability, and availability. For central banks, the following requirements are important: AML and CFT compliance, market neutrality, resilience, cooperation with market participants, universal access, cost efficiency, and interoperability. In Section 1, we argued that privacy features should be at the core of a CBDC system, but also that regulatory constraints need to be addressed. Thus, in our DSR approach, we focus on these two core requirements. As CBDC implementations are currently still at an early stage, it is not (yet)

Table 1.: Literature review on CBDC requirements based on central bank statements. We included requirements that have been referred to in at least four of the five studies.

	ECB (2020a)	BoC (2020)	Fed (2021)	BoE (2020)	BIS et al. (2020)
Liability of the central bank	●	●	●	●	●
Market neutrality	●	●	●	○	●
Security	●	●	●	●	●
Convenience and ease of use	●	●	●	●	●
Transaction speed and fast settlement	●	●	●	●	●
Privacy protection	●	●	○	●	●
Usability	●	●	●	○	●
Low cost	●	●	○	●	●
Regulatory compliance	●	●	●	●	●
Resilience	●	●	●	●	●
Cooperation with market participants	●	●	●	●	○
Universal access	○	●	○	●	●
Cost efficiency	●	●	●	●	○
Interoperability	●	●	○	●	●

feasible to address all requirements simultaneously in one artifact. This hypothesis was also confirmed in the expert interviews that we conducted. The already proposed CBDC approaches do not enable fully private transactions, and related work in cryptography lacks a concrete design that can be used for the rigorous evaluation of their regulatory compliance and feasibility from the perspective of stakeholders (see Section 4). In this light, we designed and instantiated a solution that uses cryptographic techniques, i.e., ZKPs, to address these requirements and to enable a discussion with stakeholders. In particular, we aimed to develop a CBDC system that ensures cash-like privacy by design, where the transaction amount and the identities of involved transaction parties are not shared with any third party. Compliance with regulation is enforced by per-transaction, balance, and turnover limits for fully private payments.

Next, we designed and developed our CBDC system in cycles that iterated between conceptualization, instantiation, and internal evaluation. We present the overall CBDC system architecture, onboarding procedure, depositing, withdrawing, and fully private transaction processes in Section 4. We then discussed our CBDC proposal in internal discussion rounds and presented it to leading experts from various fields. An overview of the interviewed experts is depicted in Table 2. As one key result, our evaluation confirmed the feasibility and adequacy of our CBDC system. The adjustments to our CBDC architecture after each evaluation cycle are discussed in Section 5. Sixth, as a final step, we disseminated our key findings to the interviewees and other stakeholders, in particular, to decision-makers in central banks and regulatory authorities and to researchers. In addition, we publish the source code of our prototype for the proposed CBDC system on GitHub.[4]

[4]The repository can be accessed at: https://github.com/applied-crypto/cbdc.

Table 2.: Overview of interviewed experts.

Cycle	Field of expertise	Int. No.	Exp. No.	Role	Organization
1	Law	01	01	Assistant Professor	University
	Law	02	02	Specialist Payment Fraud	Europol
	CBDC	03	03	Senior Financial Sector Expert	ex-IMF
	Payments	04	04	Head of Payments	Central Bank
	Payments		05	Head of Digitalization and Payment Systems	
	CBDC	05	06	Chief Economic Advisor	SFB Technologies
	Cryptography	06	07	Global Managing Director Digital Assets	Accenture
2	IT	07	08	Technical Lead Digital Currencies	CBDC-developing company
	Economics		09	Business Lead Digital Currencies	
	Business		10	Product Manager Digital Currencies	
	Computer Science		11	Data Engineer Digital Currencies	
	Information Systems	08	12	Senior Researcher	Research Institute
	Law	09	13	Banking Supervision Expert	Banking Association
	Payments		14	Lead Digitalisation	
	Economics		15	Chief Economist	
	Law		16	Legal Lead	
	CBDC		17	CBDC Expert	
	Payments		18	CBDC Expert	
	Law	10	19	Professor	University
	IT / Business	11	20	Head of DLT Product	Bank
3	Economics	12	21	Alternate Member of the Governing Board	Central Bank
	Computer Science		22	Professor	University
	CBDC	13	23	Senior Economist	International Organization
	Payments		24	Senior Financial Market Analyst	
	Cryptography	14	25	Professor	University
	Economics		26	PhD Candidate	Télécom Paris
	CBDC	15	27	Head of Blockchain	Association
	Payments	16	28	Market Infrastructure Specialist	Central Bank
	Computer Science		29	IT Application Development Specialist	
4	Digital Identities	17	30	Head of SSI Consortium (IDUnion)	Main Incubator
	Law	18	31	Expert on CBDC and AML Regulation	University
	Business	19	32	Senior Manager Marketing & Public Affairs	Federal Printer
	CBDC		33	Senior Project Manager CBDC	
	Digital Identities		34	Senior Consultant Trusted Services	
	Cyptography		35	Senior Principal Security Systems	
	Computer Science		36	Technological Expert CBDC	
	Business		37	Senior Account Manager	
	Business		38	Regional Sales Director	
	Business		39	Senior Business Development Manager	
	Computer Science	20	40	CBDC Technology Expert	Central Bank
	Computer Science	21	41	Research Group Lead	Research Institute
	Economics		42	PhD Candidate	
	Finance	22	43	Research Group Lead	Consumer Protection Org.
	Economics		44	Advisor	

4. Our CBDC Design

4.1. Related Work

To conceptualize our design, we investigated cryptography-related research with a focus on privacy-oriented digital payment systems. Chaum (1983) founded the research stream on *e-cash* that aims to develop cryptography-based payment systems that are private by design and make payments untraceable. He proposes a design in which users need to exchange their received *digital banknotes* for new ones in a compulsory interaction with a trusted PSP, e.g., a bank. To make the spending and receiving of a specific banknote unlinkable, *blind signatures* hide the serial numbers of unique and thus distinguishable digital banknotes. However, copying and thus double-spending these digital banknotes cannot be prevented technically. Instead, to hold users accountable, the cryptographic protocol allows retrieving a user's identity that is hidden in the digital banknote from combining the information obtained in two different payments with the same digital banknote. Despite being computationally very efficient, this approach implies that, while the sender remains anonymous, the receiver is identified by its PSP, and the payment amount is transparent. Moreover, the design cannot practically enforce per-transaction limits. Moreover, turnover and balance limits are cumbersome to implement because in a CBDC system that involves multiple PSPs, as currently planned in most jurisdictions, a synchronization among PSPs would be required to detect double-spending attempts and to prevent users from visiting multiple PSPs to circumvent these limits. Furthermore, it is not clear how to implement basic programmability functionalities, such as interest payments, on top of this design.

Sander and Ta-Shma (1999) address some limitations of Chaum's approach, for instance, hiding the transaction amount by dividing it into discrete shares. Nevertheless, the PSP still learns the transaction amount paid and received in this epoch and the identity of the receiver. The first e-cash system that hides the identity of both sender and receiver and also the transaction amount without the need for a trusted third party was proposed by Camenisch, Hohenberger, and Lysyanskaya (2006). In addition to hiding the identity of the receiver and the transaction amount, this payment system guarantees the sender's anonymity as long as they do not double-spend and exceed a pre-defined per-transaction limit. Technically, the proposal is based on ZKPs that are mathematically tailored specifically to this use case. However, as this approach is still designed for digital banknotes and different payments in which the same individual is involved cannot be identified and not even linked, turnover limits cannot

be enforced. Additionally, it is challenging to extend this model to incorporate basic programmability features.

While none of the previous academic work seems to have received considerable adoption, the first practical emergence of privacy-oriented digital payment systems happened in the context of cryptocurrencies. Researchers developed privacy-oriented modifications of Bitcoin (Nakamoto, 2008) as the use of a public distributed ledger increases the need for a privacy-enhancing design. The two arguably most relevant academic studies in this context are Zerocoin, which hides transaction parties but not the transaction amount (Miers, Garman, Green, & Rubin, 2013), and Zerocash (Ben-Sasson et al., 2014), which additionally hides the transaction amount and laid the foundation for the well-known cryptocurrency Zcash. Both approaches use general-purpose ZKPs, in particular zk-SNARKs, to demonstrate that transactions are valid and respect agreed-upon monetary policies (e.g., no double-spends) without revealing transaction details. Garman et al. (2016) acknowledge that "from an investigative standpoint, Zerocash is no different than cash" (Garman et al., 2016, p. 81). However, they also point to the conflict between regulation and private payment systems, as both Zerocoin and Zerocash do not take into account legal frameworks that restrict the use of anonymous payments to comply with AML and CFT regulation. Consequently, Garman et al. (2016) sketch potential extensions of the Zerocash model to address regulatory constraints, such as considering specific jurisdictions involved in payments, per-transaction limits, the tracing of specific coins, and the payment of appropriate taxes. Since exceeding the limits enforced by ZKPs would imply that no transactions are possible and thus reduce the utility of the payment system substantially, they propose a tiered approach for a private payment system so that any transaction above the spending limit should be additionally signed by an authority that conducts KYC and AML checks. Similarly, Bontekoe (2020) follows the approach of modifying Zcash by adding an account-based system based on a previous KYC-process. This setup allows limiting the transactions of an account within a specific epoch that can ensure balance and turnover limits and, thus, regulatory compliance. The study also describes the possibility of enabling transactions that exceed a limit by verifiably encrypting the associated transaction data and allowing for subsequent checks through a dedicated authority.

Both Garman et al. (2016) and Bontekoe (2020) also mention that users' digital identity management based on digital certificates can take a valuable role in this context. Literature that considers the role of identity in privacy-oriented payment systems also emphasizes that a mechanism for revoking identities and accounts is required (Choi et

al., 2021). From a regulatory perspective, this resonates well with sanctions lists such as the Office of Foreign Assets Control (OFAC)'s Specially Designated Nationals.

In parallel to these developments, cryptographic research in the context of CBDCs continued pursuing Chaum's initial *asymmetric* approach to transaction privacy, offering anonymity for the sender but not for the receiver (Chaum et al., 2021; Tinn & Dubach, 2021). Chaum et al. (2021) base their design on blind signatures and hence employ similar techniques as earlier work by Chaum (1983) from a technical perspective, whereas the approach by Tinn and Dubach (2021) is based on zk-SNARKs. Veneris, Park, Long, and Puri (2021) propose a system for CBDC that makes transaction amounts transparent yet enables anonymity through identity escrow (which is consequently not trustless). They also add a hardware-based solution that provides essentially cash-like privacy but requires regular online settlement. In these designs, the privacy protection of the sender covers many practical requirements, especially when a consumer purchases a product from a business that needs to disclose its accounting transparently. However, as discussed in Section 1, particularly for payments between end-users, trustless, cash-like privacy may be desirable.

To date, several CBDC designs have been proposed that aim for regulatory compliance and, at the same time, preserve users' privacy – at least to some extent (see Table 3). However, these designs differ substantially in the extent to which privacy is guaranteed. In this context, *privacy by design* and *compliance by design* play a central role, representing systems where no trust in the operator is needed. Privacy by design and compliance by design can be achieved through extending a payment system like Zcash that allows fully private transactions with the possibility of person-related monthly turnover or per-transaction limits, as proposed by Bontekoe (2020); Garman et al. (2016). Yet, since the perception that enabling fully private transactions in digital form fundamentally contradicts AML and CFT regulation is still widespread, and regulators and central banks have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Armelius et al., 2021; Auer & Boehme, 2021), we present a holistic approach for a *privacy/compliance-by-design* CBDC in detail. We also show how digital identities can replace an efficient and privacy-preserving KYC process.

A frequently stated caveat when designing a private payment system is that it is necessary to consider not only the processing and storage of transaction data but also the metadata from the corresponding communication. For example, anonymizing the parties involved in a payment is hardly useful if the parties' static IP addresses are attached to transaction requests. In this context, sidechannel attacks have already been

Table 3.: Comparison of privacy-oriented CBDC solutions.

		ECB (2019) Anonymity Vouchers	Chaum et al. (2021) Blind signatures	Tinn and Dubach (2021) P-hybrid CBDC	Choi et al. (2021) Banking$^+$ and Cash$^+$	Veneris et al. (2021) CBDL	Our approach UAS model with ZKPs
Privacy	Anonymity of the sender	●	●	●	◐	●	●
	Anonymity of the receiver	●	○	○	○	●	●
	Privacy of transaction amount	○	◐	○	○	●	●
	Trustless privacy (privacy by design)	○	●	●	○	◐	●
Regulation	Regulation by design	○	●	○	○	◐	●
	Per-transaction limits	●	●	○	○	●	●
	Turnover and balance limits	●	●	○	○	●	●
	Anonymous onboarding	○	○	○	◐	○	●
Other	Offline payments	○	○	○	○	●	○
	Account-based system	○	○	●	◐	●	●
	DLT-based system	●	○	●	●	●	○
	Involvement of PSPs	●	●	◐	●	●	◐

used to de-anonymize shielded transactions in Zcash (Tramèr, Boneh, & Paterson, 2020). Consequently, related work on private payment systems (e.g. Tinn & Dubach, 2021) state that privacy on the networking layer must also be provided. However, given the existence of onion routing mechanisms as implemented by the Tor network (Dingledine, Mathewson, & Syverson, 2004), this problem can be considered solved and we will not refer to it further in the remainder of this paper.

4.2. High-Level Overview of our CBDC Architecture

Following Garman et al. (2016), our overall CBDC design is based on a two-tiered approach that supports both fully private and transparent CBDC payments. On the transparent side, the CBDC is distributed to users via PSPs – i.e., we use an intermediated model – and access is linked to an identity – i.e., we use an account-based model. Transaction data is stored in a centralized ledger. However, our system could also accommodate a distributed ledger. The disadvantage of a distributed ledger is that it may introduce challenges regarding performance, and privacy issues may arise when storing transaction- or account-related information on the transparent side. Thus, modifications, such as storing only PSPs' balances in an obfuscated way, may be necessary on

the transparent side. In general, the transparent side is flexible to implement central bank-specific designs, e.g., introducing a maximum limit on transactions in general or using a direct instead of an intermediated model.

We propose an account-based model that allows the funnelling of all of a user's transactions into one single account, which has many similarities to the approaches by Garman et al. (2016) and Bontekoe (2020). Recently, the security of an account-based approach in combination with ZKPs has been studied in more depth also in Wüst, Kostiainen, and Capkun (2021), including several improvements in terms of performance. Yet, none of these publications incorporates a connection to digital identity management, which — as we will discuss later — is likely indispensable to address some of the challenges related to sharing access to the privacy pool. In our system, users maintain their accounting privately, and a transaction corresponds to updating their private account and sending a ZKP that proves to the central bank that the transaction's expected policies have been met. The transaction amount then essentially corresponds to the delta, i.e., the difference of the previous and the new account states' balances. In essence, users store and manage their own accounts and prove the correctness of their local accounting to the operator of the ledger, i.e., the central bank.

4.3. The Privacy Pool in Detail

Most cryptocurrencies, such as Bitcoin, record every transaction, including the sender, receiver, transaction amount, and authorization (in terms of the sender's digital signature) on a public ledger (Zhang, Xue, & Liu, 2019). Additionally, transactions either point directly to one or more previously received but so far unspent coins (unspent transaction output (UTXO) model) or to the sender's and receiver's pseudonymous accounts with a public balance (account-based model). This setup generally allows one to track the transaction history of digital banknotes and corresponding metadata as well as additional information retrieved from exchanges to identify the involved parties in most transactions (Meiklejohn et al., 2016). Consequently, even if the ledger is not public, such a construction would allow the operator of the ledger to link transactions and correlate even pseudonymous accounts with real-world identities and, thus, to retrieve personal information.

In contrast, Zcash only stores cryptographic hashes of transactions in a ledger that hides details and ensures unlinkability. The payment details are only known to the sender and receiver. In particular, for every transaction (including its details), Zcash

applies two different one-way cryptographic functions to generate two unique but different outputs, i.e., (1) a *commitment* and (2) a *nullifier*. The commitment and nullifier hence essentially hide the same transaction details but are computationally infeasible to correlate without knowledge of the transaction details. To spend a previously received transaction in the form of a commitment, the corresponding nullifier is then published together with a ZKP that the nullifier corresponds to a previously published commitment (proof of knowledge of the joint pre-image) (Ben-Sasson et al., 2014). The central bank can then check that the associated commitment has not been spent before by checking whether this particular nullifier has been published before. The ZKP-based construction hence certifies that the sender has access to digital banknotes that have not been spent previously and that the amount of money spent equals the amount received, without the need to disclose *any* additional information about the transaction, such as sender, receiver, amount, or a direct reference to the previous transaction in which the banknote was received (Ben-Sasson et al., 2014).

The design of our privacy pool is based on the construction of Zcash with an append-only ledger that stores commitments and nullifiers. Adding a per-transaction limit to Zcash would be an easy task (Garman et al., 2016). However, our approach contains a crucial modification, replacing the UTXO-based model with an unspent account state (UAS) model to (privately yet provably) funnel all of a users' transactions into one account and hence enable a user to prove that balance and turnover limits are also satisfied. Each commitment and nullifier thus represents a unique account state. By publishing a commitment and a new nullifier, the previous account state is invalidated and a new one is created. To validate transactions, the central bank receives the associated commitments, nullifiers, and proofs of the correct update of the account state using ZKPs. In particular, the central bank verifies the validity of the ZKPs and that the nullifier has not been revealed before and then adds the transaction to the ledger by including the commitment and nullifier in the existing associated ledger. Due to the succinctness property of the ZKPs used, the workload for the central bank in verifying the transactions is very small (Ben-Sasson et al., 2013). The pre-image resistance of hash functions (with a pre-image of high entropy) and the unlinkability of commitments and nullifiers ensure that information that is revealed to the central bank is not sensitive (Ben-Sasson et al., 2014). Consequently, only the account owners know their respective transaction and account details, such as the amount, current balance, and turnover. Nevertheless, they can still prove compliance with predefined rules by using ZKPs. This approach ensures *privacy by design* and enables the creation of *proofs of compliance with limits by default*. In such a system, entities that maintain

the ledger only need to be trusted with respect to *integrity* and not with respect to protecting the users' privacy.

As common in the context of blockchains, for storing the commitments and enabling efficient proofs of inclusion for a commitment, we use Merkle trees. A Merkle tree is a cryptographic data structure that represents many entries by one identifier, i.e., the Merkle root (Merkle, 1987). This setup allows for proofs of inclusion that can be checked only by comparing with the Merkle root. The Merkle root changes with each new entry in the tree. By using these Merkle trees in combination with ZKPs, it is possible to prove that a transaction proposal refers to an existing commitment in the ledger without pointing to (and thus revealing) it. If the ledger is not public, for instance in a setup based on a centralized ledger or a permissioned blockchain, it is necessary to prevent the correlatability of commitments and nullifiers that may occur through querying the Merkle path of a specific commitment, e.g., through querying Merkle subtrees.

To ensure the integrity of the payment system and the compliance with turnover, balance, and per-transaction limits, our UAS model contains the following information that is stored in digital wallets and consequently is only known to their respective holder:

- Identity information: A public key and a digital certificate that includes the account owner's identity information (digital ID card, see, e.g., (Sedlmeir, Smethurst, Rieger, & Fridgen, 2021)).
- Balance: The balance of the account holder.
- Epoch turnover: An accumulation of all amounts of spending transactions in the current epoch (whose length can be specified by the regulator).
- Epoch reset: The last reset of the epoch (the last time the epoch turnover was set to zero).

This structure should be seen as an initial proposal that one can flexibly extend when more account details need to be checked for compliance, e.g., one could include the nationality or the type of the account, i.e., private or corporate.

For our CBDC system, we determine three core transaction types for interacting with the privacy pool, illustrated in Figure 3. In the following, we outline and discuss these three types of transactions that are related to the privacy pool – onboarding, semi-private transfers, and fully private transfers. All three types of transactions involve private inputs, public outputs, and a ZKP that connects them and proves the correctness of the local accounting to the central bank. The account owner provides

private inputs, which contain the already mentioned account information. This data is sensitive, therefore, it stays hidden. However, the private inputs are represented by the public outputs using one-way functions. Hence, the account owner shares the public output (commitment, nullifier, etc.) with the central bank along with a ZKP that ensures that the public outputs were computed from the private input according to the rules of the payment system.

4.3.1. Onboarding

Each account is tied to an individual cryptographic key pair and, using a digital certificate, to a government-issued identity. We assume that each user has a single digital ID card that is also bound to a cryptographic key pair. Many countries already integrate keypairs into their physical ID cards and even provide dedicated mobile apps to store the ID card in digital form, e.g., Germany and Estonia. To onboard a new user, the user has to create a cryptographic proof that they possess a valid ID that is not expired or revoked and that the initial commitment is deterministically derived from the ID card via a one-way function and contains the correct initial account entries. Therefore, each onboarding procedure of one individual is based on the same key pair and always results in the same commitment. In detail, a ZKP for onboarding and, thus, opening a new account would be structured as follows:

> **Onboarding**
>
> INPUT: Key pair and digital certificate (government-issued ID card), timestamp
>
> OUTPUT: Commitment to initial account state, timestamp, ID issuer's public key
>
> CHECKING THAT:
>
> - *the account holder controls a valid ID card (signature, binding, non-revocation)*
> - *the identity lodged in the account corresponds to the ID card*
> - *the account's initial balance and epoch turnover are set to zero*
> - *the account's initial epoch reset equals the timestamp*
> - *the commitment equals the hash of the signed initial account state*

The onboarding process works as follows: First, using their legit key pair and digital ID, users create the onboarding transaction, consisting of the initial commitment and the ZKP and send it to the central bank. Second, the central bank verifies the

Figure 3.: Transaction types in the privacy pool.

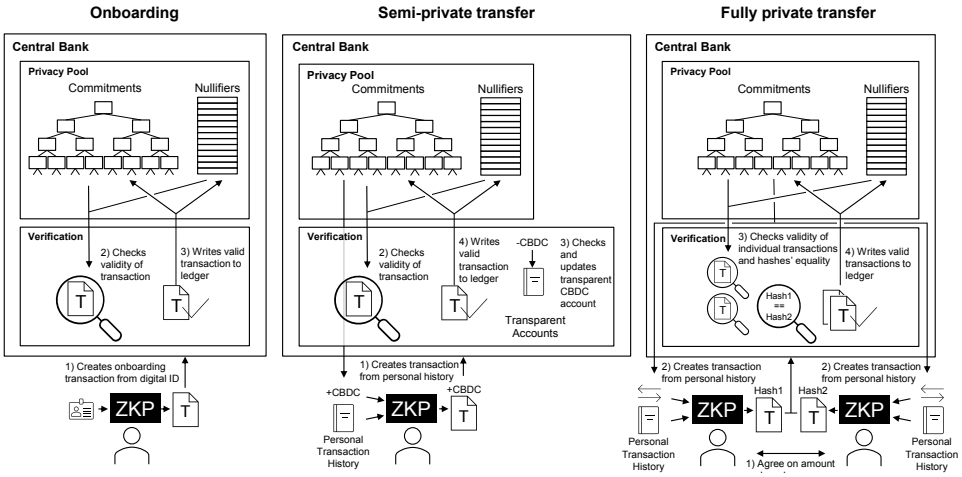

validity of the ZKP that refers to the commitment and adds the commitment to the current state of the ledger. Since the private key is only used for the generation of a signature that serves as private inpupt for the generation of the ZKP, and since the commitment represents the encryption of a hash, the central bank does not learn anything about the onboarded user. The central bank can also detect multiple attempts to create an onboarding transaction, as the commitment is deterministically derived from the ID card.

It is important to note that the anonymity of the onboarding process is *not* required to guarantee the anonymity of a user's subsequent transactions in the privacy pool due to the unlinkability of commitments and nullifiers. Consequently, the central bank could even demand that users that register on the privacy pool present some of their identity attributes, e.g., their nationality or the issuing authority of their digital ID.

Notably, it is not even necessary to check whether the commitment has already been used for previous onboarding to ensure that every user only has a single account: the commitment can only be spent once, independent on how often it is included in the Merkle tree because the corresponding nullifier is also unique. Nevertheless, to avoid attacks that spam the ledger with correct yet useless transactions, it may be useful to check for such collisions.

4.3.2. Semi-private transfers

Semi-private transfers describe the exchange of funds between an account in the privacy pool and a transparent account. These transfers include deposits and withdrawals from the same user so that a user can transfer CBDCs from their transparent account to their privacy pool account, or vice versa. In a semi-private transfer, a user combines an update of their account in the privacy pool with an update of their transparent account that is confirmed by the PSP. As it has to be ensured that the total money supply in the system is unchanged when money is deposited or withdrawn, the transaction amount, i.e., the difference in balances between the spent and created account state, must be disclosed to check whether it is equal to the counter-transaction on the transparent side.

In the following, we consider the depositing process as an example of a semi-private transfer. In more detail, the ZKP would be specified as follows:

Semi-private transfer

INPUT: Key pair, Merkle path of the previous commitment,

previous account state, amount

OUTPUT: Merkle root, nullifier, new commitment, (deposit/withdrawal) amount

CHECKING THAT:

- *the previous commitment is contained in the tree represented by the Merkle root*
- *the previous account state belongs to the previous commitment*
- *the nullifier equals the hash of the previous account state*
- *the new account state is correct (e.g., new balance = old balance + amount)*
- *the new account state complies with the rules*

  (e.g., positive balance, epoch turnover below turnover limit)
- *the commitment equals the hash of the signed new account state*

First, the user creates a new commitment and nullifier and attaches a ZKP proving that the account update is legitimate and corresponds to the public outputs. Second, the central bank verifies the ZKP and checks if the public outputs match the requirements, i.e., whether the Merkle root specified by the public outputs matches the Merkle tree of commitments in the ledger, whether the nullifier is not already included in the nullifier list maintained by the central bank, and whether the amount equals the transparent counter-transaction's amount. If all these requirements are satisfied, the central bank adds the new commitment and the nullifier to the ledger and notifies the client about the successful transaction.

The Merkle tree of commitments is append-only, and the mechanism that protects against double-spending, i.e., using the same old account state multiple times, is facilitated by the list of nullifiers. Consequently, the Merkle root does not necessarily have to be the most recent one. This allows users to create transactions that are accepted even if the Merkle root of the ledger has changed in the meantime due to a transaction by another user. Specifically, this option decreases the computational burden for the central bank, as it is sufficient to recompute the Merkle tree in larger epochs. Nevertheless, the epochs should not be too long, as the sequential processing of transactions by the same user requires a new Merkle tree that includes their previous commitment.

Processing this transaction, the central bank inevitably learns the amount of CBDC that is transferred from a transparent account to the privacy pool. However, since the

commitments are unlinkable, it is not possible for anyone except the holder of the account to further trace these CBDC units. The same applies to the nullifier that invalidates a previous account state from the Merkle tree of commitments without revealing which specific commitment it refers to. Thus, it is impossible to determine whether a specific semi-private transaction (deposit, withdrawal) was already followed by another semi-private or fully private transaction and, particularly, whether a deposit has already been spent.

In addition to using the transparent CBDC accounts for depositing money in the privacy pool, a user could use CBDC-specific automated teller machines (ATMs) to deposit cash into the privacy pool: The user would create a transaction proposal that contains the desired amount and a ZKP similar to the depositing process described above and send it to the ATM, e.g., via Bluetooth, Wi-Fi Direct, or NFC. Then, the user inserts the corresponding amount of cash directly into the ATM. The ATM confirms the receipt of cash and forwards the user's commitment, nullifier, and ZKP to the central bank. The central bank processes the data as described above. However, the user's transparent CBDC account is not involved in this case, making the depositing and withdrawal process anonymous. Instead, the transaction is conducted via the CBDC account linked to the ATM, which could be provided by a PSP or the central bank directly.

4.3.3. Fully private transfers

The transfer of funds in the privacy pool enables the private bilateral exchange of CBDC. A transfer consists of two transaction proposals, i.e., an individual payment instruction provided by both the sender and receiver. First, the sender and receiver need to agree on the amount of CBDC that should be transferred and on a common randomly generated number, i.e., a nonce, to increase entropy and to be able to link the two update proposals. Next, each participant creates their individual transaction proposal, including the corresponding commitment, nullifier, and ZKP, as they would in the semi-private case, with the difference that the transaction amount is not revealed but hidden by a hash (salted with the nonce). In detail, the ZKP of each of the two involved users looks as follows:

Fully private transfer

INPUT: Key pair, Merkle path of a previous commitment,

previous account state, (transaction) amount, nonce, role (sender/receiver)

OUTPUT: Merkle root, nullifier, new commitment, hash of amount | nonce, role

CHECKING THAT:

- *the previous commitment is contained in the tree represented by the Merkle root*
- *the previous account state belongs to the previous commitment*
- *the nullifier equals the hash of the previous account state*
- *the new account state is correct*

  (e.g., new turnover = old turnover + amount if the role is "sender")
- *the new account state complies with the rules*

  (e.g., positive balance, epoch turnover below turnover limit)
- *the commitment equals the hash of the signed new account state*
- *the hash of amount | nonce was computed correctly*

Either the sender or the receiver batches both individual transaction proposals (including their public outputs and ZKP) and sends them to the central bank. Afterward, the central bank validates the integrity of this data through verifying the ZKP that the user generated and compares the outputs with the current state of the ledger. In particular, the central bank checks whether both Merkle roots correspond to the Merkle root of the current ledger and that the two nullifiers are not yet part of the nullifiers' list. Then, the central bank verifies whether the hashes of the value concatenated with the nonce are the same in both transactions. This check guarantees that the amount deducted from one account is equal to the amount added to the other account, without the central bank learning the actual amount due to the pre-image resistance of hash functions. Furthermore, the central bank verifies that the roles specified in the two update proposals are different, i.e., there is one sender and one receiver. Finally, the central bank adds the two new commitments to the Merkle tree and the two nullifiers to the list and notifies the client application about the successful transfer.

Overall, the central bank only learns that a valid transaction was conducted and which two new commitments and nullifiers are involved. However, these hashes are neither related to each other in any further way nor could they be associated with the

hashes of previous or future transactions due to the unlinkability guarantees achieved through the system's construction based on commitments, nullifiers, and ZKPs. Thus, the transfers facilitated by the UAS do not provide any information about the individuals themselves, their account balances, or the amount of the transaction, and are therefore fully private.

5. Evaluation

As suggested by Hevner et al. (2004), Peffers et al. (2007), and Peffers, Tuunanen, and Niehaves (2018), we asked key stakeholders to evaluate our IT artifact and to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs. To obtain valuable insights from expert interviews, we followed the recommendations for conducting qualitative interviews by Myers and Newman (2007). In this context, we minimized social dissonances between the researcher's team and the interviewees by first introducing all interview participants to each other. To obtain a rich collection of perspectives, we talked to professionals from various backgrounds, including experts from regulation, cryptography, central banking, identity, and payments. We also made use of specific interview models, such as the waterfall technique, by first motivating our research project, providing context and general definition. Next, we provided a high-level architecture overview and a technical deep dive to improve disclosure in our interviews.

We also discussed the potential risks that may occur using our CBDC with experts and demonstrate adequate mitigation measures. In addition, we presented our overall CBDC system, including the onboarding procedure, semi-private, and fully private transfers. In total, we conducted 22 interviews with a total of 44 international experts with profound expertise in the fields of law, economics, technology, and others (see Table 2) to reinforce the rigor of our methodological approach. Additionally, we sought to gain insights from key stakeholders regarding their key requirements for a CBDC and to receive feedback on our proposed CBDC system. Both the diversity and CBDC-specific expertise of the interviewees provided us with valuable feedback to iteratively adjust and improve our CBDC system. In addition, feedback was collected via internal discussions within the research team's organizations and presentations at various events to test the general feasibility of our approach, thereby continuously improving the artifact (Peffers et al., 2007).

During the first evaluation cycle, the experts confirmed that our design is highly innovative, that it addresses the relevant need for privacy, and that the implemented

per-transaction and turnover limits on fully private transactions fit smoothly into existing regulatory frameworks from an AML and CFT compliance perspective. Many experts were surprised or even impressed by the technical capabilities of ZKPs and the maturity of the design (experts 1, 2, 3, 5, 6). In addition, the capability of our CBDC system to flexibly accommodate possible future regulatory changes, e.g., dynamically adaptable thresholds, was considered highly valuable (expert 1). Expert 2 confirmed that our approach of enforcing turnover limits on anonymous transactions would be in line with the 5th European Anti-Money Laundering Directive. The expert further stressed that tracing these small-scale transactions is neither required nor desirable from a law enforcement perspective. Expert 3 emphasized that, in some jurisdictions, mistrust towards the government is considerably high, indicating that a privacy-by-design approach may be helpful to improve broad adoption of a digital currency. He also pointed out that, for the usability of the payment system, it is important to reconcile and settle transparent transactions seamlessly when the limits in the privacy pool are exceeded. However, experts 4 and 5 noted that, although our adaption of the Zcash architecture is a "good trick", it may still be challenging to convince skeptical users that a solution based upon this model is, in fact, fully private. Against this background, providing the code open-source to facilitate independent audits by cryptographers and consumer protection organizations and taking large educational efforts may be required to increase trust in such a cryptography-based privacy-by-design solution.

Experts 3 and 5 noted that our design could also interact smoothly with the current account-based banking systems. Moreover, expert 6 considered our approach of "digital cash" to be intuitive, particularly illustrated through the possibility of depositing and withdrawing CBDC via an ATM. Expert 6 also mentioned that our design, in which the central bank performs the highly automatable task of verifying ZKPs and maintaining the ledger but does not need to conduct resource-intensive KYC processes, fits the roles that central banks aim for in CBDC systems (see, e.g., European Central Bank (2020a)).

However, the experts in the first evaluation cycle also raised two major concerns related to identity management and addressing the needs of corporations additional to those of private end-users, as they may require different limits and identity concepts. Specifically, expert 1 noted that the concept needs to be able to handle the recovery of funds in the case of theft or lost access to the mobile phone. Also, it must provide a way to disable a privacy pool account in the case of blacklisting, e.g., if an individual is put on a sanctions list. Moreover, while our approach to provide one wallet for both a digital identity and payments was considered efficient and appealing from a

user perspective (experts 6 and 7), experts 1, 6, and 7 consider the identity-based concept complex and difficult to implement in practice. According to the experts, it is unlikely that such a digital identity can be bootstrapped in the short-term, and expert 7 even considered the availability of a standardized, unique digital identity across multiple European member states a task that is as complex as introducing a CBDC. According to the expert, a particular challenge is that many citizens in the EU have multiple nationalities and, thus, various ID cards, which makes ID cards less suitable for guaranteeing that any citizen can only open and control one account in the privacy pool. Consequently, experts 1, 3, and 7 suggest adding intermediary-based onboarding procedures as another venue to our design.

We incorporated this feedback in our design through the following modifications:

- We added the opportunity for a joint (centralized or decentralized) ledger managed by PSPs that contains identifiers that are deterministically derived from citizens' identity, such as Hash(first name | last name | date of birth). A PSP can then sign an onboarding transaction that proves the possession of a digital certificate issued by the PSP instead of an ID card. As we already discussed in Section 4.3, the anonymity guarantees of private payments do not depend on the privacy of the onboarding process in our solution. Consequently, while the PSP learns that a specific user has registered for the privacy pool, this approach does not compromise the opportunity to conduct fully private transactions.
- We incorporate periodic proofs of non-expiration and non-revocation of the ID card (or the PSP-provided digital certificate) into our design whenever the epoch reset is performed. This modification ensures that once per epoch the user needs to prove that their ID card is still valid and hence allows to block accounts connected to an ID card that has been already revoked, e.g., in the event of loss or theft or the inclusion of an individual on a sanctions list.

In the next cycle, the experts in interview 7 pointed out that metadata, such as IP addresses, need to be taken into account for analysing privacy, and hence that pseudonymization is not sufficient to ensure privacy. This diagnosis confirms our path of ensuring the perfect unlinkability of transactions. They further noted that the *verifiability* of transactions is an important feature, i.e., a user should have the opportunity to prove that a payment did indeed happen, e.g., in the case of a lawsuit. In fact, our design already provides this capability, as a user stores their account history on their local device and can consequently reveal two consecutive previous account states in combination with the transaction confirmation signed by the central bank to demon-

strate that a payment was indeed conducted with the claimed details. Furthermore, the bilateral communication between the sender and recipient preceding a transaction, where they agree upon, a transaction amount and a transaction ID, can also be used to make the parties accountable bilaterally if desired by the involved parties, e.g., by revealing parts of their ID cards to the counterparty. The experts also appreciated the capability of our approach to account for embargo lists that prevent a sanctioned user from registering and further using their account through periodic checks of expiration and non-revocation of their ID card.

In interview 9, one expert noted that a balance limit might not be necessary, as a transaction cannot be larger than the turnover in a specific epoch. Nevertheless, the feasibility of balance, per-transaction, and turnover limits can account for the particular needs of various regulators. For instance, if any of these limits is not required in a particular jurisdiction, our system can be implemented easily without such a limit. The experts also pointed out that reversing a transfer must be possible if both parties agree to do so. Moreover, they emphasized the importance to publish the source code of a solution to be able to gain broad acceptance by the public and allow for auditing by consumer protection organizations. This in turn, will, ultimately increase trust in centrally-operated payment systems and also address the concerns regarding education on privacy-by-design approaches as raised by experts 4 and 5. Considering the opportunity to prove that a payment happened, the experts in interview 9 and expert 19 added that, while this proof is indeed a desirable feature, it is also crucial to implement the possibility to delete these records to prevent measures that aim to force users to reveal them (e.g., considering coercive detention or even torture).

While confirming that a limit-based approach is suitable to make fully private payments regulatorily compliant, expert 19 argued that the current limits for cash are relatively low, and further reducing these limits may be difficult to justify, particularly in view of privacy-oriented regulatory norms (see Section 1) and "shadow economies" that may appear when limits are too low to be practical. Moreover, the expert acknowledged that the compliance by design may even help justify higher limits and that, compared to the asymmetric privacy approach by Chaum et al. (2021), our design provides true cash-like privacy. Expert 20 confirmed the suitability of ZKPs for a privacy-oriented yet regulatorily compliant form of money from a technical perspective. The expert also pointed to similar approaches based on the Ethereum blockchain that aim to create private, account-based forms of money but yet do not address regulatory constraints. He also emphasized that due to the complexity of the privacy-oriented cryptographic

tools, such as ZKPs, only a few stakeholders that work on CBDC are indeed aware of their technical capabilities.

An issue that was already briefly mentioned by experts 2 and 3 in the first cycle and also caused intensive discussions with experts 12, 15, 18, and 19 in the second cycle refers to the integration of businesses in our CBDC system. A business should not be allowed to spend large amounts of money in the privacy pool as it could potentially evade documentation, such as avoiding tax declarations, or support money laundering on large scale. We discussed two different options to incorporate businesses in our CBDC design: The first approach is to allow them to open a private account via a business ID, e.g., provided by tax authorities or ingrained in trade registers. Our design is sufficiently flexible in assigning other limits to businesses and distinguishing between receiving and spending transactions or withdrawals of a business' privacy pool account to their transparent account. The second approach is based on the concept of asymmetric privacy considered in Chaum et al. (2021) and Tinn and Dubach (2021), tailored to business-to-customer interactions. Indeed, by extending the semi-transparent transactions to not only include deposits and withdrawals between the accounts of one single entity, it is possible to hide the identity of the buyer and only disclose the transaction amount and the receiving business. One essential advantage of this approach may be the opportunity to be able to choose whether to keep the sender's or the receiver's identity private. For example, when the purchase of a potentially embarrassing but legal product needs to be refunded, it may be desirable to hide the identity of the end-user that previously paid anonymously.

The third design cycle further confirmed the suitability of a ZKP-based system for enabling fully private transactions and aligning with regulatory constraints through enforcing limits (experts 22, 25, and 27). Besides, expert 25 emphasized that in general, ZKPs are a "perfect tool" to align privacy and compliance requirements, but also acknowledged that, from a cryptographic perspective, it may be useful to use other forms of ZKPs, such as zk-STARKs (Ben-Sasson et al., 2018), in a potential implementation, specifically as they are regarded post-quantum secure. Expert 28 highlighted that a flexible design that allows for remuneration of CBDC deposits is desirable from a central bank perspective, which illustrates that basic programmability features based on our account-based design is advantageous. Expert 23 also acknowledged that, in contrast to solutions such as the ECB's anonymity vouchers (European Central Bank, 2019), our design can provide true, cash-like privacy. On the other hand, expert 22 expressed concerns that the presumably low limits for fully private transactions might

imply that, despite a small share of transactions being fully private, most transactions will eventually be transparent, and hence privacy is not considerably improved overall.

One focus of this design cycle was the mitigation of risks that can potentially arise from our design. On the one hand, expert 22 warned that criminals could abuse the fully private payment system by getting access to several user accounts by purchasing accounts from other users on a black market or via blackmailing. In such a case, the effective limits could be circumvented by possessing a considerable number of accounts in the privacy pool. Although these problems can be mitigated through digital IDs that are bound to secure hardware (expert 22) and connecting the ID to other ID systems is a "great idea" (expert 23), expert 22 still pointed out that the use of secure hardware for storing keys conflicts with recovery capabilities. Experts 28 and 29 employed with a central bank also raised security concerns, particularly related to the high level of obfuscation of transaction-related information inside the privacy pool. Specifically, central banks require strong guarantees that, even if the implementation of the ZKP has a security gap, the extent to which this gap allows illicit activities or harms the monetary system remains marginal. In this regard, they also referred to an implementation error in Zcash that was detected only in 2019 and that would have allowed to "create money out of thin air" (Fortune, 2019).

We addressed the concerns raised in the third cycle by

- conceptualizing backup capabilities with the use of secure hardware through the precautionary creation of a transaction that withdraws all funds that a user has deposited from the privacy pool to their transparent account. The user can then store this recovery backup in the cloud, potentially in encrypted form. This transaction does not provide a proof of non-revocation and non-expiration, as this would quickly be outdated, but these proofs can be given through an in-person visit to the PSP that manages the recovery.
- mitigating risks by
    - pointing out the all-or-nothing transferability that the combination with a digital ID bound to secure hardware allows, i.e., if a user wants to sell their private account, this implies that they need to give away their complete digital identity, meaning that they can no longer use their digital ID card, including all credentials associated with it, such as digital credit cards, diplomas, or health insurances that are bound to this ID.
    - periodically closing and clearing the privacy pool so that users need to transfer their private funds to their transparent account. This process is

also helpful to improve performance, as the Merkle tree and the list of commitments do not need to grow quasi infinitely.

- suggesting a hybrid approach with moderate limits for fully private transactions that are primarily meant for customer-to-customer interactions and for semi-private transactions that are primarily meant for business-to-customer interactions.

By implementing these changes, the central bank can detect some of the most critical issues imposed by potential flaws in the implementation of ZKPs, e.g., if after closing the privacy pool, the money supply would be higher than expected. Via small limits for fully private payments, the severity of the impact of selling accounts, which cannot be excluded with certainty, can be mitigated.

As our first three design cycles raised the concern that the most fragile component of our construction is its dependence on a universal digital ID, we conducted another design cycle that – besides discussing the appropriateness of our risk mitigation measures – specifically includes experts on digital ID systems. Expert 30 agreed that our proposed risk mitigation strategies and the combination with a digital ID might be a useful approach, although the integration with secure hardware may be considered inflexible from an end-user perspective, as they cannot access their privacy pool from multiple devices. He also suggested checking the validity of the user's ID in every transaction by default to prevent misuse. Moreover, the experts in interview 19 employed at a federal printer in an advanced economy confirmed that the digital ID-based approach is not only more elegant but will also be possible relatively soon as many federal printers are working on digital IDs and general-purpose ZKPs that are used in such systems.

Within this group, expert 35 confirmed that our proposed backup capabilities seem suitable to recover funds when losing the mobile phone. The experts also stressed that the first digital IDs that integrate with secure hardware on users' devices will be available soon and that they are potentially the only way to efficiently impede the theft of digital IDs or their sharing or selling on a black market. The expert also found the combination of fully private transactions between end-users and semi-private transactions between end-users and businesses very suitable. Moreover, expert 35 pointed out that the security of implementing ZKPs has increased substantially due to cryptographic progress in the last years, so ZKP-related security gaps discussed previously are relatively unlikely today if state-of-the-art guidelines are considered. Further, the group noted that detecting security flaws in well-audited ZKP is significantly less promising

for criminals than counterfeiting paper-based money. Finally, expert 40 confirmed that our risk mitigation design may be a promising proposal to consider for central banks and could not identify obvious shortcomings. He particularly appreciated the coupling to a hardware-bound digital identity. However, as with every design proposal for a critical infrastructure, he emphasized that a thorough risk analysis would be required that is beyond the scope of this work. Nevertheless, the expert highlighted that our proposal is well presented and visualized also for non-technical experts, presenting a solid discussion basis for more in-depth analyses, and that our ZKP-based flexible design is a promising approach for the future. Experts 41 and 42 confirmed the suitability as well, emphasizing the insufficient communication between institutions that work on CBDC designs and the state of the art in cryptographic research, so the knowledge transfer via our DSR approach is highly valuable. Furthermore, experts 41 and 42 confirmed that our design incorporates privacy by design well and poses an attractive solution from the perspective of privacy-seeking users. Experts 40 and 41 also noted that, although our software-based solution cannot provide offline payments, our proposal avoids the inherent centralization of risks that comes with a hardware-based approach that may be particularly relevant in regions that do not have local businesses that develop secure hardware (expert 41). For example, there is currently no provider for secure hardware on smartphones where the manufacturing takes place in Europe. Expert 41 also pointed out that abuse cannot be excluded completely even when using hardware-bound digital identities. For instance, users could install proxies on their smartphone that transact on behalf of criminals. Yet, abuse is considerably more complex to organize than it is today with cash.

Since no considerable needs for improvements were brought forward by the experts in the fourth design cycle, and the experts interviewed in this cycle represented diverse fields, including central banks, digital identity issuers, and academics with interdisciplinary expertise from cryptography and economics, we concluded that we have reached a high level of saturation in the design and development of our artifact (Peffers et al., 2007). The critical feedback we received from key stakeholders positively influenced the design of our CBDC system, allowing us to continuously improve our artifact and thus to ultimately answer our research question.

6. Conclusion

In this paper, we followed a rigorous DSR approach to develop and evaluate a CBDC system based on an unspent-account-state model. To this end, we make use of re-

cent advancements in cryptography, especially related to ZKPs. Contrary to common beliefs (e.g. Armelius et al., 2021; Auer & Boehme, 2021), we demonstrate that a software-based CBDC system can support full privacy while addressing constraints related to AML and CFT regulation by imposing limits on anonymous payments. In our system, both privacy and compliance are provided *by design*, i.e., end-users do not have to trust third parties for preserving privacy and conducting compliance checks, as transaction data are stored only on the end-users' devices (trustless privacy). We assess the feasibility and suitability of our technical artifact in 22 interviews with 44 leading experts from various fields, including regulation, computer science, cryptography, central banking, and payments.

In addition to enabling full privacy and regulatory compliance, our ZKP-based CBDC system provides novel applications for end-users that go beyond existing forms of fully private money and especially cash. For instance, when faced with a legal accusation such as money laundering, our system enables users to reveal their payment history and provide evidence for its integrity and completeness. Thereby, users can address accusations, e.g., by proving that a payment did or did not happen.

Our proposed solution can also be used in more general applications, e.g., for account-based IT artifacts that enable full privacy by design while addressing certain compliance requirements. For example, future systems for documenting and exchanging carbon certificates between organizations and/or individuals will likely require high privacy guarantees, as well. Moreover, our solution can be applied in decentralized settings. Illustrated by the progress in the fields of blockchain technology in general and decentralized finance in particular, one can observe a tendency towards decentralizing the financial system. Specifically, tokenization promises to simplify the exchange of ownership (Sunyaev et al., 2021), but so far, existing blockchain-based solutions for managing and exchanging these tokens have not focused on stakeholders' (or even regulatory, e.g., the general data protection regulation (GDPR)) privacy requirements and regulation has rather focused on service providers that provide access to token-related services than end-users directly. Privacy-by-design approaches are the only way to establish full privacy in such blockchain-based systems because there are, by definition, no trusted third parties, and data is replicated among many different nodes. Our approach of leveraging ZKPs makes them a digital substitute for physical hardware-based approaches without a single point of failure to achieve integrity and compliance while storing data only on end-users' devices, as the correct local accounting is ensured through providing cryptographic proofs instead of hardware-based attestation.

Detecting an error in a well-audited cryptographic protocol seems less likely than a successful attack on a single secure hardware device.

Our artifact provides a starting point that balances privacy and compliance and may provide many avenues for future research. So far, we have presented our design to experts and instantiated the core logic in a technical implementation. Expert 22 also stated that it is important to focus on key requirements as most design criteria tend to have trade-offs and that it is almost impossible in one research project to implement interfaces to end-users and businesses. Thus, our focus on full (cash-like) privacy and regulatory compliance is a reasonable first step. However, there are additional important CBDC design dimensions, including security, scalability, and cost, that have to be considered. Thus, there is a need for future research for a rigorous evaluation of the extent to which our IT artifact can potentially also address these other important CBDC design dimensions. Specifically, besides more detailed analyses on performance, future analyses may involve other aspects related to user experience. For example, the implications of the added complexity of a two-tiered approach, the limited recovery options, and the integration of multiple devices require innovative solutions that shield complexity from the user and survey their impact on usability. Finally, future research could study the interplay of our design with potential extensions, such as using secure hardware for facilitating offline payments, that may, however, potentially come with restricted privacy guarantees to account for mitigating the related security challenges.

References

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, *54*(2), 442–92.

Allen, S., Čapkun, S., Eyal, I., Fanti, G., Ford, B. A., Grimmelmann, J., ... others (2020). *Design choices for central bank digital currency: Policy and technical considerations* (Tech. Rep.). National Bureau of Economic Research. Retrieved 2022-01-10, from https://www.nber.org/system/files/working_papers/w27634/w27634.pdf

Armelius, H., Claussen, C. A., & Hull, I. (2021). *On the possibility of a cash-like CBDC.* Retrieved 2022-01-10, from https://www.riksbank.se/globalassets/media/rapporter/staff-memo/engelska/2021/on-the-possibility-of-a-cash-like-cbdc.pdf

Auer, R., & Boehme, R. (2021). *Central bank digital currency: The quest for minimally invasive technology.* Retrieved 2022-01-10, from https://www.bis.org/publ/work948.pdf

Auer, R., & Böhme, R. (2020). *The technology of retail central bank digital currency.* Retrieved 2022-01-10, from https://www.bis.org/publ/qtrpdf/r_qt2003j.pdf

Bank for International Settlements, Bank of Canada, European Central Bank, Bank of Japan, Sveriges Riksbank, Swiss National Bank, ... Federal Reserve Board of Governors (2020). *Central bank digital currencies: Foundational principles and core features.* Retrieved 2022-01-10, from https://www.bis.org/publ/othp33.pdf

Bank of Canada. (2020). *Contingency planning for a central bank digital currency.* Retrieved 2022-01-10, from https://www.bankofcanada.ca/2020/02/contingency-planning-central-bank-digital-currency/

Bank of England. (2020). *Central bank digital currency: Opportunities, challenges and design.* Retrieved 2022-01-10, from https://www.bankofengland.co.uk/paper/2020/central-bank-digital-currency-opportunities-challenges-and-design-discussion-paper

Bech, M. L., & Garratt, R. (2017). Central bank cryptocurrencies. *BIS Quarterly Review September*.

Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. (2018). *Scalable, transparent, and post-quantum secure computational integrity.* Retrieved 2022-01-10, from https://eprint.iacr.org/2018/046.pdf

Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In *IEEE symposium on security and privacy* (pp. 459–474).

Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., & Virza, M. (2013). SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Annual cryptology conference* (pp. 90–108).

Biryukov, A., Khovratovich, D., & Pustogarov, I. (2014). Deanonymisation of clients in Bitcoin P2P network. In *Proceedings of the 2014 ACM SIGSAC conference on computer*

and communications security (p. 15–29). ACM.

Boar, C., & Wehrli, A. (2021). *Ready, steady, go? – Results of the third BIS survey on central bank digital currency.* Retrieved 2022-01-10, from https://www.bis.org/publ/bppdf/bispap114.pdf

Bok, S. (1989). *Secrets: On the ethics of concealment and revelation.* Random House Digital, Inc.

Bontekoe, T. (2020). *Balancing privacy and accountability in digital payment methods using zk-SNARKs* (Master's thesis, University of Twente). Retrieved 2022-01-10, from http://essay.utwente.nl/83617/1/Bontekoe_MA_EEMCS.pdf

Camenisch, J., Hohenberger, S., & Lysyanskaya, A. (2006). Balancing accountability and privacy using e-cash. In *International conference on security and cryptography for networks* (pp. 141–155).

Chaum, D. (1983). Blind signatures for untraceable payments. In *Advances in cryptology* (pp. 199–203).

Chaum, D., Grothoff, C., & Moser, T. (2021). *How to issue a central bank digital currency.* Retrieved 2022-01-10, from https://arxiv.org/ftp/arxiv/papers/2103/2103.00254.pdf

Cheng, J., Lawson, A. N., & Won, P. (2021). *Preconditions for a general-purpose central bank digital currency.* Retrieved 2022-01-10, from https://www.federalreserve.gov/econres/notes/feds-notes/preconditions-for-a-general-purpose-central-bank-digital-currency-20210224.htm

Choi, K. J., Henry, R., Lehar, A., Reardon, J., & Safavi-Naini, R. (2021). *A proposal for a Canadian CBDC.* Retrieved 2022-01-10, from https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3786426_code244292.pdf?abstractid=3786426&mirid=1

Dinev, T., Xu, H., Smith, J. H., & Hart, P. (2013). Information privacy and correlates: an empirical attempt to bridge and distinguish privacy-related concepts. *European Journal of Information Systems*, *22*(3), 295-316.

Dingledine, R., Mathewson, N., & Syverson, P. (2004). Tor: The second-generation onion router. In *13th USENIX security symposium.* USENIX Association.

Dold, F. (2019). *The GNU Taler system: Practical and provably secure electronic payments* (Doctoral dissertation, Université Rennes 1). Retrieved 2022-01-10, from https://tel.archives-ouvertes.fr/tel-02138082/file/DOLD_Florian.pdf

European Central Bank. (2019). *Exploring anonymity in central bank digital currencies.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.mipinfocus191217.en.pdf

European Central Bank. (2020a). *Report on a digital euro.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/euro/html/digitaleuro-report.en.html

European Central Bank. (2020b). *Study on the payment attitudes of consumers in the euro area.* (https://www.ecb.europa.eu/press/pr/date/2020/html/ecb

.pr201202~0645677cf6.en.html, accessed 2021-04-06)

European Central Bank. (2021a). *A digital euro to meet the expectations of Europeans.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/press/key/date/2021/html/ecb.sp210414_1~e76b855b5c.en.html

European Central Bank. (2021b). *ECB publishes the results of the public consultation on a digital euro.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210414~ca3013c852.en.html

European Central Bank. (2021c). *Eurosystem launches digital euro project.* Retrieved 2021-07-17, from https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210714~d99198ea23.en.html

European Convention. (2000). *Charter of fundamental rights of the European Union.* Retrieved 2022-01-10, from https://www.europarl.europa.eu/charter/pdf/text_en.pdf

European Court of Human Rights. (1950). *European convention of human rights.* Retrieved 2022-01-10, from https://www.echr.coe.int/documents/convention_eng.pdf

European Union. (2018). *Directive (EU) 2018/843 of the European parliament and of the council.* Retrieved 2022-01-10, from https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32018L0843&from=EN

Fauzi, P., Meiklejohn, S., Mercer, R., & Orlandi, C. (2019). Quisquis: A new design for anonymous cryptocurrencies. In *International conference on the theory and application of cryptology and information security* (pp. 649–678).

Fortune. (2019). *Zcash discloses vulnerability that could have allowed 'infinite counterfeit' cryptocurrency.* Retrieved 2022-01-10, from https://fortune.com/2019/02/05/zcash-vulnerability-cryptocurrency/

Garman, C., Green, M., & Miers, I. (2016). Accountable privacy for decentralized anonymous payments. In *International conference on financial cryptography and data security* (pp. 81–98).

Garratt, R. J., & Van Oordt, M. R. (2021). Privacy as a public good: A case for electronic cash. *Journal of Political Economy*, *129*(7).

Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, *18*(1), 186–208.

Gregor, S., & Hevner, A. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*, 337-355.

Grothoff, C., & Dold, F. (2021). *Why a digital euro should be online-first and bearer-based.* Retrieved 2021-07-18, from https://taler.net/papers/euro-bearer-online-2021.pdf

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*(1), 75–105.

Iivari, J. (2015). Distinguishing and contrasting two strategies for design science research. *European Journal of Information Systems*, *24*(1), 107-115.

Kahn, C. M. (2018). *Payment systems and privacy.* Retrieved 2022-01-10, from https://

papers.ssrn.com/sol3/papers.cfm?abstract_id=3315945

Kahn, C. M., McAndrews, J., & Roberds, W. (2005). Money is privacy. *International Economic Review*, *46*(2), 377–399.

Kannengießer, N., Lins, S., Dehling, T., & Sunyaev, A. (2020). Trade-offs between distributed ledger technology characteristics. *ACM Computing Surveys*, *53*(2).

Kiff, J., Alwazir, J., Davidovic, S., Farias, A., Khan, A., Khiaonarong, T., ... others (2020). *A survey of research on retail central bank digital currency.* Retrieved 2022-01-10, from https://www.imf.org/-/media/Files/Publications/WP/2020/English/wpiea2020104-print-pdf.ashx

Kubach, M., & Sellung, R. (2021). On the market for self-sovereign identity: Structure and stakeholders. In *Open identity summit 2021.* Gesellschaft für Informatik eV.

Lane, T. (2020). *Money and payments in the digital age.* Retrieved 2022-01-10, from https://www.bis.org/review/r200311d.pdf

March, S. T., & Smith. (1995). Design and natural science research on information technology. *Decision Support Systems*, *15*(4), 251–266.

Mastercard. (2021). *Mastercard and island pay launch world's first central bank digital currency-linked card.* Retrieved 2021-07-18, from https://www.mastercard.com/news/press/2021/february/mastercard-and-island-pay-launch-world-s-first-central-bank-digital-currency-linked-card/

Mattke, J., Hund, A., Maier, C., & Weitzel, T. (2019). How an enterprise blockchain application in the US pharmaceuticals supply chain is saving lives. *MIS Quarterly Executive*, *18*(4).

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2016). A fistful of Bitcoins: Characterizing payments among men with no names. *Communications of the ACM*, *59*(4), 86–93.

Merkle, R. C. (1987). A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques* (pp. 369–378).

Miers, I., Garman, C., Green, M., & Rubin, A. D. (2013). Zerocoin: Anonymous distributed e-cash from Bitcoin. In *IEEE symposium on security and privacy* (p. 397-411).

Myers, M., & Newman, M. (2007). The qualitative interview in is research: Examining the craft. *Information and Organization*, *17*, 2-26.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system.* Retrieved 2022-01-10, from http://bitcoin.org/bitcoin.pdf

Odlyzko, A. (2004). Privacy, economics, and price discrimination on the internet. In *Economics of information security* (pp. 187–211). Springer.

Partala, J., Nguyen, T. H., & Pirttikangas, S. (2020). Non-interactive zero-knowledge for blockchain: A survey. *IEEE Access*, *8*, 227945-227961.

Peffers, K., Tuunanen, T., & Niehaves, B. (2018). Design science research genres: Introduction to the special issue on exemplars and criteria for applicable design science research. *European Journal of Information Systems*, *27*(2), 129-139.

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77.

Pocher, N., & Veneris, A. (2021). *Privacy and transparency in CBDCs: A regulation-by-design AML/CFT scheme.* Retrieved 2022-01-10, from https://www.eecg.utoronto.ca/~veneris/21icbc2.pdf

Rossi, M., Mueller-Bloch, C., Thatcher, J. B., & Beck, R. (2019). Blockchain research in information systems: Current trends and an inclusive future research agenda. *Journal of the Association for Information Systems*, *20*(9), 1388–1403.

Sander, T., & Ta-Shma, A. (1999). Flow control: A new approach for anonymity control in electronic cash systems. In *International conference on financial cryptography* (pp. 46–61).

Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering*, *63*(5), 603–613.

Silfversten, E., Favaro, M., Slapakova, L., Ishikawa, S., Liu, J., & Salas, A. (2020). *Exploring the use of Zcash cryptocurrency for illicit or criminal purposes.* RAND.

Sunyaev, A., Kannengießer, N., Beck, R., Treiblmaier, H., Lacity, M., Kranz, J., ... Luckow, A. (2021). Token economy. *Business & Information Systems Engineering*, 1–22.

Tefft, S. K. (1980). *Secrecy a cross-cultural perspective.* Human Sciences Press.

The Guardian. (2021). *New Zealand's central bank says its systems have been hacked.* Retrieved 2022-01-10, from https://www.theguardian.com/world/2021/jan/11/new-zealands-central-bank-says-its-systems-have-been-hacked

Tinn, K., & Dubach, C. (2021). *Central bank digital currency with asymmetric privacy.* Retrieved 2022-01-10, from https://www.mcgill.ca/engineering/files/engineering/central_bank_digital_currency_with_asymmetric_privacy_mcgill_tinn_dubach.pdf

Tramèr, F., Boneh, D., & Paterson, K. (2020). Remote side-channel attacks on anonymous transactions. In *29th {USENIX} security symposium* (pp. 2739–2756).

United Nations. (1948). *Universal declaration of human rights.* Retrieved 2022-01-10, from https://www.un.org/en/udhrbook/pdf/udhr_booklet_en_web.pdf

Veneris, A., Park, A., Long, F., & Puri, P. (2021). *Central bank digital loonie: Canadian cash for a new global economy.* Retrieved 2022-01-10, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3770024

Wüst, K., Kostiainen, K., & Capkun, S. (2021). *Platypus: A central bank digital currency with unlinkable transactions and privacy preserving regulation.* Retrieved 2022-01-10, from https://eprint.iacr.org/2021/1443

Zhang, R., Xue, R., & Liu, L. (2019). Security and privacy on blockchain. *ACM Computing Surveys*, *52*(3), 1–34.

Zwick, D., & Dholakia, N. (2004). Whose identity is it anyway? Consumer representation in the age of database marketing. *Journal of Macromarketing*, *24*(1), 31–43.

The Importance of Where Central Bank Digital Currencies Are Custodied

The importance of where Central Bank Digital Currencies are custodied: exploring the need for a Universal Access Device

Filippo Zatti[a] and Rosa Giovanna Barresi[b,c]

[a]0000-0002-1528-3045, Department of Economics and Management, BABEL, University of Florence; [b]0000-0002-2845-2194
[b] Barresi Law Firm; [c] BABEL

ABSTRACT

Issuance of digital currency is not new, but DLT enables to consider monetary policy and financial inclusion in a new way. CBDC is so disruptive it can change money. As digital asset trading evolved into a market, developers released digital wallets for their storage. While these attempts meet the users' needs, their technical level is open to improvements. The development of a secure wallet infrastructure can solve critical problems, like ensuring equal access to banking service, also offering a novel approach to digital identity. Whenever a financial transaction is carried out, we must face problems with fair competition, ownership, and management of our personal data. Infrastructure is needed to protect and support our digital rights. This technology is ready to develop a secure UAD, a single key tool for protecting and representing us and the organisation to keep it in a network. UAD would enable the full benefit of DLT and restore the separation of payment institutions and investment banks at the core of political agendas, store SSI related information, facilitate micro payments, and tackle using money for illicit purposes more effectively. UAD is a citizen-centric device, a symbol of the new digital currency and a tool for social inclusion.

1. The idea behind CBDC

The creation of a crypto space can help overcome established economic and political patterns that have, in many cases, proved inadequate to ensure economic and financial stability and social progress.

However, in order for the process to be implemented to achieve expected results, it is necessary to be immediately aware that emerging perspectives require a thoughtful and functional transformation of the legal institutions on which the new reality lies. Before blockchain technology arose, digital wallets were actually a phenomenon, but in the previous context their function was exclusively related to the need for digital representation of physical goods - including currency and payment instruments - to facilitate online transactions. The ability to manage goods that do not have a physical twin - as certain cryptoassets - creates the need for accessible, functional, secure forms of custody. We are therefore confronted with a new need, rather than increasing the efficiency of payment systems linked to Web 2.0. In Web 3.0, digital wallets change their function with the correspondent effect for technical, economic and legal aspects (Turi 2020).

CONTACT Filippo Zatti. Email: filippo.zatti@unifi.it (par. 1, 2 and 6).
Rosa Giovanna Barresi. Email: rgbarresi@gmail.com (par. 3, 4 and 5).

One first situation in which a custody system with characteristics such as digital wallets is needed is the management of CBDC.

The project to issue digital currencies with distributed ledger technology is no new. About seven years ago, when Bitcoin emerged and FinTech began to rise, the Bank of Canada launched the Jasper project. It aims to examine how the new technology can improve the efficiency of the payment system, even if it is aware that such opportunities have difficulties fulfilling its mandate. And, not only in terms of price stability, but also in terms of maintaining trust in the payment system. The study also includes CBDC issuance accessories such as the so-called "Universal Access Device" (hereinafter UAD). It is questionable whether the central bank should implement such an instrument, or whether the involvement of some private entities is possible or desirable. The second critical aspect concerns the usability of instruments also, by those who escape financial inclusion, which is not challenged by the traditional banking system by choice, personal, social, or economic conditions. For this reason, UAD solutions should be extensible (universality), cost-effective (accessibility) and differ from those currently used for cryptoassets (smartphones and similar dedicated devices). Finally, universality requires that this type of tool be perceived or managed from a legal view as a necessary and identifiable tool of the person and jointly able to ensure complete compliance with the rules on the protection of personal data. This is central to the development of UAD solutions. These types of devices should ensure a high level of privacy and a low risk of illegal use of CBDC (AML and CFT compliance). One solution could be multi-party computations and zero-knowledge proof. However, other proposals are coming.

The difference between CBDC and stablecoin is not only in terms of the modalities of issue and custody, but also in terms of the role of money in the economic and financial system, in terms of disintermediation, in terms of the modalities of data management related to monetary transactions. Legal questions arise that concern CBDC closely before the UAD architecture and how it is applied in individual money systems (Mancini-Griffoli et al. 2019). But the definition of UAD is already becoming strategic in the light of the profound economic and legal changes that are taking place, and that is the prospective geo-economic level. The transformation of the monetary system, the erosion of the fiat currencies' monopoly, and the evolution of technology are three factors that require the attention of the public authorities to converge on the UAD. Probably, in the new world, the currency will no longer represent sovereignty. Instead, it will be the monopoly of the infrastructures that hold it. Therefore, the UAD will no longer be only digital wallets, and their function may be much more relevant than the accessory one reserved to them today.

2. UAD and CBDC: union of necessity or convenience?

The idea of a form of digital money is not recent. Even in the 'prehistory' of cryptocurrencies, the traceability of payments was raised (Chaum 1983). DigiCash (Ecash), E-gold, Liberty Reserve are the significant examples of the attempts made before Bitcoin and cryptos focused on privacy like Monero, Zcash, and Mimblewimble.

Outside the legal framework, cash only partially replaces a payment and custody method from a technical view with cryptocurrency. Contrary to crypto, cash can be used "offline", fully fungible, not fully evading tax control, and is resistant to accumulation. However, both physical and digital currencies as cryptocurrencies ensure a potential transaction will be successful regardless of custody or third-party relationship that could block or delay it or require verification before it can take place (Goodell et al. 2020). It would lead to a move to DLT-based solutions instead of centralized solutions. Here, a preference for a token-based system rather than payment accounts might be natural (Auer and Boehme 2020). Regardless of the architecture of the DLT system, portability and custody of digital currency in the various forms they can take today (CBDC, stablecoin, cryptocurrency) is emerging in all the economic and legal relevance (Adrian and Mancini-Griffoli 2019). In a token-based payment system, in which tokens perform a different economic function than cash, along with an uncertain legal definition, custody is no longer marginal (Zatti 2019).

3

It encompasses aspects that relate not only to business activities, but also to the dimension of sovereignty, understood in its broadest sociological conception, rather than resorting to a purely political one. CBDC issuing projects can be interpreted according to a 'conservative' ideological vision if it is opposed to the idea from which Bitcoin originates and 'first- style' Libra is promoted. The neutral-like solutions risk being more political than those that are so since origin. However, the process underway, it must be recognized, is valuable because it brings back the attention to the debate started after the imminence of the 2008-2009 global financial crisis and dormant until now. It is no coincidence that currency occupies a central place in economic law. It brings with it a crucial economic function of modern economies, namely promoting exchange. From a legal view, money is functionally defined in every legal system, namely in terms of guaranteeing trade and the stability of its value. First, the international dimension of exchange, and then globalization, not to mention the monetary agreements that followed each other, have given currency a meaning that goes far beyond economic function and civil law issues. The creation of a DLT-based system for digital coins as tokens is in line with the evolution of the world's geo-economic framework. In the current context, the digital currency goes beyond sovereignty because it tends to take effect other than the jurisdiction in which it is issued and regulated. Dealing with the issue from an economic-monetary view is based on understanding what is happening, even though the pandemic has further accentuated it: an increase in the public debt of nation-states; an international financial system once again facing a deep economic recession; and the US — People's Republic of China dyarchy. It is no coincidence that CBDC's oldest and most advanced project (named DCEP in China) is the People's Central Bank to date, aiming to increase the weight of renminbi in the international payment system at the expense of the US dollar (Sender 2020). And it's not just accidental in recent times, in the US, there's been an acceleration on various fronts of the digital currency (Cheng et al. 2021). Significant, even if it is a bill first introduced to Congress in April 2020, as a measure in the face of the COVID-19 pandemic, is the 'Automatic Boost to Communities Act' (aka ABC). Here, 'digital dollars' and the 'digital dollar account wallets' are introduced as well as the 'digital dollar cash wallets'.

The proposed solution in the bill appears only partially the most commonly used for digital currency as direct credit to the central bank and private intermediaries. The proposed model creates an indirect claim against the US Treasury, which issues digital dollars that record them as liabilities in each bank's Federal Reserve accounts. But the feature of this model is the creation of 'state wallets' authorized and regulated by the US Treasury and operated by the Federal Reserve. The models proposed to date focus on the type of infrastructure underlying digital money and on aspects of monetary policy management, as a stimulating study by the BIS compared the approaches adopted by the People's Bank of China, the Swedish Central Bank and the Bank of Canada (Auer et al. 2020).

So far, the wallet's management methods, not only operational but also legal, do not sound perceived as primary. But wallets, as has been said, are a crucial and relevant aspect of the digital currency.

As far as the Chinese project is concerned, different types of wallets in distinctive identification forms are considered to have varied types of anonymity of the user and access modalities in accordance with KYC restrictions. This approach would enable users of digital assets to remain anonymous in peer exchange. However, it allows the central bank to track data on prudential regulation, money laundering and other criminal offences. In the Swedish pilot project, there is use of payment accounts with intermediaries. However, in the case of modest money, an alternative option remains open: prepaid cards that can use tokens (Bindseil 2020). But both Chinese and Swedish approaches do not tackle cap management and how to identify these thresholds. In addition, anonymity with the central bank is guaranteed in both cases on the traditional regulatory model, where intermediaries are entrusted with the responsibility of verifying KYC and monitoring the correct use of CBDC.

Finally, as in the introduction to this paper, Canada's central bank, taking both models into account, is also oriented towards a two-tier solution that also meets the needs of the unbanked through a specific cost-effective solution that allows CBDCs to be used without a smartphone. It is the proposal for a UAD, which is specifically discussed here in this paper. A

solution of this type would allow a series of advantages for the end-user of no small importance.

These benefits would be both in terms of reducing payment burdens and in terms of the possibility of excluding citizens from the applications of the current digital revolution. Concurrently, it would be a critical issue for the banking system and, more generally, for financial intermediaries. What and how many advantages and disadvantages depend on the type of architecture chosen in the respective monetary areas. However, it is interesting to note both the Italian and German banking associations are pushing for a programmable digital Euro, supported by the ECB, in stable coin or CBDC. The reasons for the introduction of the programmable digital euro are cross-border payments, particularly outside SEPA and micro-payments, and payment automation.

Consider the recent study of the BIS: the architecture of CBDC varies depending on the type of economy served (Auer et al. 2020). Instead, the regulation of wallets is equally likely to be related to the economy's regulatory model. Such a correlation may seem more understandable by resorting to the Libra/Diem project's transformation since its inception. Concerns emerged about financial system stability following the risk of a private monopoly of corporate money (corporate - sovereignty). Together, we would have the potential to create as many wallets as Facebook accounts. Data conveyed by the social network would have been further increased by the information resulting from Libra/Diem's use. Digitization of currency and economy in general requires politicians, regulators, and academics to measure the compatibility of a global private monetary system with a limited dimensional expression of political sovereignty (Pistor 2019). In the future currency, a battle is played out that goes well beyond the purely technical aspects. Digital currency is technically aligned to any other asset (to be understood in a broad and physical sense) that can be represented through digital assets. Therefore, in that economic and monetary system, the provision of token custody and deposit services is not considered as if it is for any other dematerialized asset, financial or not. Thus, wallet passes from being considered an accessory service of payment, as it happens for the instruments in which the legal tender currency and the bank currency can be deposited or transformed, to a must custody device.

In the token-based system, wallets play a key role because through them direct currency exchanges can take place, avoiding the passage by intermediaries. It makes sense no longer to think of a regulation model like the one currently in place for dematerialized currencies, which means that the potential changes and crucial functions of wallets within the system will be taken into account.

The system proposed in the ABC Act albeit initially in the introduction of universal basic income, is not only emblematic for aspects related to the distribution of the authorization and control powers of wallets and the private configuration of the instrument, but also for aspects inherent in the operation of the instruments introduced. The bill provides for operations that can be consolidated at the end of the current pandemic. The Digital Dollar Accounts Wallet, "FedAccounts", will be available alongside the related financial services (debit cards, access to online accounts, automatic debit cards, mobile banking, ATM), which are physically managed by postal services. At the latest, on January 1, 2022, the Treasury will offer a second opportunity to those entitled to so-called BOOST payments: to receive payments in digital dollar wallets or via temporary cards that allow them to activate account wallets. Commissions do not affect digital dollar wallets, as also minimum or maximum balances, and may not be closed or tied due to profitability. Postal services will provide access to digital dollar wallets and related services to US citizens, residents of the US, and businesses.

All should be in the disadvantaged areas of the country, either due to income or in a disaster zone or in metropolitan areas with medium income, challenged or poorly served, as the Federal Financial Institutions Examination Council says. The ABC Act also provides adequate cover for losses resulting from fraud or security breaches. Moreover, digital dollar wallets must comply with the Bank Secrecy Act and the Privacy Act, which entailed the applicable federal tax and Internal Revenue Code of 1986.

Once the stimulus plan is over, the Treasury will develop and administer a digital dollar cash wallet system, named eCash. It will be available for those who request it to store, send, and receive digital coins or other digital currency issued by the legal tender Treasury. Furthermore, integration and interoperability between the digital dollar account wallets managed by the Federal Reserve and digital dollar cash wallets shall be promoted. The latter should also be hosted on 'easy-to-find, inexpensive digital devices, including smartphones, and managed through software with any necessary open-source licensed software or hardware'. With regard, then, to issues concerning privacy protection, the Secretary will set up a Digital Financial Privacy Board with the task of supervising, monitoring, and reporting on the architecture and implementation of the digital dollar cash wallet system. Then, once it is fully operational, it will ensure the continuous supervision of the system's management itself. This system must be designed in such a way as to reproduce the characteristics of transactions with physical currency.

The relationship between cryptocurrencies and wallets differs from those between CBDCs and digital currency wallets in the UAD. It reproduces legal and economic criticality that underlie the legal nature of cryptocurrency and blockchains' governance. As long as private monetary systems cause satellites on the "monetary planets", if we are allowed this bold metaphor, the strengthened schemes on which modern financial systems move are seldom involved. Moving to a fiat digital currency is complex in both technical and emerging economic matters and inevitably in legal reforms. Perhaps we face further changes in the concept of sovereignty. Homogenization of goods and rights due to the transformation of computer data necessarily shifts the spotlight on the device allowing its storage and use.

The aim is to launch a debate on the critical issues that are simply highlighted here, and to find a solution that can be reconciled with all the problems raised. What we will try to make emerge is that dematerialization is something different from digitization. Digitization creates an exchange system based not on reality, but on codification. When we accept this idea, even before the type of architecture of CBDC, the regulation of the instrument with which rights and assets are managed takes new importance.

3. Exploring UAD about technical design

Several trends are driving the development of a UAD: while physical movement of cash is becoming expensive, paying with plastic is getting fashionable.

More seriously, as natural calamities (floodings, earthquakes, or pandemics) (Cœuré et al. 2020) can impair cash distribution and credit card payments, people need a robust and reliable way of exchanging money.

A UAD-based infrastructure should handle low-cost transactions through NFC (Near-Field Communication), even when banking services are not available (Bank of Japan, 2020). The same infrastructure should be used for delivering money to the needy in a secure way. A pan-European UAD infrastructure should rely on two components: a standard hardware platform and a secure management network (Bank of Japan 2020). Hardware supports several applications running simultaneously, almost as a smartphone, but in a highly secure environment. Whereas to available e-wallets, a UAD could handle together with a CBDC wallet, applications for NFC payments, and also store SSI artefacts.

In contrast to common smartphones, the network will comparatively constantly monitor the UAD hardware and will operate only certified applications. From a functional viewpoint, a UAD is a state-of-the-art smartphone, less the (costly) multimedia components, and plus a 'connected secure element'. This highly secure component marks a major step in mobile application security and is presently installed only on top-level smartphones. It hosts on the same chip both network identification (the information usually stored inside a SIM card) and secure storage and computing functionality. Any attempt to tamper with that chip would result in disabling the smartphone. As the secure storage facility can be used for storing digital certificates and SSI-related information, the whole chip is a secure way of linking network identification data with user personal information. Thanks to its functional specialization, the UAD will vastly improve on smartphone standards in terms of unit cost, weight, autonomy,

and usability. Through careful design, the UAD shall guarantee equal access to state-issued money, even to people who do not own a smartphone or may find it difficult to use (the elderly, the unbanked, the digitally impaired). With its limited weight and its square form factor, it will allow simple single-hand operation and thumbprint identification. Also, it will support different User Interfaces, according to national sensibilities: following the principles of User Interface (Ux) design, those 'skins' will be tested and refined by national focus groups.

Many of these features will be shared among all citizens: for example, an improved display designed for visually impaired people will benefit anyone who uses the device with insufficient illumination (Miedema et al. 2020). From a system engineering standpoint, the role of the secure management network is as relevant as the hardware architecture in ensuring reliable and secure operation of the device. It will track software releases from certified producers and will ensure automated and timely updates of the software. Tampered devices will be disabled automatically and, as in the case of mobile telephones, the judiciary power (through an appropriate command chain) may request to monitor the activity of a specific device or to investigate its previous use.

4. Economic aspects affecting use of UAD

4.1. The struggle for digital supremacy

The conflict between the US and China has been part of the international scene for more than five years. While their battle for digital supremacy is only one aspect of this larger conflict, its outcomes are having a significant impact on the EU. The European Commission (EC) has striven to safeguard the rights of its citizens in various ways:

1. International Treaties: The EU-US Privacy Shield was signed with the Obama administration and established a legal framework for processing sensitive data outside the EU.
2. Investigations: Internet companies like Google and Apple have been investigated and fined for anti-competitive behavior levied.
3. Technological innovation: promoting the development of start-ups by financing and other incentives.

These measures were approved by European partners and helped enforce laws and regulations approved by the European Parliament. Moreover, the CJEU may invalidate international agreements after judges have established they violate the TFEU. This happened on July 16, 2020, when the CJEU invalidated the decision of the EC on the adequacy of the protection provided by the EU-US Data Protection Shield. This international agreement was through a voluntary self-certification scheme, whereby US companies committed to complying with European data protection requirements, to store and process European residents' data in the US (Floridi 2020).

About the investigations on Internet companies, it is extremely difficult to restore a free market once they have reached a near-monopoly position. Also, there are few chances of reaching a verdict when dealing with corporations that can bear the cost of a multi-year dispute. For example, in August 2016 the EC established Ireland granted illegal State aid to Apple through selective tax breaks for the years 2003-2014, giving Apple 13 billion Euro in a tax advantage (European Commission 2020). But on July 15, 2020, the General Court of the European Union of the First Instance annulled that decision.

Conventional actions for encouraging a free market can prove inadequate to contain the expansion of Internet companies. WhatsApp Pay obtained the authorization to launch a mobile service payment in Brazil by exploiting a flaw in a start-up friendly legislation. As there are more than 120 mln Brazilian users of WhatsApp, the new payment application would have reached a monopoly position in the payment services industry. Urged by commercial banks on June 24, 2020, the Brazilian central bank suspended the authorization, only one week after the launch of the service (Finextra 2020).

In the end, European digital sovereignty should be safeguarded by more incisive actions, as the construction and protection of digital infrastructure. Digital infrastructure standards as safeguards of Europeandigital sovereignty

7

4.2. Digital infrastructure standards as safeguards of European digital sovereignty

It has been widely known that proprietary standards can hinder the development of the market and reduce the level of protection of privacy. We should develop digital infrastructure standards before proprietary approaches have taken control of our life. Market regulators can carry out their activities only after a fair playing field has been established, by building a European digital infrastructure, shared by all service providers.

Seeking again back to history, the European GSM phone standard was crucial in developing a global cellular phone market. It started in 1985 as a limited agreement on cellular technology between France, Germany, and Italy, and gradually evolved into a global standard. When European Telecommunications Standardization Institute (ETSI) was set up in 1987, membership was still limited to European organizations. The opening of ETSI to non-European members sent a clear message that GSM technology was to be shared and everyone could have their say in its specifications. After that, foreign phone manufacturers cooperated in developing the standard and a global market developed. As a further result, on June 15, 2017, roaming charges were completely eliminated in the EU.

As an example of the problems caused by a proprietary standard, the EC launched an investigation against Apple Pay (European Commission 2020a). According to the Commission, "Apple is 'abusing' its control of the wallet through a policy that blocks third-party payment services to access its NFC hardware, which enables contactless payments" (Toplin 2020). The commission-supplied data showed Apple Pay's growing market penetration in mobile payment services, but Apple is opposing that claim, pointing to the COVID-19 pandemic as the cause of the switch in customers' preferences.

As cryptography techniques are evolving, it appears that the FinTech sector will have to rely on some UAD, to store personal information in a secure and reliable way. The definition of a European Standard for such a device would enable fair competition between companies and will safeguard the digital identities of European citizens.

4.3. From the individual right to identity to the concept of digital sovereignty

The EU Charter of Fundamental Rights states everyone has the right to protect their personal data (Art. 8). An analogy can further explain the concept of personal data as a negotiable commodity. In almost any jurisdiction the trading of human organs is forbidden (both the selling and the buying of them), since this activity causes an irreparable consequence for the person involved. Following this analogy, as the loss of personal identity is irreparable, it should not be traded too. Once sensitive data (for example, one's socio-economic profile as a consumer) has been released, the activity of finding and deleting all their copies is virtually impossible. The need to proactively protect this personal right is the basis of digital sovereignty. Water and power have been granted public service status, as they are necessary to provide for the fundamental rights of the person. In the same way, the infrastructures needed to proactively protect the digital identity should be recognized as a public service.

According to Art. 16 of the TFEU "The European Parliament and the Council... lay down the rules relating to the protection of individuals with regard to the processing of personal data by Union institutions, bodies, offices, and agencies". In this respect, the UAD and comparable infrastructures should be recognized as tools for the protection of the digital identity.

5. Crucial legal issues to face before introducing UAD

5.1. Basic compliance analysis for UAD-based CBDC

Many arguments support the case for issuing a CBDC: lowering transaction costs and improving the efficiency of present payment systems while pursuing financial inclusion (Auer et al. 2020).

Until now, the ECB has not disclosed any plans to issue a CBDC in the near future. On the other hand, the CBDC issue has been addressed in many publications of the ECB (ECB 2020a; ECB 2020b; ECB 2020c; Panetta 2020; ECB 2021a), and for the sake of our discussion, the role

of the UAD will be described in a European context. Among the many design options for implementing a CBDC, the most popular approach is to complement physical cash in supporting citizens in their day-to-day expenses: a retail oriented, legal tender CBDC, distributed by bank intermediaries (Auer and Boehme 2020; Pfister 2020; ECB 2021b). A wholesale oriented CBDC, on the contrary, would target the more specific issue of inter-bank settlements. Following this (hypothetical) indirect model of issuance, the central bank will issue CBDC money to commercial banks and Payment Interface Providers (PIPs) and only these two organizations will supply that money to citizens.

The introduction of UAD can illustrate the legal issues that underlie the operation of such a retail CBDC. From a practical standpoint, the UAD is a secure, physical object, carrying the identification of its owner (it is so flexible that specific users may be provided with ad-hoc versions of the software). All the events the UAD may encounter during its lifecycle must be managed in accordance with the basic AML, CFT, and GDPR. CBDC design should find a new balance between allowing data portability, protecting privacy and mitigating the risks of money laundering and illicit financing (BIS 2020; Jackson and Tayhar 2020).

Also, that solution should be robust enough to withstand fraud and cyberattacks and to accommodate the usual judicial proceedings like a seizure, bankruptcy, auditing and monitoring of accounts, etc. Replacement of lost, stolen or malfunctioning units must also be supported. Six roles can be identified to support the functioning of a UAD-based CBDC:

1. The ECB, as the issuer of CBDCs.
2. The judiciary power, as the issuer of judicial proceedings.
3. An independent authority as a link between the ECB and the judiciary power tasked with monitoring and investigation.
4. The executive power (usually the Ministry of Interior) as the custodian of the citizens' identities and issuer of ID cards.
5. Commercial banks and Payment Interface Providers (PIPs) as the executors of the customer identification (KYC) process.
6. The UAD Management Network, executing orders received from the previous five.

As already described, a UAD is a safe way to store a network identification (an irreplaceable SIM card) in the same place and some User Personal Data (a set of self-sovereign identification artefacts guaranteed by the Ministry of the Interior).

The executive power, judiciary power, the PIPs, and the commercial banks identify the users by their Personal Data (in the same way they are doing now).

The suggestion to delegate the KYC process to commercial banks and PIPs is in accordance with a proposal from the Bank of England (Bank of England 2020; Allen et al. 2020).

The only organization that will manage both network identification and personal data will be an independent authority: a renewed European-level organization will be tasked with the monitoring of the network and AML/CFT investigations. This in accordance with a study published by the ECB itself: 'That authority checks the identities of users involved in large-value transactions and prevents CBDC from being transferred to embargoed users' (ECB 2019). At the lower level, the UAD Management Network has only access to the network identification data stored on the SIM card.

5.2. The management of the UAD as a conflict of competence issue

Should a UAD be adopted for supporting a retail CBDC, in which European institutions should be responsible for its operation? This thought experiment may help clarify the more general question of whether the European legal framework specifies who should manage any sensitive infrastructure.

According to the Art. 127, par. 2 of TFEU, the ECB, through the Eurosystem, has the responsibility 'to promote the smooth operation of payment systems'. So, it could be assumed the legislation would grant to the ECB/Eurosystem the authority to manage the UAD as a payment system. According to an ECB study, 'the central bank is the only entity allowed to issue CBDC units and remove them from circulation' (ECB 2019). However, in the analysis of conflicting competences, it is important to identify the overriding goal among all those

assigned to the conflicting institutions. From this view, the UAD lifecycle goes way beyond the CBDC circulation process, touching on personal data. This involves again Art. 16 TFEU. In this respect, the guidelines for the deployment of the UAD should be probably set at a higher level. This may open the way for cooperating with many European institutions also with the private sector.

As a first example, the European Blockchain Partnership (EBP) (European Commission, 2018) was established on April 10, 2018. Later, other nations (including Italy) have joined the initiative, aimed at creating EBSI (European Blockchain Services Infrastructure), a European digital infrastructure for a secure and reliable data transfer between European-based institutions and organizations. ESSIF, the European Self Sovereign Identity Framework, is a special use case for EBSI.

ENISA, the EU Agency for Cybersecurity, established in 2004, is a centre of expertise for cybersecurity in Europe. While its focus is to prevent, detect, and respond to information security problems, it also helps in drafting EU policy and law on data protection and network security.

Europol (European Union Agency for Law Enforcement Cooperation) was founded in 1998 to deal with criminal intelligence and combat serious international organised crime and terrorism. In the US, the mission of safeguarding the financial and critical infrastructure is mandated to the Secret Service (now a Federal Agency under the Department of Homeland Security). As Europol acts as a worldwide contact point, the two agencies cooperate on counterfeiting and financial crimes.

6. From the Bancor to the 'Universal Wallet': two different monetary systems, the same ontological concept of sovereignty

Historically, reforms of the banking system followed severe political or economic upheavals.

The mind quickly turns to Bretton Woods and the idea of a currency as a unit of account that would allow a trade to be tracked by considering assets and liabilities of balances of payments. It is not by chance that the idea of a Bancor came back into the debate among economists after the 2008-2009 global financial crisis. However, in this context, the international financial system has shown resilience in dealing with a severe crisis. The COVID-19 pandemic is again testing the Bretton Woods system. This time, resilience may not have the same result, as non-conventional monetary policy instruments already existing must be confronted with a significant increase in national public debt. Time has probably come to imagine a new monetary system structure. The emerging monetary system still has a predominance of the dollar, supported by the so-called C6, and then, cascading, bilateral swaps, regional pools, IMF, national currency and credit. Looking ahead, however, it could be a transition from the dollar to another currency, like the renminbi (Jenkins et al. 2014).

Another hypothesis could be to benefit from the impact of the introduction of DLT in the payment system, considering various types of currencies with different functions, public and private. Imagine in this context a super-sovereign currency that could be the new foundation of the global monetary system. Hence, we will find various strategies for crisis management and global control, disengaging sovereign currencies from international trade, and balance of payment consequences. A complex political address, but this DLT application could facilitate its transparency and decentralization functions.

Continuous acceleration of projects that lead to the digital currency, in particular CBDC, in a short time creates not only the conditions for more efficient currency circulation, but also the definition of a monetary system with the features described above. This efficiency can be exported to international transactions that access operators can use to store and use applicable digital currencies. In this context, the 'wallet' is more relevant than its contents. The goal of regulation will move from currency to wallet, to which any type of digital currency and other information of this type can necessarily flow. It could go beyond regulation and imagine situations where monetary policy directly affected the wallet.

There are many scenarios that can be opened up and merit a thorough investigation by politicians, regulators, academics, and industry. The aim of this paper is to start the debate and to show both a direction and a specific solution: UADs. UAD would enable the full benefits of DLT, solve financial inclusion problems, split payment institutions and investment banks (putting them at the centre of political agendas), encourage micro-payments, and limit using money for illegal purposes. All this obviously requires careful work to identify UAD management problems and feasibility analysis in terms of cost/profit ratio. Likewise, it is essential to carefully review the legal issues we have only mentioned here. With a particular focus on the EU, the problems would cover both the regulatory framework and the link to existing GDPR, AML, CFT, PSD2 (Hossein Nabilou 2019).

However, it should be noted there are technical, design, legal and economic resources in the EU that are capable of developing and implementing this project. But there could be other positive effects if interest in this type of problem converges outside the EU too.

So far, we have tackled the crisis 'making money'. Experience has taught us this lesson: liquidity is effective, but not decisive because it evokes sovereignty. Separating currencies from sovereignty ('breaking money') and transferring them to the infrastructure by which liquidity is stored and used could be an effective and decisive solution to the governance of the world's monetary system after the pandemic.

It is worth checking.

References

Adrian, T., & Mancini-Griffoli, T. (2019). *The rise of digital money*. International Monetary Fund.

Allen, S., Čapkun, S., Eyal, I., Fanti, G., Ford, B., Grimmelmann, J., Juels, A., Kostiainen, K., Meiklejohn, S., Miller, A., Prasad, E., Wüst, K., & Zhang, F. (2020). Design Choices for Central Bank Digital Currency: Policy and Technical Considerations. *NBER WORKING PAPER SERIES*. https://doi.org/10.3386/w27634

Auer, R., & Boehme, R. (2020). The technology of retail central bank digital currency. *BIS Quarterly Review*.

Auer, R., Cornelli, G., & Frost, J. (2020). Rise of the central bank digital currencies: Drivers, approaches and technologies. *BIS Working Papers*, *No 880*.

Bank of England. (2020). Central Bank Digital Currency: Opportunities, challenges and design. *Discussion Paper*.

Bank of Japan. (2020). *The Bank of Japan's Approach to Central Bank Digital Currency*. https://www.boj.or.jp/en/announcements/release_2020/data/rel201009e1.pdf

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

BIS. (2020). Central banks and payments in the digital era. *Annual Economic Report*.

Chaum, D. (1983). Blind Signatures for Untraceable Payments. In D. Chaum, R. L. Rivest, & A. T. Sherman (Eds.), *Advances in Cryptology* (pp. 199–203). Springer US. https://doi.org/10.1007/978-1-4757-0602-4_18

Cheng, J., N Lawson, A., & Wong, P. (2021, February 24). *Preconditions for a general-purpose central bank digital currency* [Institutional]. Board of Governors of the Reserve Federal System. https://www.federalreserve.gov/econres/notes/feds-notes/preconditions-for-a-general-purpose-central-bank-digital-currency-20210224.htm

11

Cœuré, B., Cunliffe, J., Bank für Internationalen Zahlungsausgleich, Bank of Canada, Europäische Zentralbank, Nihon Ginkō, Sveriges Riksbank, Schweizerische Nationalbank, Bank of England, Federal Reserve System, & Board of Governors. (2020). *Central bank digital currencies: Foundational principles and core features report no. 1 in a series of collaborations from a group of central banks*. https://www.bis.org/publ/othp33.pdf

ECB. (2019). Exploring anonymity in central bank digital currencies. *IN FOCUS*, *Issue no 4*.

ECB. (2020a). *Unleashing the euro's untapped potential at global level*. https://www.ecb.europa.eu/press/key/date/2020/html/ecb.sp200707~3eebd4e721.en.html

ECB. (2020b). *Digital euro trade mark registration request to EUIPO*. https://euipo.europa.eu/eSearch/#details/trademarks/018311625

ECB. (2020c). *Report on a digital euro*. https://www.ecb.europa.eu/pub/pdf/other/Report_on_a_digital_euro~4d7268b458.en.pdf

ECB. (2021a). *ECB digital euro consultation ends with record level of public feedback*. https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210113~ec9929f446.en.html

ECB. (2021b, March 25). *Digital central bank money for Europeans – getting ready for the future*. https://www.ecb.europa.eu/press/blog/date/2021/html/ecb.blog210325~e22188c522.en.html

European Commission. (2018, April 10). *European countries join Blockchain Partnership* [Text]. Shaping Europe's Digital Future - European Commission. https://ec.europa.eu/digital-single-market/en/news/european-countries-join-blockchain-partnership

European Commission. (2020a, June 16). *Antitrust: Commission opens investigation into Apple practice regarding Apple Pay* [Text]. European Commission - European Commission. https://ec.europa.eu/commission/presscorner/detail/en/ip_20_1075

European Commission. (2020b, July 15). *Statement by Executive Vice-President Margrethe Vestager following today's Court judgment on the Apple tax State aid case in Ireland* [Text]. European Commission - European Commission. https://ec.europa.eu/commission/presscorner/detail/en/statement_20_1356

Finextra. (2020, June 24). *Brazil Central Bank ices Facebook's WhatsApp payment service*. Finextra Research. https://www.finextra.com/newsarticle/36083/brazil-central-bank-ices-facebooks-whatsapp-payment-service

Floridi, L. (2020). The Fight for Digital Sovereignty: What It Is, and Why It Matters, Especially for the EU. *Philosophy & Technology*, *33*(3), 369–378. https://doi.org/10.1007/s13347-020-00423-6

Goodell, G., Al-Nakib, H. D., & Tasca, P. (2020). Digital Currency and Economic Crises: Helping States Respond. *ArXiv:2006.03023 [Cs, Econ, q-Fin]*. http://arxiv.org/abs/2006.03023

Hossein Nabilou. (2019). Testing the waters of the Rubicon: The European Central Bank and central bank digital currencies. *Journal of Banking Regulation*. https://doi.org/10.1057/s41261-019-00112-1

Jackson, H. E., & Tayhar, M. E. (2020). *Fintech Law—The case studies*. https://projects.iq.harvard.edu/files/fintechlaw/files/fintech_law_the_case_studies.pdf

Jenkins, P., Bernes, T. A., Mehrling, P., & Neilson, D. H. (2014). China's Engagement with an Evolving International Monetary System: A Payments Perspective. *Cigi and Institute for New Economic Thinking*, *Special Report*.

Mancini-Griffoli, T., Martinez Peria, M. S., Agur, I., Ari, A., Kiff, J., Popescu, A., & Rochon, C. (2019). Casting light on Central Bank Digital Currencies. In C. Brummer (Ed.), *Cryptoassets: Legal, regulatory, and monetary perspectives*. Oxford University Press.

Miedema, J., Minwalla, C., Warren, M., & Shah, D. (2020, June). *Designing a CBDC for universal access* [Staff Analytical Note 2020-10 (English)]. Bank of Canada. https://www.bankofcanada.ca/2020/06/staff-analytical-note-2020-10/

Panetta, F. (2020, October 22). *On the edge of a new frontier: European payments in the digital age*. https://www.ecb.europa.eu/press/key/date/2020/html/ecb.sp201022~d66111be97.en.html

Pfister, C. (2020). *Central Bank Digital Currency*. https://publications.banque-france.fr/en/central-bank-digital-currency

Pistor, K. (2019, August 5). *The Right Response to the Libra Threat | by Katharina Pistor & Co-Pierre Georg*. Project Syndicate. https://www.project-syndicate.org/commentary/regulating-private-money-facebook-libra-by-katharina-pistor-and-co-pierre-georg-2019-08

Sender, H. (2020, August 3). China's new digital currency takes aim at Alibaba and Tencent. *Financial Times*.

Toplin, J. (2020, June 17). *Apple Pay is facing an antitrust investigation by The European Commission*. Business Insider. https://www.businessinsider.com/apple-pay-faces-european-antitrust-probe-2020-6

Turi, A. N. (2020). Currency Under the Web 3.0 Economy. In A. N. Turi (Ed.), *Technologies for Modern Digital Entrepreneurship: Understanding Emerging Tech at the Cutting-Edge of the Web 3.0 Economy* (pp. 155–186). Apress. https://doi.org/10.1007/978-1-4842-6005-0_5

Zatti, F. (2019). *Entangled in Cryptoassets' Legal Nature and Governance: Searching for Clear Boundaries or Working for their Removal?* (SSRN Scholarly Paper ID 3335793). Social Science Research Network. https://doi.org/10.2139/ssrn.3335793

Platypus: A CBDC with Unlinkable Transactions and Privacy Preserving Regulation

Platypus: A Central Bank Digital Currency with Unlinkable Transactions and Privacy-Preserving Regulation

Karl Wüst*
CISPA Helmholtz Center
for Information Security
Germany

Kari Kostiainen
Department of Computer Science
ETH Zurich
Switzerland

Noah Delius
Department of Computer Science
ETH Zurich
Switzerland

Srdjan Capkun
Department of Computer Science
ETH Zurich
Switzerland

ABSTRACT

Due to the popularity of blockchain-based cryptocurrencies, the increasing digitalization of payments, and the constantly reducing role of cash in society, central banks have shown an increased interest in deploying central bank digital currencies (CBDCs) that could serve as a digital cash-equivalent. While most recent research on CBDCs focuses on blockchain technology, it is not clear that this choice of technology provides the optimal solution. In particular, the centralized trust model of a CBDC offers opportunities for different designs.

In this paper, we depart from blockchain designs and instead build on ideas from traditional e-cash schemes. We propose a new style of building digital currencies that combines the transaction processing model of e-cash with an account-based fund management model. We argue that such a style of building digital currencies is especially well-suited to CBDCs. We also design the first such digital currency system, called Platypus, that provides strong privacy, high scalability, and expressive but simple regulation, which are all critical features for a CBDC. Platypus achieves these properties by adapting techniques similar to those used in anonymous blockchain cryptocurrencies like Zcash to fit our account model and applying them to the e-cash context.

1 INTRODUCTION

Recent research on digital currencies has mostly focused on blockchains such as Bitcoin [32] instead of traditional e-cash systems such as [14]. This is mostly due to the popularity of blockchains for permissionless digital currencies, i.e., digital currencies that do not rely on a trusted central authority.

Inspired by the popularity of blockchains, several central banks, such as Swedish central bank [39] and the Bank of England [9], have expressed interest in creating a digital version of their currency. The People's Bank of China [44] has already deployed a digital yuan into trial use. Recently, several central banks, together with the Bank of International Settlements, have outlined the principles and core features of such a central bank digital currency (CBDC) [8].

A CBDC has a different trust model and different requirements from permissionless cryptocurrencies. Namely, the central bank is generally a trusted authority and the consensus process should not be open to everyone. Nevertheless, decentralized ledgers have often been proposed for such central bank digital currencies [5, 17, 42], since they offer benefits over traditional e-cash such as increased robustness due to the distributed consensus, as well as transferability due to the ledger-based system. While traditional e-cash [14] provides privacy for the sender, it leaks the transaction amounts since the coins need to be deposited immediately for double-spending protection. In ledger-based systems, coins are not deposited, but instead value is transferred, which allows for private transactions such as in Zerocash [38].

However, e-cash systems have several advantages compared to ledger-based systems. Namely, e-cash systems are easier to scale, mainly because they do not require byzantine agreement between independent parties due to their centralized nature. This has particular implications on their sharding potential. For example, depositing coins can easily be sharded based on the serial number of the deposited coin. Since the coins are signed by the central authority, there is no need to check (potentially cross-shard) if the coin was produced as an output of a previous transaction (such as in ledger-based systems), and instead it suffices to check that the coin is signed by the central authority and that the serial number has not been seen before. Further, the requirements for clients can potentially be reduced compared to ledger-based systems, since in ledger-based systems, clients keep up to date with the whole ledger or use a lightweight client, which reduces their privacy [24] without the use of additional mechanisms that require additional trust assumptions [30, 43] or expensive cryptographic protocols [29].

We want to leverage the different trust model of central bank digital currencies and combine the benefits of ledger-based digital currencies and traditional e-cash schemes. Namely, we assume an authority that is trusted for the integrity of the currency (e.g. double-spending protection) but is not trusted for privacy, a setting that has been proposed by several central banks [8, 16]. We want to make use of the performance benefits from traditional e-cash schemes, but combine them with a transaction mechanism inspired by anonymous ledger-based cryptocurrencies like Zerocash [38] that provides anonymity for the sender and recipient as well as secrecy of the transaction amounts. In addition, the mechanism should be easy to extend with regulation mechanisms for e.g. money laundering protection similar to [23, 42].

To achieve these goals, as the first main contribution of this paper, we propose a new style of building digital currencies that combines the transaction processing model of e-cash payments

*Work partially done while at ETH Zurich

with an account-based model for managing users' funds (which is also used in some ledger-based systems like Ethereum [41]). We argue that this style of building digital currencies is particularly well-suited to CBDCs and allows us to achieve strong privacy, high scalability, and simple but expressive regulation, which are all desirable features for a CBDC.

As the second main contribution of this paper, we design the first digital currency system, called Platypus[1], that follows this design pattern. Platypus is also inspired by previous anonymous blockchain-based cryptocurrencies such as Zerocash [38]. In Platypus, each participant owns an account that is represented by a commitment, called *account state commitment*, to a serial number and a balance and which is signed by the central bank. A transaction then consists of updating the commitments of both, the sender and recipient. The sender and recipient reveal the serial number of their current account states, prove in zero-knowledge that they are in possession a corresponding state commitment signed by the central bank, and that the sum of their balances remains invariant.

Such a design provides advantages over anonymous ledger-based designs as well as over UTXO-based designs (e.g. Zcash). The main advantage over ledger-based designs lies in the scalability of such an approach. Since transactions do not need to be ordered on a ledger and the system is centralized, transaction validation can be sharded almost arbitrarily using standard database techniques such as two-phase commit [28] and thus there is no inherent limit on its throughput. In addition, the requirements for clients are reduced significantly. In systems like Zcash, clients need to download and decrypt every transaction stored on the ledger if they want to benefit from Zcash's privacy guarantees. If a currency is to be widely used, as expected from a CBDC, downloading and decrypting every transaction quickly becomes infeasible for most users, who may want to use this currency on a mobile device.

Another advantage of the account-based design is simplified enforcement of regulatory rules. If enrollment in the system is bound to real identities, this account-based design can simplify regulatory rules similar to those proposed by Garman et al. [23] and Wüst et al. [42]. In particular, it enables enforcement of regulatory rules that, e.g., limit the amount of funds that a particular user can posses at a time (as mentioned e.g. by the Bank of England [9]), or that require disclosure of the user's identity if a certain limit for receiving funds within a period of time is exceeded. However, in contrast to [23], this enforcement is more flexible since it can be applied on the sender or recipient side (instead of only the sender) and it can be done more efficiently since it does not require aggregation over multiple transaction outputs. In contrast to [42], it preserves full transaction unlinkability.

This design also provides advantages over more traditional e-cash schemes like the original proposal by Chaum [14] and optimizations using similar principles [11, 12], in which a bank issues blinded coins to a user, who then spends them at one or more merchants, who deposit them back in the bank. Namely, one of the main advantages is that our account-based approach does not require spending of individual coins, which has two important effects.

[1]Similar to how its namesake combines features from different animals, Platypus combines ideas from e-cash, blockchains, and bank accounts.

First, this leads to a more compact scheme, since the transaction size and the verification cost do not increase with the transaction value. In traditional e-cash schemes, each coin is spent individually which means that the transaction size depends on the transaction value whereas the size can stay constant in an account-based design. Even in divisible e-cash [13] – which shrinks the transaction size to logarithmic in the value – each coin is still deposited individually, i.e. the verification cost at the bank is linear in the transaction value.

Second, and more importantly, this improves privacy: In traditional e-cash systems, the amount that a merchant receives always leaks, since they need to deposit the received coins at the bank. In the case of an online e-cash system, this immediately leaks all transaction values of the merchant, in an offline system, this leaks the amount that is received between two deposits. Thus, traditional e-cash systems only provide *payer privacy*, but not *recipient privacy* or *value privacy*. With our account-based design, the size of a transaction is independent of its value and the funds are immediately deposited in the blinded account of the recipient. This ensures the anonymity of the sender and recipient, the confidentiality of the transaction value, and the unlinkability of transactions.

Contributions. In this paper, we make the following contributions:

- A *new pattern of building digital currencies* that combines the transaction processing model of e-cash with an account-based fund management model.
- A *new digital currency design* called Platypus that provides unlinkable transactions, high scalability, and privacy-preserving regulation mechanisms which are all critical features for a CBDC.
- A *security analysis* that shows that Platypus provides integrity and strong privacy guarantees.
- *An implementation and evaluation* that show that transaction creation is fast and Platypus can be easily scaled.

2 OVERVIEW

In this section we provide an overview of Platypus. We start by explaining our motivation and goals, followed by the trust assumptions and our system model. After that, we explain the main ideas of Platypus.

2.1 Motivation & Goals

Recently, multiple central banks together released a report detailing the principles, motivations and risks of CBDCs [8]. This serves as a good basis for technical decisions in the design of a CBDC since it directly provides the view of the involved central banks.

One of the main motivations outlined in [8] is continued access to central bank money, i.e. the function of a CBDC as a form of a "digital banknote". Currently, both, access and the use of cash are declining in many jurisdictions, which creates the risk that some businesses and households lose access to risk-free central bank money. A CBDC could step in to fill this void to ensure the confidence in a currency.

Cash does not only provide risk-free central bank money, but it also provides very strong privacy guarantees. In a cash payment, third parties neither learn the identities of the parties nor the value. This is a property that should also be mirrored by a CBDC [5, 8]. A

working paper from the Swiss National Bank [16] explicitly mentions "mass surveillance" as one of the potential threats of a CBDC, which exemplifies the need for strong privacy guarantees and a consultation from the European Central Bank [21] showed that privacy is the most important feature of a CBDC for the survey respondents.

A CBDC could also increase resilience and the diversity of payment systems, improve financial inclusion, and simplify cross-border payments if the CBDCs of multiple countries are interoperable [8]. Lastly, even though this is not stated as one of the main motivations of [8], a CBDC could be used for "programmable monetary policy" to e.g. provide so-called "helicopter drops" that distribute funds to the public combined with an expiration date for spending these funds.

CBDCs also create some risks for financial stability [5, 8]. In particular, it can lead to a form of bank runs, since it provides a convenient way (in contrast to paper money) of storing their funds as central bank money. One of the potential mitigations for this risk is to explicitly design the currency as a cash-like system that, e.g., enforces limits on how much currency can be held by a single party at a time. Because of this, allowing the enforcements of such limits is one of the central regulatory goals for such a digital currency.

Another regulatory requirement for CBDCs is the enforcement of anti-money-laundering (AML) legislation [5, 8]. However, this partially conflicts with the goal of improved payment privacy. This conflict can be solved by allowing anonymous payments up to a given limit per unit of time above which the recipient needs to disclose their identity to a regulator. This idea has also been proposed by the European Central Bank in the form of "anonymity vouchers" [20] as well as by previous work [23, 42].

Based on these motivations and ideas, we focus on a *retail CBDC* that can be used as a digital equivalent of cash since this is the main use case considered by central banks [8]. Other settings for CBDCs exist, e.g., replacement of a settlement layer or to replace bank transfers. In such settings other properties may become relevant, such as offline receiving that we discuss in Appendix B.

Given our focus, our main goal is to provide a digital currency that is maintained by a central bank and provides fully anonymous transactions, i.e., where the transaction values are secret, the sender and recipient cannot be identified and transactions are unlinkable to previous or future transactions. In addition, this solution should make use of the benefits allowed by the trust model in which a central authority is trusted for integrity (as proposed, e.g., in [16]) and should provide significant performance benefits over other anonymous digital currencies such as Zerocash [38].

As a secondary goal, we want this digital currency to be easily and efficiently extendable with regulation mechanisms similar to those described by Garman et al. [23] and Wüst et al. [42] to make it viable for the use as fiat currency.

2.2 System Model & Trust Assumptions

Motivated by the considerations in Section 2.1, we consider the setting in which a *central bank* wants to issue and maintain a digital currency, as shown in Figure 1. Such a centralized design is proposed by a working paper of the Swiss National Bank [16] and suggested as one possible option in a report from a group of central banks [8].

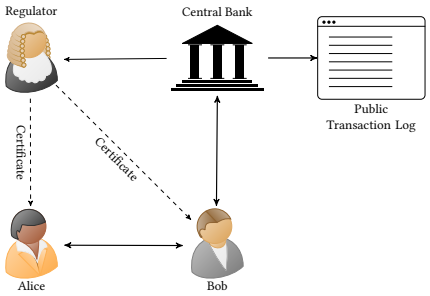

Figure 1: Platypus System model. Platypus consists of a central bank that is responsible for transaction validation, a regulator that issues certificates to clients and receives transaction information relevant for compliance with regulatory rules, as well as clients that participate in the system. The central bank also publishes a log of all transactions.

In addition to this central bank, we assume that there exists a *regulator* (e.g., a government agency), which is responsible for enforcing regulatory requirements, such as anti-money-laundering (AML) legislation. While such a regulator is not necessary for the functioning of the core protocol, it would likely be an integral part in any deployment of a CBDC in practice (see Section 4).

Our system also contains *clients* that can act as payment senders and payment recipients. We assume that these clients are considered untrusted, i.e., they may behave arbitrarily.

Since central banks are responsible for monetary policy, we assume that the central bank is trusted for the integrity of the currency and the regulator trusts the central bank to comply with regulatory requirements. The central bank is responsible for the issuance of new money and preventing double-spending is in its own interest as double-spending would effectively increase the currency in the system.

However, based on our considerations in Section 2.1 and the potential threat of mass surveillance [16], we assume that the central bank is not trusted for privacy, i.e., it might be interested in deanonymizing the sender or recipient of a transaction or recover transaction values.

We consider full protection against network-based deanonymization attacks (e.g., linking an IP address to multiple transactions) to be out of scope of this paper. To protect against such attacks, clients can use protections such as anonymity networks like Tor [3] if desired. However, to provide some resilience against such attacks, we design the system such that the recipient and sender of a transaction cannot be easily linked together by the central bank, even without the use of anonymity networks and provide some discussion about network-based transaction linking in Appendix B. In the normal case (cooperating recipient) this is achieved by only having the recipient communicate with the central bank. For other cases, and to simplify account recovery (see Appendix B), the central bank publishes transactions in a publicly accessible log, which clients can use to look up their previous transactions. This log can be mirrored by third parties, similar to block explorers in blockchain systems.

Finally, we assume that clients communicate with each other through secure channels and that all cryptographic primitives used are secure according to the standard definitions for their security: we assume that commitments are computationally binding and hiding, that signatures are unforgeable, that the zero-knowledge proof systems are zero-knowledge and provide soundness, and that encryption is CPA-secure.

2.3 Platypus Design

Platypus uses a hybrid between an account model and an e-cash design, in which each participant is responsible for keeping track of their own *account state* which is kept as objects similar to coins in an e-cash system. However, in contrast to e-cash, where a client usually has multiple coins that can be used in a transaction, a client has a single account state which is consumed in every transaction and a replaced with a new one. This account state is represented by an *account state commitment* state_i to the account balance bal_i and to a serial number serial_i. The account state commitment is produced by a previous transaction and is signed by the central bank. To sign these state commitments, the central bank uses its secret key sk_C (corresponding to public key pk_C). For enforcement of regulatory policies, the account state may contain additional information as described in Section 4.

Figure 2 shows how a transaction is processed in the normal case where both the sender and recipient have already participated in the system. In step ①, Alice initiates a transaction, in which she sends a value of v_{Tx} to the recipient, Bob, by creating a *sender account update*. Alice creates a commitment, called *transaction commitment*, to the value v_{Tx} using a random blinding factor blind_{Tx}, denoted by $\text{comm}_{\text{Tx}} = \text{comm}(v_{\text{Tx}}, \text{blind}_{\text{Tx}})$. She then creates a new state commitment state^A_{i+1} that commits to a fresh pseudorandomly (based on her longterm key) chosen serial number serial^A_{i+1} and a value $\text{bal}^A_{i+1} = \text{bal}^A_i - v_{\text{Tx}}$ where bal^A_i is the balance committed to in her current state commitment state^A_i. Alice then creates a non-interactive zero-knowledge proof zkp^A_{i+1} which proves that she performed these steps correctly.

Note that, for this zero-knowledge proof, both the previous account state commitment state^A_i and the central bank's signature are secret values, i.e. they are not revealed in this transaction. This ensures that this transaction is not linkable to the previous transaction in which state^A_i was created and which contains state^A_i and the bank's signature. The zero-knowledge proof potentially also needs to prove compliance with regulatory rules, if a regulation mechanism as described in Section 4 is in place.

This zero-knowledge proof zkp^A_{i+1}, as well as the transaction commitment comm_{Tx}, the serial number of the old state serial^A_i, and the new account state commitment state^A_{i+1} are then sent to the recipient, Bob. Alice also provides the random value blind_{Tx} required to open the commitment comm_{Tx}, such that Bob can use it to create a zero-knowledge proof for his own account update.

To complete the transaction (step ②), Bob then creates a creating a *receiver account update*, for which he proceeds similarly to Alice, with the difference that his zero-knowledge proof zkp^B_{j+1} reuses the transaction commitment comm_{Tx} and proves that his account balance in his new state state^B_{j+1} increases by exactly v_{Tx} compared to his previous state state^B_j with serial number serial^B_j.

Once Bob has created the proof zkp^B_{j+1}, he sends the transaction commitment comm_{Tx}, Alice' and his serial numbers (serial^A_i, serial^B_j), both of their new state commitments (state^A_{i+1}, state^B_{j+1}), and both zero-knowledge proofs (zkp^A_{i+1}, zkp^B_{j+1}) to the central bank.

The central bank then executes the transaction (step ③) by verifying both zero-knowledge proofs and checking that neither of the serial numbers (serial^A_i, serial^B_j) have been used in previous transactions. If this is the case, the central bank adds both serial numbers to the set of used serial numbers, signs the two new state commitments (state^A_{i+1}, state^B_{j+1}) with their private key sk_C and sends the signatures $\sigma^A_{i+1} = \text{Sign}(sk_C, \text{state}^A_{i+1})$ and $\sigma^B_{j+1} = \text{Sign}(sk_C, \text{state}^B_{j+1})$ back to Bob, who checks if the signatures are valid and, if so, accepts the payment (step ④). Bob then forwards σ^A_{i+1} to Alice, who verifies the signature and updates her stored state information, which completes the payment (step ⑤).

The central bank keeps a record of all recent (i.e., for some specified time interval chosen by the central bank) transactions, which they publish in a publicly accessible way. In particular, for each transaction, the bank publishes all values received from Bob, as well as the bank's signatures on the new account state commitments. This allows Alice to check the set of recently published transactions for the serial number of her old account state to find and receive the transaction containing the signed new state commitments, even if Bob does not forward this information to her. Note that the published set of transactions does not need to be ordered (in contrast to a ledger or blockchain) and can be mirrored by arbitrary parties. To enable efficient account backups and recovery, additional encrypted (with a key of the owner) information about the contents of the transaction can be included in the transaction. We describe backups and account recovery in Appendix B.

The centralized design of Platypus also simplifies sharding, since standard database sharding techniques can be used for checking and updating the serial numbers of the used account states. We describe the sharding potential in more detail in Appendix B.

3 PLATYPUS BASE TRANSACTION DETAILS

In this section, we describe the details of the base transactions in Platypus, i.e., the creation of transactions without regulation mechanisms. We defer the explanation of regulation mechanisms to Section 4 to improve readability and to make the system design easier to understand.

Platypus makes use of zero-knowledge proofs in its transactions. These zero-knowledge proofs can be instantiated with different proof techniques, but the statements that are proven are independent of these techniques. In Section 6, we implement and evaluate Platypus using zk-SNARKs.

Some of these proof-techniques (including that by Groth [26] used in our implementation) require a *trusted setup* to generate a common reference string, which has been criticized in the context of decentralized cryptocurrencies like Zcash, and extraordinary efforts were made to keep it secure when Zcash was originally launched [34]. It is important to note here that, at least in most

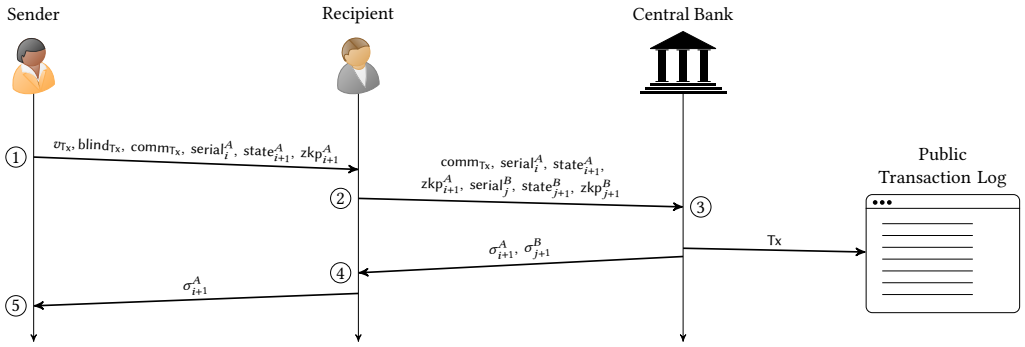

Figure 2: Platypus Base Transaction. ① *Transaction Initiation.* The sender, Alice, creates the transaction commitment as well as the proofs for the update of her account state commitment and sends all of this to the recipient Bob, together with the blinding value of the transaction commitment. ② *Transaction Completion.* Based on this, Bob creates the proofs for the update of his account state commitment and sends the new account state commitments and both proofs to the central bank. ③ *Transaction Execution.* The central bank verifies the proofs, checks that the revealed serial numbers have not been used before, and based on this either accepts or rejects the transaction. If the bank accepts, it signs the new account states and sends the signatures back to Bob. Simultaneously, the central bank also publishes the full transaction (i.e. everything received from Bob plus the new signatures) on a public transaction log. ④ *Payment Acceptance.* If the signatures are valid, Bob accepts the payment and forwards the signature on Alice' state to her, which completes the payment (⑤ *Payment Completion.*).

constructions including [26], a compromise of this trusted setup does not affect the zero-knowledge property of the proofs. Instead it "only" affects soundness, which in the context of digital currencies allows the creation of money, but does not affect privacy. In our context, i.e., a CBDC, the central bank is the entity that creates money and is trusted for the integrity of the currency, which means that it can therefore be trusted to perform the setup without any additional assumptions. Of course, in practice, it could be preferable to distribute the setup between multiple parties or to use a proof system that does not require a trusted setup.

3.1 System Setup

To set up the system, the central bank creates a private/public key pair (sk_C, pk_C) that is used for signing account state commitments and publishes its public key pk_C. In addition, if a proof system is used that requires the setup of a common reference string (see above), the central bank runs the trusted setup procedure (possibly in conjunction with other parties). In addition, the central bank sets a parameter $\mathsf{bal}_{\mathsf{max}}$ which is a maximum limit on account balances to prevent value overflows and can be set to a value larger than all realistic values for account balances. Similarly, the entity responsible for regulation generates the parameters required by the regulation, which we describe in Section 4.

User Enrollment. When a user U enrolls in the system, they create a secret key $sk_U = (sk_{U1}, sk_{U2})$, consisting of two randomly chosen keys sk_{U1}, sk_{U2}. These keys can later be used to pseudorandomly derive serial numbers and blinding values for their account states using pseudorandom functions $f_{sk_{U1}}$ and $g_{sk_{U2}}$. Pseudorandom serial numbers prevent possible attacks that could destroy funds [37]. Blinding values could instead also be chosen randomly. However, using pseudorandom values for both simplifies the creation of backups for an account (see Appendix B). If regulation is in place, users may also need to register their identity with the regulator (see Section 4).

To create an enrollment transaction, U derives pseudorandom values $\mathsf{serial}_1^U, \mathsf{blind}_1^U$ from their secret key as $\mathsf{serial}_1^U = f_{sk_{U1}}(0)$ and $\mathsf{blind}_1^U = g_{sk_{U2}}(0)$ and uses them to create a new state commitment $\mathsf{state}_1^U = \mathsf{comm}(\mathsf{serial}_1^U, \mathsf{bal}_1^U, \mathsf{blind}_1^U)$ for an account with no balance, i.e., $\mathsf{bal}_1^U = 0$. U then creates a non-interactive zero-knowledge proof zkp_1^U which proves that the account state commitment corresponds to an account with balance zero, i.e., zkp_1^U proves the following statement:

Given public value

$$\mathsf{state}_1^U$$

I know secret values

$$sk_{U1}, \mathsf{serial}_1^U, \mathsf{blind}_1^U$$

such that

$$\mathsf{state}_1^U = \mathsf{comm}(\mathsf{serial}_1^U, 0, \mathsf{blind}_1^U)$$
$$\mathsf{serial}_1^U = f_{sk_{U1}}(0)$$

U then sends state_1^U and zkp_1^U to the central bank. The bank checks if the proof is correct and then signs the account state commitment state_1^U and sends the signature back to U.

3.2 Transaction Creation

Here, we describe how a transaction between a sender (Alice) and a recipient (Bob) is created. We assume that clients keep all values secret unless mentioned otherwise and that they communicate through secure channels.

Alice' current account state is represented by a commitment $\text{state}_i^A = \text{comm}(\text{serial}_i^A, \text{bal}_i^A, \text{blind}_i^A)$, and similarly Bob's current account state is represented by $\text{state}_j^B = \text{comm}(\text{serial}_j^B, \text{bal}_j^B, \text{blind}_j^B)$ if he already has an account. In addition, both are in possession of a signature from the central bank on their account state commitment, denoted by $\sigma_i^A = \text{Sign}(sk_C, \text{state}_i^A)$ and $\sigma_j^B = \text{Sign}(sk_C, \text{state}_j^B)$, respectively. The commitment can be created using any hiding and binding commitment scheme. The steps correspond to the steps shown in Figure 2.

① Transaction Initiation:

(i) To create a transaction to Bob with value v_{Tx}, Alice chooses a fresh random value blind_{Tx} and creates a commitment $\text{comm}_{\text{Tx}} = \text{comm}(v_{\text{Tx}}, \text{blind}_{\text{Tx}})$.

(ii) Alice also derives pseudorandom values $\text{serial}_{i+1}^A, \text{blind}_{i+1}^A$ from her secret key as $\text{serial}_{i+1}^A = f_{sk_{A1}}(\text{serial}_i^A)$ and $\text{blind}_{i+1}^A = g_{sk_{A2}}(\text{blind}_i^A)$ and creates a new account state $\text{state}_{i+1}^A = \text{comm}(\text{serial}_{i+1}^A, \text{bal}_i^A - v_{\text{Tx}}, \text{blind}_{i+1}^A)$.

(iii) Alice then creates a non-interactive zero-knowledge proof zkp_{i+1}^A that proves the following statement:
Given public values

$$\text{serial}_i^A, \text{comm}_{\text{Tx}}, \text{state}_{i+1}^A, \text{bal}_{\max}, pk_C$$

I know secret values

$$sk_{A1}, \text{bal}_i^A, \text{bal}_{i+1}^A, \text{blind}_i^A, \sigma_i^A, v_{\text{Tx}}, \text{blind}_{\text{Tx}}, \text{serial}_{i+1}^A, \text{blind}_{i+1}^A$$

such that

$$\begin{aligned}
\text{True} &= \text{Vrfy}(pk_C, \text{comm}(\text{serial}_i^A, \text{bal}_i^A, \text{blind}_i^A), \sigma_i^A) \\
\text{comm}_{\text{Tx}} &= \text{comm}(v_{\text{Tx}}, \text{blind}_{\text{Tx}}) \\
\text{state}_{i+1}^A &= \text{comm}(\text{serial}_{i+1}^A, \text{bal}_{i+1}^A, \text{blind}_{i+1}^A) \\
\text{bal}_{\max} &\geq \text{bal}_{i+1}^A \\
\text{bal}_{i+1}^A &= \text{bal}_i^A - v_{\text{Tx}} \\
\text{serial}_{i+1}^A &= f_{sk_{A1}}(\text{serial}_i^A)
\end{aligned}$$

(iv) Alice then sends $v_{\text{Tx}}, \text{blind}_{\text{Tx}}, \text{comm}_{\text{Tx}}, \text{serial}_i^A, \text{state}_{i+1}^A, \text{zkp}_{i+1}^A$ to Bob.

② Transaction Completion:

(i) After receiving the partial transaction from Alice, Bob derives pseudorandom values $\text{serial}_{j+1}^B, \text{blind}_{j+1}^B$ from his secret key as $\text{serial}_{j+1}^B = f_{sk_{B1}}(\text{serial}_j^B)$ and $\text{blind}_{j+1}^B = g_{sk_{B2}}(\text{blind}_j^B)$ and uses them to create a new account state $\text{state}_{j+1}^B = \text{comm}(\text{serial}_{j+1}^B, \text{bal}_j^B + v_{\text{Tx}}, \text{blind}_{j+1}^B)$.

(ii) If Bob already has an account, Bob creates a non-interactive zero-knowledge proof zkp_{j+1}^B that, similar to Alice' proof (with the difference of proving that his balance *increased* by the transaction value), proves the following statement:
Given public values

$$\text{serial}_j^B, \text{comm}_{\text{Tx}}, \text{state}_{j+1}^B, \text{bal}_{\max}, pk_C$$

I know secret values

$$sk_{B1}, \text{bal}_j^B, \text{bal}_{j+1}^B, \text{blind}_j^B, \sigma_j^B, v_{\text{Tx}}, \text{blind}_{\text{Tx}}, \text{serial}_{j+1}^B, \text{blind}_{j+1}^B$$

such that

$$\begin{aligned}
\text{True} &= \text{Vrfy}(pk_C, \text{comm}(\text{serial}_j^B, \text{bal}_j^B, \text{blind}_j^B), \sigma_j^B) \\
\text{comm}_{\text{Tx}} &= \text{comm}(v_{\text{Tx}}, \text{blind}_{\text{Tx}}) \\
\text{state}_{j+1}^B &= \text{comm}(\text{serial}_{j+1}^B, \text{bal}_{j+1}^B, \text{blind}_{j+1}^B) \\
\text{bal}_{\max} &\geq \text{bal}_{j+1}^B \\
\text{bal}_{j+1}^B &= \text{bal}_j^B + v_{\text{Tx}} \\
\text{serial}_{j+1}^B &= f_{sk_{B1}}(\text{serial}_j^B)
\end{aligned}$$

(iii) Finally, Bob sends the values $\text{comm}_{\text{Tx}}, \text{serial}_i^A, \text{state}_{i+1}^A, \text{zkp}_{i+1}^A, \text{serial}_j^B, \text{state}_{j+1}^B, \text{zkp}_{j+1}^B$ to the central Bank.

③ Transaction Execution:

(i) The central Bank checks that none of the serial numbers $\text{serial}_i^A, \text{serial}_j^B$ appear in its stored set of previously used serial numbers and that both zero-knowledge proofs $\text{zkp}_{i+1}^A, \text{zkp}_{j+1}^B$ verify. If this is not the case, then the central bank rejects the transaction and informs Bob.

(ii) Otherwise, the central bank accepts the transaction and adds both serial numbers to the set of previously used serial numbers, signs the new state commitments as $\sigma_{i+1}^A = \text{Sign}(sk_C, \text{state}_{i+1}^A)$ and $\sigma_{j+1}^B = \text{Sign}(sk_C, \text{state}_{j+1}^B)$ and sends them to Bob. In addition, the central bank publishes the transaction (i.e., all values received from Bob plus σ_{i+1}^A and σ_{j+1}^B) on a publicly available log.

④ Payment Acceptance: Bob checks that the signatures received from the central bank are valid, accepts the payment and stores σ_{j+1}^B to update his account if this is the case, and forwards σ_{i+1}^A to Alice. Otherwise, he rejects the payment and informs Alice.

⑤ Payment Completion: Alice checks that the signature received from Bob is valid. Otherwise, or if she has not received a signature from Bob after a timeout, she inspects the central bank's public transaction log to retrieve the transaction and the signature on her new account state commitment. She then stores σ_{i+1}^A to update her account and the payment is completed.

4 REGULATION IN PLATYPUS

As described in Section 2.1, a CBDC requires the possibility to enforce regulatory policies. In particular, a CBDC should enable rules that ensure the financial stability of a system, e.g., to prevent bank runs, as well as rules that allow enforcement of anti-money-laundering legislation or allow the detection of tax evasion [5, 8].

The design of Platypus explicitly simplifies the implementation of such compliance policies through its account-based design. This account-based design allows storing additional information within an account state, which enables efficient zero-knowledge proofs through which the account holder can prove compliance with a given rule. In particular, it improves efficiency over previous designs such as that of Garman et al. [23] that require proofs over the state of the whole system (inclusion of several UTXO in a Merkle tree) instead of a proof of a signature. In contrast to the design by Garman et al., which only allows proofs on the state of the sender, it also allows proofs about the state of the recipient. In addition, it improves privacy compared to designs like PRCash [42] that require linking several transactions together for efficiency.

In this section, we describe a generic framework for enabling such regulatory policies. In Appendix A, we describe in detail two examples for such policies that are in line with the goals of a CBDC as stated by several central banks [8]. Namely, the first example puts limits on how much currency a user can hold without declaring it to authorities. The second example limits how much currency a user can receive anonymously within a given time period. The goals of this second example are similar to that of the *anonymity vouchers* as proposed by the European Central Bank [20] as well as the previous proposals by Garman et al. [23] and Wüst et al. [42]. In Appendix A, we also show that our regulation mechanism ensures compliance with the enforced policies.

4.1 Regulation Framework

Meaningful compliance rules need to be bound to a recognized identity. Otherwise, a user could establish a large number of pseudonymous identities to circumvent these rules. This requires an entity responsible for establishing these identities.

In addition, many practical rules do not simply prevent someone from taking an action but instead require them to disclose information under certain conditions. We therefore also assume the existence of a government agency that is responsible for receiving such information and operating on it, e.g., within the legal system. For simplicity, we assume that these roles are taken on by a single entity that we call the *regulator*. However, in practice, the responsibilities could be split, e.g., one entity could be responsible for establishing identities and a separate agency could hold the responsibility for each compliance rule. As part of the system setup, the regulator creates one key pair for issuing certificates (sk_{RC}, pk_{RC}), another key pair that is used for encryption (sk_{RE}, pk_{RE}), and publishes both public keys.

Enrollment. To enable regulation, users need to explicitly enroll in the system and establish identities. To do this, and to later be able to prove their identity, each user U generates a random secret value, called *secret identity* si_U from which their *public identity* $pi_U = \mathsf{PubID}(si_U)$ is derived, i.e., essentially a private/public key pair with the sole purpose of identifying the user. The user then needs to receive a *certificate* $\sigma_R^U = \mathsf{Sign}(sk_{RC}, (pi_U, params))$, i.e., a signature from the regulator on the user's public identity, as well as potentially some other individual parameters ($params$) that can be used for different rules on an individual basis. For example, the certificate could contain a holding limit that is individual to each user. This can be useful to, e.g., allow retail businesses to hold a larger amount of currency than users can hold in private accounts.

To issue this certificate, the user proves knowledge of the secret identity corresponding to their public identity pi_U, which is then, together with the individual regulation parameters $params$, signed by the regulator after confirming the real identity of the user (e.g., by the user physically going to an office of the responsible government agency). The user's certificate and secret identity can then later be used in zero-knowledge proofs for anonymous identification. To ensure that this identity cannot be used for multiple accounts, the public identity is always included in the account state commitment and the user proves equality of the public identities committed to in the old and new state commitments.

Structure of regulated transactions. For the enforcement of some regulatory rules, it can be useful to keep track of information involving the user's transaction history, which then allows the user to create proofs involving this information when creating a transaction. To enable this, such *auxiliary information* can be committed to in the account state commitment of the user. The account state of a user U is thus represented by a commitment $\mathsf{state}_i^U = \mathsf{comm}(\mathsf{serial}_i^U, \mathsf{bal}_i^U, pi_U, \mathsf{aux}_i^U, \mathsf{blind}_i^U)$ where aux_i^U denotes the required auxiliary information.

The zero-knowledge proof of the base protocol (see Section 3), is then extended such that the user also proves in zero-knowledge that they comply with the regulation rules. This includes proving that they know a private identity for which they have a certificate from the regulator. The regulator can require the user to disclose some information (e.g., the user's identity, balance, transaction value etc.) under certain specified conditions. We denote the computation of this information with a function RegInfo which takes the user's state information as input and either outputs either a fixed dummy value (if the condition is not triggered) or the information that is required to be disclosed (if the condition is triggered), e.g., the user's identity pi_U and account balance bal_{i+1}^U.

The user encrypts the output of this function with the regulator's public encryption key pk_{RE}, resulting in a ciphertext E_{i+1}^U and proves that the computation of RegInfo and the encryption was correctly performed. Encrypting a dummy value if the transaction is fully compliant ensures that all transactions are indistinguishable to parties other than the regulator independent of triggering the condition. The ciphertext E_{i+1}^U is sent to the bank as part of the transaction, which forwards it to the regulator, who can then decrypt it (and discard it, if it is the dummy value).

Below, we show the general structure of the updated proof statement for regulated transactions. The function updateAux is used to update the auxiliary information aux and the function checkOther is a predicate that can contain additional checks that would cause the generation of the proof to fail. E.g., this could be used to impose hard limits on the amount that a user can hold in their account (see Appendix A). These functions, as well as RegInfo are dependent on the policies that are enforced. All of these functions can also depend on additional public information aux_{pub} such as the current date. To improve readability, the differences to the proof of the base transaction (see Section 3.2) are shown in purple.

Given public values

$$\mathsf{serial}_i^U, \mathsf{comm}_{\mathsf{Tx}}, \mathsf{state}_{i+1}^U, E_{i+1}^U, \mathsf{bal}_{\max}, pk_C, pk_{RC}, pk_{RE}, \mathsf{aux}_{pub}$$

I know secret values

$$si_U, pi_U, \sigma_R^U, params, \mathsf{aux}_i^U, \mathsf{aux}_{i+1}^U$$
$$sk_{U1}, \mathsf{bal}_i^U, \mathsf{bal}_{i+1}^U, \mathsf{blind}_i^U, \sigma_i^U, v_{\mathsf{Tx}}, \mathsf{blind}_{\mathsf{Tx}}, \mathsf{serial}_{i+1}^U, \mathsf{blind}_{i+1}^U$$

such that

$$\mathsf{True} = \mathsf{Vrfy}(pk_C, \mathsf{comm}(\mathsf{serial}_i^U, \mathsf{bal}_i^U, pi_U, \mathsf{aux}_i^U, \mathsf{blind}_i^U), \sigma_i^U)$$
$$\mathsf{comm}_{\mathsf{Tx}} = \mathsf{comm}(v_{\mathsf{Tx}}, \mathsf{blind}_{\mathsf{Tx}})$$
$$\mathsf{state}_{i+1}^U = \mathsf{comm}(\mathsf{serial}_{i+1}^U, \mathsf{bal}_{i+1}^U, pi_U, \mathsf{aux}_{i+1}^U, \mathsf{blind}_{i+1}^U)$$
$$\mathsf{bal}_{\max} \geq \mathsf{bal}_{i+1}^U$$
$$\mathsf{bal}_{i+1}^U = \mathsf{bal}_i^U + v_{\mathsf{Tx}}$$
$$\mathsf{serial}_{i+1}^U = f_{sk_{U1}}(\mathsf{serial}_i^U)$$
$$pi_U = \mathsf{PubID}(si_U)$$

$$\text{True} = \text{Vrfy}(pk_{RC}, (pi_U, params), \sigma_R^U)$$
$$\text{aux}_{i+1}^U = \text{updateAux}(\text{aux}_i^U, \text{aux}_{pub}, pi_U, params, v_{\text{Tx}}, \text{bal}_{i+1}^U)$$
$$E_{i+1}^U = \text{Enc}(pk_{RE}, \text{RegInfo}(\text{aux}_i^U, \text{aux}_{pub}, pi_U, params, v_{\text{Tx}}, \text{bal}_{i+1}^U))$$
$$\text{True} = \text{checkOther}(\text{aux}_i^U, \text{aux}_{pub}, pi_U, params, v_{\text{Tx}}, \text{bal}_{i+1}^U)$$

Depending on the type of compliance rule in place, not all parts are necessary. For example, a holding limit (see Appendix A) does not require committing to any auxiliary data aux_i^U. We show the proof statement for a recipient here, but this can be equally applied to the sender (with the only difference being an increase vs. decrease of the balance). Two example policies that illustrate regulation proofs are discussed in Appendix A.

5 SECURITY ANALYSIS

In this section, we analyze the security of Platypus, in particular its integrity and privacy guarantees.

5.1 Transaction Integrity

We first discuss the integrity of our system. Since Platypus is a digital currency system, this entails that only authorized parties should be able to spend funds or create funds and funds should not be spendable more than once. In particular, the system should provide *transaction unforgeability* and *balance invariance*. We define these two properties below and show that our system provides them.

Transaction unforgeability essentially ensures that only authorized parties can create transactions that spend their respective funds and that the transaction values and intended recipients cannot be changed by an adversary. Balance invariance ensures that an adversary cannot spend funds multiple times or increase the supply of the currency. We capture the first of these properties with the following *transaction forgery* game:

Definition 5.1 (Transaction Forgery Game)**.** Given our system, the game consists of an interaction between an adversary $\mathcal{A}$ and a challenger C with access to an oracle O that simulates honest parties in the system. The game proceeds as follows:

(1) C initializes the system with a security parameter λ, which is used by the system to in turn initialize all used primitives, such as the signature scheme or the zero-knowledge proof system. C also initializes the oracle O.
(2) $\mathcal{A}$ can then generate arbitrary private keys and associated accounts with a balance chosen by $\mathcal{A}$, which O enrolls in the system by signing the associated account state commitments.
(3) $\mathcal{A}$ can also ask O to initialize additional clients with balances chosen by $\mathcal{A}$. O initializes them with the specified balance by signing an according account state commitment and then sends the signed account state commitment and serial number for each of them to the adversary.
(4) $\mathcal{A}$ can use his accounts to create arbitrary transactions, interact arbitrarily (i.e. send or receive transactions) with any account managed by the oracle, or can ask the oracle to create transactions between accounts managed by the oracle which are created and forwarded to the adversary. All transactions created in interaction with O are added to a query set Q.
(5) For each of these transactions, the adversary can then decide to submit them to O for execution, where O acts as central bank, performs the same checks as the central bank and either accepts or rejects the transaction.
(6) The adversary wins the game if they can create a transaction that is accepted by the oracle (simulating the central bank) in the *transaction execution step* that does not appear in the query set Q and is either
   - a transaction in which $\mathcal{A}$ controls neither the sender nor the recipient account
   - a transaction in which $\mathcal{A}$ controls the recipient account, but not the sender account and no transaction with the same sender serial number and the same transaction value, and for which the adversary controls the recipient account, exists in Q

Claim 5.1 (Transaction Unforgeability)**.** *No computationally bounded adversary $\mathcal{A}$ without access to the simulation trapdoor of the zero-knowledge proof system can win the transaction forgery game with non-negligible probability.*

Proof Sketch. Assume such an adversary $\mathcal{A}$ exists. Then there are two possible cases to distinguish: Either 1) the adversary forges a valid account update for the sender that is not part of any transaction in Q, or 2) he reuses a valid sender account update from a transaction $\text{Tx}_Q \in Q$.

In Case 1, $\mathcal{A}$ either a) creates a valid account update for an account not controlled by $\mathcal{A}$ without knowing the respective secret values, b) gains knowledge of the secret values, or c) creates a valid account update for a non-existing account.

In case 1a), $\mathcal{A}$ must be able to create a zero-knowledge proof that is accepted by the central bank without knowing the secrets, thus violating our assumption that the zero-knowledge proof system is sound. In 1b) $\mathcal{A}$ must be able to compute the sender's secret values based on previously seen transactions, in particular also the blinding value used to create the previous account state. Since this blinding value is only used for the account state commitment, which is never opened, such an adversary could be used to distinguish commitments to two different pairs of serial numbers and account balances, which violates our assumption that the commitment scheme is hiding.

In case 1c), $\mathcal{A}$ either needs to produce a signature from the central bank on a forged account state commitment or they need to produce a proof of knowledge of such a signature without having knowledge of it. If $\mathcal{A}$ can produce either of them, then this adversary $\mathcal{A}$ can also be used to either break soundness of the zero-knowledge proof system or to win the signature forgery game, which violates our assumptions.

Now consider case 2. Then $\mathcal{A}$ either a) does not control the recipient account for the transaction Tx_Q from Q, or b) controls the recipient account for Tx_Q. In case 2a) $\mathcal{A}$ does not know the blinding value used to create the transaction commitment and needs to either find a transaction $\text{Tx}'_Q \in Q$ for which the transaction commitment is the same as in Tx_Q (to reuse its recipient state update) which is negligible, or $\mathcal{A}$ needs to create a recipient account update that uses the transaction commitment from Tx_Q which is analogous to case 1.

In case 2b) $\mathcal{A}$ controls the recipient account of Tx_Q and therefore needs to create a transaction Tx' with a different recipient account update that changes the transaction value. In this case, the adversary knows the blinding value used to create the transaction

commitment since he controls the recipient account used in Tx_Q. However, since the commitment scheme is binding, $\mathcal{A}$ cannot open the commitment to any value other than the originally committed value, and since we assume the proof system to be sound, $\mathcal{A}$ can therefore not create any recipient account update that changes the recipient's balance by any other value. Thus, $\mathcal{A}$ cannot create any such transaction Tx' without violating either the binding property of the commitment scheme or soundness of the proof system.

Since all possible cases violate at least one assumption, Platypus provides transaction unforgeability. □

Claim 5.2 (Balance Invariance)**.** *No computationally bounded adversary without access to the simulation trapdoor of the zero-knowledge proof system can create a transaction that increases the available funds in the system or spends funds more than once.*

Proof Sketch. There are multiple cases to distinguish. An adversary can either 1) attempt to use the same sender account state in multiple transactions, 2) attempt to use a sender account state that never resulted from a transaction accepted by the central bank, or 3) attempt to create a transaction that increases the balance of the recipient by more than it decreases the balance of the sender.

First, let us consider the case where an adversary attempts to use the same account state multiple times as sender in a transaction. Similar to traditional e-cash schemes like [14] as well as Zerocash [38], double spending is prevented using serial numbers that uniquely define an account state and can only be used once. Once the serial number serial_i^A has been revealed for one account state commitment state_i^A, the same account state can no longer be used for future updates, since reusing the account state would require proving that the same account state commitment opens to a different serial number serial'^A_i. If the adversary can create such a proof, then either the proof system is not sound or the commitment scheme used to create the state commitment is not binding, both of which contradict our assumptions. No client can therefore use the same account state for more than one transaction.

Now consider the case where an adversary creates a transaction that uses a sender account state that has never been the result of a transaction accepted by the central bank. This would immediately allow the adversary to win the transaction forgeability game and thus violates at least one of our assumptions.

Lastly, consider the case where an adversary attempts to create a transaction that increases the account balance of the receiver by more than the value subtracted from the account balance of the sender. Since the value of each transaction is committed to using the transaction commitment $\mathsf{comm}_{\mathsf{Tx}}$, which is created using a hiding and binding commitment scheme, no computationally bounded party can open the commitment to a transaction value other than what was committed to originally. Since the proof of the transaction sender proves that their account balance was decreased by exactly the committed value and the proof of the transaction recipient proves that their balance was increased by exactly this value, any adversary that could increase the recipients balance by a different value could be used to either break soundness of the zero-knowledge proof system or to break the binding property of the commitment scheme. Thus, the account balance of the recipient is increased by exactly the amount that the balance of the sender is decreased and the transaction does not increase the total amount of funds available in the system. □

5.2 Transaction Privacy

Here, we consider the privacy guarantees provided by Platypus. In particular, we consider privacy towards parties other than the regulator and show that accepted transactions in our system are indistinguishable. In Appendix A we discuss what additional information the regulator receives. We do not consider network-level attacks on anonymity here, as they are out of scope of this paper, but we provide a short discussion of such attacks in Appendix B.

We capture the privacy guarantees with the following *transaction indistinguishability* game:

Definition 5.2 (Transaction Indistinguishability Game)**.** Given our system, the game consists of an interaction between an adversary $\mathcal{A}$ and a challenger C with access to an oracle O that simulates honest parties in the system. The game proceeds as follows:

(1) C initializes the system with a security parameter λ, which is used by the system to in turn initialize all used primitives, such as the signature scheme or the zero-knowledge proof system. C also initializes the oracle O.
(2) $\mathcal{A}$ can then generate arbitrary private keys and associated accounts with a balance chosen by $\mathcal{A}$, which O enrolls in the system by signing the associated account state commitments.
(3) $\mathcal{A}$ can also ask O to initialize additional clients with balances chosen by $\mathcal{A}$. O initializes them with the specified balance by signing an according account state commitment and then sends the state commitment and serial number for each of them to the adversary.
(4) $\mathcal{A}$ can use his accounts to create arbitrary transactions, interact arbitrarily (i.e. send or receive transactions) with any account managed by the oracle, or can ask the oracle to create transactions between accounts managed by the oracle which are created and if they result in a valid transaction, they are executed (i.e. the states of the involved parties are updated) and forwarded to the adversary.
(5) In the challenge phase, $\mathcal{A}$ chooses parameters (i.e. sender, recipient, value) for two transactions Tx_0 and Tx_1, such that the adversary controls neither the sender nor the recipient account and the transaction value does not exceed the sender's balance and sends these parameters to C.
(6) C chooses a bit $b \in \{0, 1\}$ u.a.r., executes the transaction Tx_b and sends the resulting transaction to $\mathcal{A}$
(7) $\mathcal{A}$ then outputs a bit b' and wins the game if $b = b'$

Claim 5.3 (Transaction Indistinguishability)**.** *No computationally bounded adversary $\mathcal{A}$ can win the transaction indistinguishability game with non-negligible advantage.*

Proof Sketch. As stated in Section 2.2, we assume that all used cryptographic primitives are secure according to their respective notions. In particular this includes that the pseudorandom function is indistinguishable from a truly random function, the commitments to different values are indistinguishable, the zero-knowledge proof system provides zero-knowledge (i.e. we have access to a simulation oracle $\mathcal{S}$ that can simulate indistinguishable proofs for any statement), and that the encryption scheme provides CPA-indistinguishability.

We now show that no efficient adversary $\mathcal{A}$ can succeed in winning the game with non-negligible advantage using a hybrid argument. To that end consider two set of distributions $T_0^0, T_0^1, \ldots, T_0^9$ and $T_1^0, T_1^1, \ldots, T_1^9$ for the challenge transactions Tx_0 and Tx_1, respectively in which we gradually replace fields in the transactions through an idealized version. That is, T_k^0 (for $k \in \{0,1\}$) is the distribution for the real transaction Tx_k, T_k^1 replaces the sender zero-knowledge proof zkp_{i+1}^A with a simulated proof (from $\mathcal{S}$), T_k^2 additionally replaces the sender serial number serial_i^A with the output of a truly random function, T_k^3 also replaces the sender's state commitment state_{i+1}^A with a commitment to randomly chosen account parameters, and T_k^4 replaces the encrypted regulation information E_{i+1}^A with the encryption of a random value. The same is repeated for the recipient's part of the transactions for the distributions $T_k^5, \ldots, T_k^8$, and finally T_k^9 also replaces the transaction commitment c_{Tx} with a commitment to a random value.

T_0^9 and T_1^9 are therefore distributions in which all fields in the transaction have been replaced with random values (sampled according to the distribution resulting from truly random inputs to the respective functions) and the zero-knowledge proofs are simulated based on these random values. A special case is the serial number, which is replaced by the output of a truly random function with a previous serial number as input. However, since all previous serial numbers are unique for transactions accepted by the central bank, the output is also truly random. Therefore T_0^9 and T_1^9 are the same distributions and thus indistinguishable for any adversary.

Assume that we have an arbitrary adversary $\mathcal{A}$ that wins our game with non-negligible advantage, i.e., that can successfully distinguish T_0^0 and T_1^0. Thus, for some non-negligible function p, we have $\left|\Pr\left[\mathcal{A}(T_0^0) = 1\right] - \Pr\left[\mathcal{A}(T_1^0) = 1\right]\right| \geq p(\lambda)$. Due to the triangle inequality, we also have:

$$
\begin{aligned}
&\left|\Pr\left[\mathcal{A}(T_0^0) = 1\right] - \Pr\left[\mathcal{A}(T_1^0) = 1\right]\right| \\
&\leq \sum_{i=1}^{9} \left|\Pr\left[\mathcal{A}(T_0^{i-1}) = 1\right] - \Pr\left[\mathcal{A}(T_0^i) = 1\right]\right| \\
&+ \sum_{i=1}^{9} \left|\Pr\left[\mathcal{A}(T_1^{i-1}) = 1\right] - \Pr\left[\mathcal{A}(T_1^i) = 1\right]\right| \\
&+ \left|\Pr\left[\mathcal{A}(T_0^9) = 1\right] - \Pr\left[\mathcal{A}(T_1^9) = 1\right]\right|
\end{aligned}
$$

Since the last term is zero (as T_0^9 and T_1^9 are the same distribution), at least one of the other terms must be non-negligible, i.e. $\left|\Pr\left[\mathcal{A}(T_k^{i-1}) = 1\right] - \Pr\left[\mathcal{A}(T_k^i) = 1\right]\right| \geq p'(\lambda)$ for some $i \in \{1, \ldots, 9\}, k \in \{0,1\}$ and some non-negligible function p'. Since the only difference between these two distributions is that one of them replaces one of the fields with a value that is indistinguishable (according to the respective notion of the used primitive), this leads to a contradiction. Therefore, Platypus provides transaction indistinguishability.

□

5.3 Availability of Funds

While we do not consider network-level attacks on availability, our system should ensure that a client cannot be prevented from using their funds by a third party. For example, Ruffing et al. [37] described an attack on Zerocoin [31], in which an attacker invalidates coins from another user by creating and immediately spending coins with the same serial number as that of an honest user, which prevents the honest user from using their funds. Since Platypus also uses serial numbers to prevent double-spending, we need to consider similar attacks. In particular, we make the following claim:

Claim 5.4. *No computationally bounded adversary can invalidate the account state of another client.*

Proof Sketch. First, note that in order to prevent a client from creating a transaction that updates their account state, either some information necessary to create the account state update needs to be withheld from the client, or the adversary needs to cause the central bank to reject the transaction. We assume that the client does not lose access to their long term keys and private information and thus they can always retrieve all necessary information from the central bank's transaction log.

Since the central bank will always accept a valid transaction unless it reuses a previously seen serial number, the adversary can only make the central bank reject an account update from a client by creating a transaction that uses the same serial number as used by the honest client (as in [37]).

To invalidate a user's account state with serial number serial_i^U, the adversary needs to create an account update that reveals the same serial number and they need to prove that this serial number was committed to in a valid state commitment for which they know the corresponding secret key. Thus, the adversary needs to create a series of account states that at some point results in the same serial number serial_i^U, i.e. they need to find a secret key sk' and an index j, such that $f_{sk'}^{\circ j}(0) = \mathsf{serial}_i^U = f_{sk_{U1}}^{\circ i}(0)$ (where $f_x^{\circ k}$ is the k-times iterated composition of f_x and k is bounded by an arbitrary but fixed value n (polynomial in the security parameter)).

Since f_x is a pseudorandom function, so is $h_{(x,k)} = f_x^{\circ k}$ for a randomly chosen key (x,k) where $k \in \mathbb{Z}_n^+$ (by induction). A successful adversary as described above would therefore need to find a key for the pseudorandom function family h that produces the given input/output pair which is infeasible. □

6 EVALUATION

6.1 Implementation

We implemented Platypus using the gnark[1] library for the zero-knowledge proofs with the BN256 curve and the Groth16 proof system [26]. Our implementation covers benchmarks for the creation and verification of the zero-knowledge proofs, as well as a simple 'end-to-end' system to measure throughput.

For the signatures and commitments, we use the gadgets as provided by the library for EdDSA [10] signatures and MiMC [4] hashes. Our prototype also uses MiMC for the pseudorandom function to generate serial numbers. Blinding values for commitments are randomly chosen in our prototype. To provide public key encryption for our regulation mechanism, our implementation uses Elgamal encryption [18]. Our implementation covers the base transaction as well as transactions with two regulation enforcement rules. These rules can be toggled individually and put limits on the amount of money that can be received within a given time interval or held in

Table 1: Performance of Platypus. This table shows proving and verification, as well as the time required for the trusted setup and the number of R1CS constraints. All measurements are averaged over 100 runs and rounded to two significant figures.

	Trusted Setup [s]	Proving [s]	Proving iPhone [s]	Verification [s]	# R1CS constraints	Tx Size [B]
Base Tx	0.73	0.11	0.19	0.000 89	11 728	672
Tx with holding limit	2.8	0.37	0.69	0.000 93	47 356	800
Tx with receiving limit	2.9	0.37	0.69	0.000 94	48 631	800
Tx with both limits	3.4	0.43	0.80	0.000 92	61 344	864

the account before the user is required to report this information to the regulator (see Section 4 and Appendix A). The proof generation and verification in the gnark library is parallelized.

Throughput Benchmark. Our throughput benchmark consists of a simple (non-optimized) server and a client that generates the full transaction, i.e., simulates both sender and recipient. The *server* exposes the functionality of the central bank as a simple REST interface with JSON payloads. Serial numbers submitted by clients in valid transaction proofs are stored in a local SQLite [2] database.

The *client library* handles all client-side functionality, i.e., generation of keys and transaction proofs, account state management, and communication with the server via HTTP.

To measure the transaction throughput that the server can handle, we use our client library to first prepare a large number of transactions (10k for our measurement). For each transaction, the client generates and enrolls two accounts and then creates a transaction between these pairs of accounts. In the measurement phase, the client library then submits these transactions to the server and measures the time required for all submitted transactions to complete.

6.2 Results

We measured the proving and verification time of our implementation for the base transaction and regulated transactions with receiving and holding limits (see Appendix A for more detail on the policies). We also measured the time required for the trusted setup, which is a one-time operation only run during system setup. As can be seen in Table 1, this setup is quite fast, taking less than ten seconds for all configurations.

Table 1 shows the results of our measurements as well as the number of R1CS constraints for our zero-knowledge proofs and the transaction sizes. R1CS is an intermediate format, used in many zero-knowledge proof systems, which represents the proof constraints and thus provides a system-independent measure of the proof complexity. We performed these measurements on a machine with an Intel® Core™ i7-7700 CPU (3.60GHz) with 4 cores and 16GB of RAM. We also measured proving time on an Apple iPhone 13 mini to benchmark performance on a mobile device. The results for proof generation and proof verification are provided per proof, i.e., both the sender and the recipient have to perform a proof generation and the central bank needs to perform two proof verifications per transaction. The client-side proof generation could be done in parallel after first communicating the values used to create the transaction commitment ($\mathsf{blind}_{\mathsf{Tx}}$ and v_{Tx}), i.e., the transaction sender and recipient can compute their proofs simultaneously to reduce the total transaction creation time. Our results show that this can be done efficiently. Transaction sizes are based on 256-bit serial numbers and commitments and show the size of the transaction before execution, i.e., when they are submitted to the central bank. After execution (in the log), they additionally include two signatures, thus, an additional 128 bytes.

With both regulation mechanisms in place, the proof generation takes 0.4 seconds on our test machine and 0.8 seconds on the iPhone (see Table 1). This makes it feasible to perform the complete transaction within one second, which is often considered an important limit for usability [33] and which makes it usable for retail payments. In addition, our numbers show that the overhead of adding additional regulation mechanisms is small. Concretely, while adding any regulation adds a significant overhead in the proving time compared to the base transaction, the additional impact of enforcing a second rule is small and only increases the proving time by a tenth of a second on the iPhone and a twentieth on our larger test machine (while verification time stays constant), which shows that Platypus can easily support enforcement of multiple regulatory rules.

The fast transaction verification is constant independent of the size of the proof statement (i.e., regardless of the regulation mechanisms in use) and allows a single machine to process a large number of transactions. In our throughput benchmark, our machine achieves a throughput of 922 transactions per second.

Even though CBDCs are not intended to replace all other forms of payments, only to complement them [8], it is interesting to consider the feasibility of such a system for all payments in a large economic area. Data from the European Union show that in 2016, the EU population performed 163 billion payments [19] for a population of just below 450 million people [22]. This corresponds to a volume of slightly more than five thousand transactions per second on average, or if we assume that all of these payments take place within only 8 hours of each day (to exclude times with a low transaction volume), a volume of about 15.5 thousand transactions per second.

Thus, to handle all transactions in the EU, a deployment of Platypus would require the equivalent of approximately 17 of our test machines (with its 922 transactions per second), which is a modest requirement for such a large economic area. Put differently, assuming the same transaction volume per person and again assuming that all transactions are concentrated on 8 hours per day, a single machine would be able to easily handle the transactions of a small country like Switzerland (≈ 300 transactions per second), Israel (≈ 320 transactions per second), or Sweden (≈ 350 transactions per second).

7 RELATED WORK

E-Cash Systems. With e-cash [14], Chaum introduced the first design for an anonymous digital currency, in which a user can withdraw a coin from a bank by generating a coin identified by a serial

number and receiving a blind signature on it, which ensures that the bank does not see the serial number. The user later unblinds this signature, which allows them to use the coin for payments. A merchant receiving a payment deposits the coin at the bank, at which point the bank checks if the serial number has already been used. If that is the case, the bank rejects the payment, otherwise it is accepted.

E-cash makes withdrawal and spending of a coin unlinkable, but it reveals to the bank the total transaction volume of a client (based on their withdrawals) and the value of each transaction for every merchant (based on their deposits). It also requires users to store information linear in the number of coins that they own. Later designs [11, 13] reduce the overhead. Camenisch et al. later also proposed an e-cash system that offers a form of regulation [12], limiting the amount that can be spent anonymously by a user per merchant. However, in all previous e-cash designs, the merchant still reveals the value of their received coins to the bank when depositing them.

Baldimtsi et al. [7] used techniques for double-spending detection for a transferable e-cash design, in which a coin can be transferred to different users without interaction with the bank. Once a coin gets deposited, the bank then checks for double-spending and identifies the offending party. This removes the issue that the merchant needs to reveal the transaction value to the bank for all received funds. Unfortunately, such a transferable e-cash scheme necessitates that coins grow in size depending on how often they were used, which makes spending less efficient than other e-cash schemes. This also affects linkability, since coins of a different size (i.e. coins that have been used a different number of times) are distinguishable.

Blockchain-based Systems. Several proposals for anonymous cryptocurrencies exist in the blockchain space. Zerocash [38] and its instantiation Zcash is currently considered to provide the strongest privacy guarantees. All of the transaction information is completely hidden and transactions are unlinkable, similar to the guarantees provided by Platypus. Garman et al. later showed how Zerocash can be extended with accountability mechanisms [23] that put restrictions on the transaction sender. One of the main drawbacks of Zerocash and the proposal by Garman et al. are the heavy client requirements which are difficult to remove or reduce in a decentralized setting [43]. This is particular due to the transaction receiving mechanism, which requires decrypting every transaction included in the blockchain as well as the requirement to prove knowledge of the path of a transaction output in a Merkle tree, which requires clients to keep this tree up to date. The second also makes scaling more difficult, since adding new transaction outputs to this tree requires all transactions to be serialized. In contrast, Platypus can take advantage of the changed trust assumptions to provide better scalability and to reduce the requirements for clients.

Other recent research has proposed schemes to provide regulation in a semi-centralized blockchain setting. PRCash [42] provides a design that uses lightweight zero-knowledge proofs to efficiently enable a receiving limit per time interval (*epoch*) for anonymous transactions. However, PRCash is based on a transaction design called mimblewimble [27] that does not provide full unlinkability for transactions and the regulation mechanism requires linking several transactions within an epoch. Platypus therefore provides better privacy and at the same time improves scalability through its centralized design. In addition, Platypus simplifies regulation compared to the designs of [23] and [42] due to its account-based design, since it does not require the inclusion of multiple UTXOs in proofs.

Parallel work by Androulaki et al. [6] proposed an auditable anonymous token management system for use in a permissioned blockchain targeted towards enterprise networks. In contrast to Platypus, which is account-based, their design uses a UTXO model, in which the UTXOs are represented as Pedersen commitments [35]. They then use a combination of a permissioned blockchain and a potentially distributed *certifier* to authorize payments. Transactions are committed to the blockchain after proving that the spender has a signature on each spent UTXO and later the newly created UTXOs get signed by the *certifier* using randomizable signatures [36]. In addition, the scheme allows for a set of *auditors*, each of which is responsible for auditing a different set of participants and which can access all information of their assigned participants. Platypus instead allows for fully anonymous transactions as long as specified conditions are not violated and is extendible with different regulatory rules. The regulation mechanisms enabled by Platypus make it more suitable for the use as central bank digital currency in which most transactions should be equivalent to cash with respect to their privacy properties [8]. In contrast, the auditability provided by the design from Androulaki et al. [6] is targeted at business-to-business usecases in which each business has their own auditor who should be able to access all of the transaction information of the business.

Another parallel design by Tomescu et al. [40] called UTT uses a similar base design to that of Androulaki et al. [6]. Namely, UTXOs (called coins) resulting from a transaction are represented as homomorphic commitments and signed with randomizable signatures [36] by a bank (that can be distributed using threshold cryptography). To later use these commitments as inputs, they are re-randomized, a nullifier (used to prevent double-spending) is deterministically computed and revealed by the sender and then the sender proves that the sum of the output coins is the same as the sum of the inputs. UTT also provides a monthly anonymity budget that limits how much money can be sent anonymously. Platypus, in contrast, enables more expressive regulation policies and can also enforce them on the side of the recipient. Similar to Zerocash [38], to receive a UTT payment, the recipient also has to scan all transactions on a ledger and perform a trial decryption for each. In addition, UTT transactions are 16× larger than ours and transaction verification time is significantly slower than in Platypus.

Finally, parallel work by Gross et al. [25] proposes the use of a modified Zerocash [38] for a “privacy pool” of a CBDC. Similar to Platypus, it replaces UTXOs with accounts, but in contrast to the e-cash style transaction execution of Platypus, it requires proofs of inclusion in a Merkle tree (like Zcash). In addition, regulation in [25] only allows hard limits (per transaction or for the account balance) and is thus less expressive and versatile than in our system.

8 CONCLUSION

Despite the prominence of blockchain-based digital currencies, they may not be the best technology choice for issuing a CBDC. Given the trust model of CBDCs (central authority) and the desirable features of a CBDC (privacy, performance, scalability, regulation), we argue that a traditional e-cash scheme can be a more suitable

starting point for designing CBDCs. With our solution Platypus we have shown that an e-cash like system can provide all these features at the same time.

We have also proposed a new style of building digital currencies that combines e-cash style transaction processing with the account-model that is common in blockchain systems like Ethereum [41] and with privacy techniques inspired by Zerocash [38]. We hope that our work can inspire other researchers to design new e-cash solutions that leverage the design pattern proposed in this paper, extend our work, and ultimately provide better CBDC designs.

REFERENCES

[1] [n.d.]. gnark Library. https://github.com/ConsenSys/gnark. (Cited on p. 10)
[2] [n.d.]. SQLite. https://www.sqlite.org/. (Cited on p. 11)
[3] [n.d.]. Tor Browser. https://www.torproject.org/. (Cited on p. 3, 15)
[4] Martin Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. 2016. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In *International Conference on the Theory and Application of Cryptology and Information Security*. (Cited on p. 10)
[5] Sarah Allen, Srdjan Čapkun, Ittay Eyal, Giulia Fanti, Bryan A Ford, James Grimmelmann, Ari Juels, Kari Kostiainen, Sarah Meiklejohn, Andrew Miller, Eswar Prasad, Karl Wüst, and Fan Zhang. 2020. *Design Choices for Central Bank Digital Currency: Policy and Technical Considerations*. Technical Report. The Brookings Institution. (Cited on p. 1, 2, 3, 6, 14)
[6] Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhiyaoui, and Björn Tackmann. 2020. Privacy-preserving auditable token payments in a permissioned blockchain system. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. (Cited on p. 12)
[7] Foteini Baldimtsi, Melissa Chase, Georg Fuchsbauer, and Markulf Kohlweiss. 2015. Anonymous Transferable E-Cash.. In *Public Key Cryptography*. (Cited on p. 12, 16)
[8] Bank of Canada, European Central Bank, Bank of Japan, Sveriges Riksbank, Swiss National Bank, Bank of England, Board of Governors of the Federal Reserve, and Bank for International Settlements. 2020. Central bank digital currencies: foundational principles and core features. https://www.bis.org/publ/othp33.htm. (Cited on p. 1, 2, 3, 6, 7, 11, 12, 14, 16)
[9] Bank of England. 2020. Central Bank Digital Currency: Opportunities, challenges and design. https://www.bankofengland.co.uk/-/media/boe/files/paper/2020/central-bank-digital-currency-opportunities-challenges-and-design.pdf. (Cited on p. 1, 2, 14)
[10] Daniel J Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. 2012. High-speed high-security signatures. *Journal of cryptographic engineering* 2, 2 (2012). (Cited on p. 10)
[11] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. 2005. Compact E-Cash. In *Advances in Cryptology - EUROCRYPT 2005 (Lecture Notes in Computer Science, Vol. 3494)*. (Cited on p. 2, 12, 16)
[12] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. 2006. Balancing accountability and privacy using e-cash. In *International Conference on Security and Cryptography for Networks*. (Cited on p. 2, 12, 16)
[13] Sébastien Canard, David Pointcheval, Olivier Sanders, and Jacques Traoré. 2015. Divisible e-cash made practical. In *IACR International Workshop on Public Key Cryptography*. Springer, 77–100. (Cited on p. 2, 12)
[14] David Chaum. 1983. Blind Signatures for Untraceable Payments. In *Advances in Cryptology: Proceedings of Crypto 82*. (Cited on p. 1, 2, 9, 11, 16)
[15] D. Chaum, A. Fiat, and M. Naor. 1990. Untraceable Electronic Cash. In *Proceedings on Advances in Cryptology (CRYPTO '88)*. (Cited on p. 16)
[16] David Chaum, Christian Grothoff, and Thomas Moser. 2021. How to issue a central bank digital currency. *SNB Working Papers* (2021). (Cited on p. 1, 3, 16)
[17] George Danezis and Sarah Meiklejohn. 2016. Centrally Banked Cryptocurrencies. In *23nd Annual Network and Distributed System Security Symposium, NDSS*. (Cited on p. 1)
[18] T. Elgamal. 1985. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory* 31, 4 (1985). (Cited on p. 10)
[19] Henk Esselink and Lola Hernández. 2017. The use of cash by households in the euro area. *ECB Occasional Paper* 201 (2017). (Cited on p. 11)
[20] European Central Bank. 2019. Exploring anonymity in central bank digital currencies. https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.mipinfocus191217.en.pdf. (Cited on p. 3, 7, 14)
[21] European Central Bank. 2021. Eurosystem report on the public consultation on a digital euro. https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf. (Cited on p. 3)
[22] eurostat. [n.d.]. Population Development and Projections. https://ec.europa.eu/eurostat/web/population-demography-migration-projections/visualisations (retrieved 2021-04-10). (Cited on p. 11)
[23] Christina Garman, Matthew Green, and Ian Miers. 2016. Accountable privacy for decentralized anonymous payments. In *International Conference on Financial Cryptography and Data Security*. (Cited on p. 1, 2, 3, 6, 7, 12, 14)
[24] Arthur Gervais, Srdjan Capkun, Ghassan O Karame, and Damian Gruber. 2014. On the privacy provisions of bloom filters in lightweight bitcoin clients. In *Proceedings of the 30th Annual Computer Security Applications Conference*. (Cited on p. 1)
[25] Jonas Gross, Johannes Sedlmeir, Matthias Babel, Alexander Bechtel, and Benjamin Schellinger. 2021. Designing a central bank digital currency with support for cash-like privacy. *Available at SSRN 3891121* (2021). (Cited on p. 12)
[26] Jens Groth. 2016. On the size of pairing-based non-interactive arguments. In *Annual international conference on the theory and applications of cryptographic techniques*. (Cited on p. 4, 5, 10)
[27] Tom Elvis Jedusor. [n.d.]. Mimblewimble. http://mimblewimble.org/mimblewimble.txt. (Cited on p. 12)
[28] Butler Lampson and Howard E Sturgis. 1979. Crash recovery in a distributed data storage system. (1979). (Cited on p. 2, 15)
[29] Zeyu Liu and Eran Tromer. 2021. Oblivious Message Retrieval. *Cryptology ePrint Archive* (2021). (Cited on p. 1, 16)
[30] Sinisa Matetic, Karl Wüst, Moritz Schneider, Kari Kostiainen, Ghassan Karame, and Srdjan Capkun. 2019. BITE: Bitcoin Lightweight Client Privacy using Trusted Execution. In *28th USENIX Security Symposium (USENIX Security 19)*. 783–800. (Cited on p. 1)
[31] Ian Miers, Christina Garman, Matthew Green, and Aviel D Rubin. 2013. Zerocoin: Anonymous distributed e-cash from bitcoin. In *2013 IEEE Symposium on Security and Privacy*. (Cited on p. 10)
[32] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008). (Cited on p. 1)
[33] Jakob Nielsen. 1994. *Usability engineering*. Morgan Kaufmann. (Cited on p. 11)
[34] Morgen Peck. 2016. The Crazy Security Behind the Birth of Zcash, the Inside Story. *IEEE Spectrum* (2016). https://spectrum.ieee.org/tech-talk/computing/networks/the-crazy-security-behind-the-birth-of-zcash (Cited on p. 4)
[35] Torben Pryds Pedersen. 1992. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *Advances in Cryptology — CRYPTO '91*. (Cited on p. 12)
[36] David Pointcheval and Olivier Sanders. 2016. Short randomizable signatures. In *Cryptographers' Track at the RSA Conference*. (Cited on p. 12)
[37] Tim Ruffing, Sri Aravinda Thyagarajan, Viktoria Ronge, and Dominique Schroder. 2018. (Short Paper) Burning Zerocoins for Fun and for Profit-A Cryptographic Denial-of-Spending Attack on the Zerocoin Protocol. In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*. (Cited on p. 5, 10)
[38] Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. 2014. Zerocash: Decentralized anonymous payments from bitcoin. In *Security and Privacy (SP), 2014 IEEE Symposium on*. (Cited on p. 1, 2, 3, 9, 12, 13)
[39] Sveriges Riksbank. 2020. The Riksbank's e-krona pilot. https://www.riksbank.se/globalassets/media/rapporter/e-krona/2019/the-riksbanks-e-krona-pilot.pdf. (Cited on p. 1)
[40] Alin Tomescu, Adithya Bhat, Benny Applebaum, Ittai Abraham, Guy Gueta, Benny Pinkas, and Avishay Yanai. 2022. UTT: Decentralized Ecash with Accountable Privacy. *Cryptology ePrint Archive* (2022). (Cited on p. 12)
[41] Gavin Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger. (2014). (Cited on p. 2, 13)
[42] Karl Wüst, Kari Kostiainen, Vedran Čapkun, and Srdjan Čapkun. 2019. PRCash: Fast, private and regulated transactions for digital currencies. In *International Conference on Financial Cryptography and Data Security*. (Cited on p. 1, 2, 3, 6, 7, 12, 14, 16)
[43] Karl Wüst, Sinisa Matetic, Moritz Schneider, Ian Miers, Kari Kostiainen, and Srdjan Čapkun. 2019. Zlite: Lightweight clients for shielded zcash transactions using trusted execution. In *International Conference on Financial Cryptography and Data Security*. (Cited on p. 1, 12, 16)
[44] Wolfie Zhao. [n.d.]. Chinese State-Owned Bank Offers Test Interface for PBoC Central Bank Digital Currency. *Coindesk*. 2020. https://www.coindesk.com/chinese-state-owned-bank-offers-test-interface-for-pboc-central-bank-digital-currency (retrieved 2021-04-14). (Cited on p. 1)

A REGULATION DETAILS

In this appendix, we provide descriptions of two example regulation policies, namely limits on how much money can be held in one account and how much money can be received within a given time period. We also show that our regulation mechanism ensures compliance with the policies that are in place.

A.1 Holding Limits

One compliance rule that is of particular interest for financial stability in an economic system, specifically to prevent bank runs, consists of limiting the amount of money that can be held in a CBDC [8, 9]. In addition, such a *holding limit* can be useful to authorities to prevent evasion of wealth tax.

A holding limit can be designed in different ways. The simplest way is to enforce a hard global limit on the amount that can be held by a single account. The only regulation mechanism required to enforce this is the establishment of real identities and proving the possession of a certificate. In addition to this, the value $\mathsf{bal}_{\mathsf{max}}$ that is used in the base transaction (see Section 3) and used to prevent overflows is set to the holding limit required by the regulatory rule which will prevent any balance from exceeding this limit.

A more flexible option could allow different holding limits for different users, for example to allow business accounts to hold more digital currency than private accounts. To do this, this individual holding limit lim_{hold} is included as part of the parameters $params$ in the user's certificate (see Section 4.1). In each transaction, the user then proves in zero-knowledge (i.e. without revealing the limit) that their new balance does not exceed this limit, i.e. the predicate checkOther checks that the user's new balance bal^U_{i+1} is less than the holding limit lim_{hold}. With such hard limits, there is no need to provide a ciphertext E^U_{i+1} with encrypted information for the regulator and the corresponding.

Lastly, it is possible to have soft limits instead of hard limits that allow holding a larger amount of currency with the requirement of revealing this information to the regulator. To enable this, the user's certificate again includes an individual holding limit as before, but the proof in the transaction changes. Instead of proving that they have not exceeded the limit in the transaction, the user encrypts their public identity and their account balance with the regulator's public key if they have exceeded the limit, or fixed dummy values otherwise. That is, the function RegInfo (see Section 4.1) returns the public identity of the user and their balance if the balance is above the limit lim_{hold} and the dummy value otherwise. This ciphertext E^U_{i+1} is added to the transaction and the user proves in zero-knowledge that they have either not exceeded the holding limit and encrypted the dummy value or that they have exceeded the limit and encrypted their public identity and their account balance.

Creating the proof in this way leaks no information to third parties, only to the regulator. The regulator can decrypt the encrypted information and disregard it if it contains the dummy values or keep it otherwise. However, to third parties all transactions are indistinguishable and they do not learn whether a transaction contains real information or dummy values.

A.2 Receiving Limits

Another example for a compliance rule that is commonly suggested for CBDCs is a limit on how much money can be received or spent by a party within a given amount of time [5, 8, 20, 23, 42]. Such a limit serves to emulate reporting requirements for cash transactions that are required for compliance with anti-money-laundering legislation or to prevent tax evasion. Since it is easy to quickly create a large number of digital transactions, these limits should cover a certain amount of time instead of only applying to a single transaction to ensure that they cannot be circumvented by simply splitting a large transaction into multiple smaller transactions.

In the following, we describe how such limits can be added for receiving currency, but the same techniques could also be directly applied for sending currency. Similar to the "anonymity vouchers" proposed by the european central bank [20] and proposed limits in previous work on blockchain-based digital currencies [23, 42], we focus on soft limits that allow for fully anonymous transactions if the total received value for each user is below a given threshold within a fixed time interval, but require reporting if the threshold is exceeded.

Similar to the previous example, the user enrolls in the system where they receive a certificate that includes a receiving limit lim_{rec} as part of the parameters $params$. The system additionally defines *epochs*, the time intervals for which the limits are defined. The length of these epochs is a parameter of the deployed system and can be arbitrary, e.g. a day, a week, or even a year, without affecting the linkability of transactions (in contrast to PRCash [42]). The current epoch number is part of the auxiliary public information aux_{pub}.

For each user, the account state includes two additional pieces of information in its auxiliary information aux^U_i, namely, the last epoch in which the user's account was updated and the cumulative sum of all funds that the user received within that epoch. With each transaction, the function updateAux updates this information accordingly.

Similar to the holding limits above, each transaction includes an encryption E^U_{i+1} (with the regulator's public key) of either the total received value in the current epoch and the recipient's identity – i.e., the function RegInfo (see Section 4.1) returns the public identity of the user and the cumulative epoch total – if the balance is above the limit lim_{rec} or dummy values otherwise. The user then proves (in zero-knowledge) that they performed this correctly, i.e. that the total value that they've received in the current epoch is below the limit or that their correct identity and the correct value were encrypted. Similar to the previous example, to third parties (including the central bank), all transactions remain indistinguishable.

A.3 Regulation Integrity

Since Platypus includes regulation mechanisms, we also need to consider the integrity of this mechanism. In particular, we make the following claim:

Claim A.1. *No client can create a transaction that is non-compliant with a regulation mechanism.*

Proof. This follows directly from the soundness of the zero-knowledge proof system. A transaction will only be valid if the transacting parties prove compliance with the regulatory rules that are in place. For example, if a receiving limit is in place, the recipient proves that either the received amount is within the limit or that the encrypted values attached to the transaction are correct encryptions of their identity and the received value with the public key of the regulator. Since the central bank will only sign updated account state commitments if the corresponding transaction is valid, and by our assumptions, the regulator trusts the central bank to verify this, no client can create a transaction that is non-compliant with the regulation mechanisms that are in place. □

A.4 Privacy towards the Regulator

For any transaction in which the client is not required by the regulation mechanism to include additional encrypted information, the regulator only receives dummy values from decrypting the fields storing regulatory information. Since the dummy values are fixed, the regulator does not gain any additional information from them and thus, these transaction are indistinguishable for the regulator (analogous to Section 5.2).

Of course, since this is the explicit goal of the regulation mechanism, the regulator can decrypt encrypted regulatory information included in a transaction and can thus distinguish them from other transactions and learn additional information about the client, their account and their account history, depending on what information the regulation mechanism requires.

B DISCUSSION

Network level attacks on privacy. As mentioned in Section 2.2, full protection against network level deanonymization attacks is out of scope for this paper. Nevertheless, we designed Platypus to provide some resilience against such attacks. In particular, the sender of the transaction communicates with the central bank through the recipient in the standard case, such that the central bank cannot link the sender and recipient based on the network connections. In exceptional cases, in which the recipient stops cooperating with the sender and does not return the signature on the sender's new account state, the sender can access the public transaction log, or a mirror of this log, to retrieve recent transactions.

While these two mechanisms provide some protection against simple deanonymization attempts, they do not fully protect against all adversaries, in particular if the adversary can see other traffic in the network. If a client is worried about such network level attacks, they can mitigate the risk by using anonymous communication networks such as Tor [3].

Backups and Account Recovery. The account state model that Platypus uses, requires users to have knowledge of their current account state. To enable efficient backups and account recovery, values such as the blinding value of the account state commitment or the serial number are pseudorandomly generated from the user's secret key. To create a backup, the user can simply store this secret key as well as their key and certificate used for regulation.

To recover the account from a backed up secret key, the user needs to retrieve their most recent account state. There are two possibilities to do this. As first option, the user can estimate a time interval in which their most recent transaction took place and retrieve all transactions from that time from the public transaction log. They can then use their secret key and generate serial numbers from it using the pseudorandom generator until they find one that matches a serial number from the log. The second option is that the user generates a list of potential serial numbers (pseudorandomly derived from their secret key), which they then use to query the public transaction log in binary search until they find the latest matching transaction.

The main drawback of the first option is that the user potentially needs to download a large amount of data, if they are unsure in which time interval their latest transaction took place. The drawback of the second option is that some of their transactions can potentially be linked if the central bank is monitoring and correlating queries to the transaction log. Since the number of transactions that could be linked with this approach is only logarithmic in the number of total transactions from the user and no other information about these transactions is revealed, it is unlikely that this presents an issue for most users in practice.

Once the user has retrieved their account state, they need to find their account balance and other values that their account state commits to (i.e. the values used for the regulation mechanism). Without these values they cannot create new transactions. The account balance is much smaller than the serial number and the blinding value and could therefore in principle be brute forced. However, this is inconvenient and can become infeasible if a significant amount of additional information (for the regulation mechanism) is also part of the account state. An easier solution is to add a *memo* field (similar to Zcash) in addition to each account state commitment as part of every transaction, which stores this information encrypted with a long-term key known to the user. The user can then simply include this key in their backup and use it to retrieve all relevant values when performing account recovery after retrieving it together with the account state commitment.

Sharding in Platypus. As mentioned in Section 2, the centralized and account-based design of Platypus simplifies sharding, as it enables the use of standard database sharding techniques.

Figure 3 shows an example of how transaction validation can be sharded. The verification of the zero-knowledge proofs can be performed in separate compute nodes independently from checking and updating the serial numbers of the used account states. While we show each shard here as one database node, each shard can, of course, also be replicated individually. Each database shard is assigned a specified subset of all serial numbers. For example, when using 4 shards, each shard could be assigned a quarter of all possible serial numbers based on the two most significant bits of the serial number. The compute nodes are independent of the transactions. When submitting a transaction, a client can connect to any compute node of the central bank (e.g. through a load balancer), which verifies the zero-knowledge proofs. If the proofs verify, the compute node checks in the database shards if the account states with the provided serial numbers have been invalidated already. Since the serial numbers are pseudorandom, most transactions will be cross-shard transactions if there are at least two database shards. However, since Platypus uses an account-based design, each transaction will never require more than two shards, one for checking the serial number of the sender and one for checking the serial number of the recipient. This is in contrast to UTXO-based systems in which an arbitrary number of shards could be involved in each transaction.

To check the serial numbers in the database shards, the compute node acts as a coordinator in a two-phase commit protocol [28] between the database shards. Each database shard checks if the serial number already exists in the database. If this is the case in one of the shards, the coordinator sends an abort to both shards. Otherwise, they both add the respective serial number to the set of used serial numbers and return a success to the compute node. Finally, the compute node signs the new account states, returns the signatures to the client and publishes the transaction on the public transaction log. Since the transaction log does not require

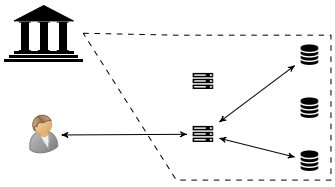

Figure 3: Sharding Potential in Platypus. The central bank can shard both computation and the storage of serial numbers internally. A client can connect to an arbitrary compute node (e.g. through a load balancer) which validates transactions independently from other compute nodes. The compute node then uses a two-phase commit to check and update serial numbers in the database shards corresponding to the serial numbers of the sender and recipient.

ordering, this step can be done concurrently by separate compute nodes without requiring any consensus protocol between them.

Offline Recipient. Most designs of blockchain-based cryptocurrencies allow a recipient to be offline when receiving funds. The sender only needs the recipient's public key to create a full transaction. One limitation of Platypus is that creating a transaction requires interaction between both participants, i.e. the recipient needs to be online to receive funds. This is similar to other e-cash schemes [7, 11, 12, 14, 15], in which the sender and recipient always need to interact. Involving the recipient in the transaction creation is necessary for two main reasons. First, it enables an account-based design with full anonymity. In an account-based system, without involvement of the recipient, some other party would need to be able to update the recipient's account and thus the account would be linkable to a public key of the recipient by that party.

Anonymous UTXO-based systems with offline receiving are possible, but require retrieving this information later, e.g. through downloading all transactions. This puts a heavy load on clients, who need to download and process all of this information. To reduce this load, clients currently either need to rely on trusted execution environments [43], private information retrieval protocols that are very expensive for the server [29], or accept reduced privacy guarantees by trusting a server to filter transactions for them. By directly communicating all necessary transaction information between the sender and recipient, this issue is side-stepped in the e-cash style transaction processing used by Platypus. This guarantees that privacy can be provided efficiently even with a large transaction volume.

Second, and most importantly, including the recipient in the transaction creation is a requirement for enabling regulatory rules affecting the recipient (similar to [42]) that allow full anonymity, even with respect to the regulator, as long as the transaction conforms to some constraints.

As an example, consider a simple holding limit (as described in Section 4) that puts a fixed limit $\mathrm{bal}_{\max}$ on the amount that each party can hold. Let us now assume that there is some mechanism that allows the central bank to check compliance with such a rule without violating any of the privacy properties and without interaction with the recipient. If a sender Alice now creates a transaction of value v_{Tx} with Bob as the recipient, there are two options, which both leak information about Bob's funds to Alice: Either the transaction is accepted, or it is rejected. In the first case, Alice knows that previous to the transaction, Bob owned less than $\mathrm{bal}_{\max} - v_{\mathrm{Tx}}$, in the second case, Alice now knows that Bob owned more than $\mathrm{bal}_{\max} - v_{\mathrm{Tx}}$.

We argue that not enabling offline receiving is a small drawback compared to the advantages of Platypus for our use cases. Recall that Platypus is intended as a "cash-like" CBDC, which is the main goal for many central banks [8, 16]. In such a setting, interaction between the recipient and the sender is the standard case, e.g. for credit card payments or for actual cash payments. Most payments are for retail payments, in which the device of the user interacts with a payment terminal or the user interacts with an online shop, or for peer-to-peer payments between friends, in which their devices can interact. Nevertheless, online receiving does not necessarily require the user to be active, but only their device. For example, if Alice wants to send some funds to Bob and Bob's device is not online, Alice can already initiate the transfer on her device. The device can then, without initiating the actual transaction at that point, contact Bob's device in the background until it becomes available. At that point, the device can initiate the actual transaction and then the payment can complete since both are online.

Bank Market Power and Central Bank Digital Currency

Bank Market Power and Central Bank Digital Currency: Theory and Quantitative Assessment*

Jonathan Chiu
Bank of Canada

Seyed Mohammadreza Davoodalhosseini
Bank of Canada

Janet Jiang
Bank of Canada

Yu Zhu
Bank of Canada

February 20, 2022

Abstract

This paper develops a micro-founded general equilibrium model of payments to study the impact of a central bank digital currency (CBDC) on intermediation of private banks. If banks have market power in the deposit market, a CBDC can enhance competition, raising the deposit rate, expanding intermediation, and increasing output. A calibration to the United States economy suggests that a CBDC can raise bank lending by 1.96% and output by 0.21%. These "crowding-in" effects remain robust, albeit with smaller magnitudes, after taking into account endogenous bank entry. We also assess the role of a non-interest bearing CBDC as the use of cash declines.

JEL Codes: E50, E58.
Keywords: Central bank digital currency; Banking; Market power; Disintermediation; Monetary policy

*This paper was circulated under the title "Central Bank Digital Currency and Banking." We are grateful to Todd Keister for his insightful comments throughout the project. We would like to thank the editor and three anonymous referees who helped us improve the paper. We also thank Jason Allen, James Chapman, Nuno Marques da Paixao, Walter Engert, Rod Garratt, Scott Hendry, Charlie Kahn, Thorsten Koeppl, Oleksiy Kryvtsov, Jiaqi Li, Gradon Nicholls, Peter Norman, Eric Smith, Maarten van Oordt, Francisco Rivadeneyra, Russell Wong, Randall Wright, our Bank of Canada colleagues and many conference participants for their comments and suggestions. The views expressed in this paper are those of the authors and not necessarily the views of the Bank of Canada.

1 Introduction

Many central banks are considering issuing central bank digital currencies (CBDCs), a digital form of central bank money that can be used for retail payments. The Bank for International Settlements surveyed 65 central banks in 2020, covering 72% of the world population and 91% of the world output. Of these central banks, 86% are engaging in work regarding a CBDC; 60% have started experiments or proofs-of-concept for a CBDC; and 14% have moved forward to development and pilot arrangements (see Boar and Wehrli 2021).

In the debate around the impact of introducing a CBDC, one frequently raised concern is that, by competing with bank deposits as a payment instrument, a CBDC could increase commercial banks' funding costs and reduce bank deposits and loans, leading to bank disintermediation. For example, Mancini-Griffoli et al. (2018) caution that a CBDC would force banks to increase their deposit interest rates, and banks would respond by increasing lending rates at the cost of loan demand. The 2018 report by the Committee on Payments and Market Infrastructures of the Bank for International Settlements raises the same concern.

This paper develops a general equilibrium model of banking and payments to assess this disintermediation concern, both theoretically and quantitatively. In this model, banks act as intermediaries, issuing loans to entrepreneurs and creating deposits, which households can use as a means of payment to trade consumption goods. Besides deposits, households have access to two other payment instruments: cash and CBDC. Cash and deposits differ in the types of exchange they can facilitate. For example, cash cannot be used in online transactions while deposits can be used via debit/credit cards or electronic transfers. A CBDC, however, is a perfect substitute for deposits in terms of payment functions and bears an interest set by the central bank.

Our main finding is that introducing a CBDC does not necessarily lead to disintermediation if banks have market power in the deposit market. In this case, the impact of a CBDC is non-monotonic in its interest rate. It expands bank intermediation if its interest rate lies in an intermediate range and causes disintermediation only if its interest rate is set too high. The main mechanism through which a CBDC "crowds in" bank intermediation works as follows. In an imperfectly competitive deposit market, banks restrain the deposit supply to keep the deposit interest rate below the level under perfect competition. A CBDC offers an outside option to depositors and sets an interest rate floor for bank deposits. This floor limits the reduction in the deposit rate and reduces commercial banks' incentive to restrain the deposit supply. If the CBDC rate is not too high, banks supply more deposits, reduce the loan rate, and expand lending.

Interestingly, a CBDC can have a positive effect on deposits, loans, and output even if it has zero market share. The mere existence of a CBDC as an outside option forces banks to match the CBDC rate and create more deposits and loans.[1] A policy implication is that one should assess the effectiveness of a CBDC based on its equilibrium effect on deposits or the deposit rate instead of its usage.

Calibrating our model to the United States economy, we find that a CBDC expands bank intermediation if its interest rate is between 0.30% and 1.49% (for reference, during the calibration period, the average 3-month T-bill rate is about 0.90%). At the maximum, it can increase loans and deposits by 1.96% and the total output by 0.21%. The CBDC leads to disintermediation, however, if its rate exceeds 1.49%. To break even, banks are forced to raise the lending rate to compensate for the interest paid on deposits. As a result, both loans and deposits decrease.[2] Finally, even a non-interest-bearing CBDC can restrict banks' market power and improve intermediation if the use of cash continues to decline. Without a CBDC, banks would limit intermediation and pay negative deposit rates. We have also extended the model to incorporate an imperfectly competitive loan market and endogenous bank entry. A CBDC can still promote bank intermediation, albeit by a smaller magnitude relative to that in the benchmark model.

Our study highlights the role of banks' market power in determining the effects of a CBDC on bank intermediation. The study is closely related to two concurrent papers. Keister and Sanches (2021) focus on the welfare implications of an interest-bearing CBDC when the banking sector is perfectly competitive. They find that, while the CBDC always crowds out bank intermediation, social welfare can still increase when the efficiency in exchange significantly improves, especially when financial frictions are not very severe.[3] In contrast, Andolfatto (2020) studies the effect of a CBDC on banking when there is a monopolistic bank. Using an overlapping generations model, he shows that a CBDC could compel the bank to increase the deposit rate, leading to an increase in bank deposits and financial inclusion. Under the assumption that the central bank offers a lending facility and a deposit

[1] This insight is closely related to that of Lagos and Zhang (2019, 2021), who show that monetary policies discipline the equilibrium outcome by setting the value of the outside option and can be effective even if the use of money approaches zero. Rocheteau et al. (2018) show a related message that money holdings can limit the bank's market power on the lending side.

[2] As suggested by Meaning et al. (2018), an important research question regarding a CBDC is "... at which point do the benefits of a new competitive force for the banking sector get outweighed by the negative consequences of the central bank disintermediating a large part of banks business models?" Our calibration exercise allows us to pin down the interest of a CBDC at which its effect on bank intermediation reverses from positive to negative.

[3] Using a related model, Williamson (2020a) shows that introducing a CBDC to compete with bank deposits can raise welfare by freeing up scarce collateral for banks that are subject to limited commitment.

facility at the same policy rate, the bank's deposit and loan decisions are made separately. The loan rate and quantity are fully determined by the policy rate and are not affected by the CBDC.

Compared to these papers, our framework is more suitable for quantifying the effects of a CBDC and accommodates various design choices as the payment landscape evolves. First, our model captures a complete spectrum of competitiveness. If the number of banks is one, the banking sector is monopolistic, as in Andolfatto (2020). If this number tends to infinity, the banking sector is perfectly competitive as in Keister and Sanches (2021). We use data to discipline the level of competitiveness, which is crucial for quantifying the effects of a CBDC. Second, we explicitly model cash, deposits and a CBDC as three imperfectly substitutable payment instruments that facilitate different types of transactions. This allows us to discuss the design of a CBDC in terms of its acceptability and its effect when the payment landscape evolves, for example, when the use of cash declines.

The economic literature on CBDCs is just emerging, with several lines of research complementary to our work. A number of studies focus on the role of CBDCs as a monetary policy tool. Barrdear and Kumhof (2021) evaluate the macroeconomic consequences of a CBDC in a dynamic stochastic general equilibrium model. Davoodalhosseini (2021) explores the usage of a CBDC for balance-contingent transfers. Dong and Xiao (2021) examine the effects of CBDC rate when CBDC and deposits are complements. Brunnermeier and Niepelt (2019) and Niepelt (2020) derive conditions under which introducing a CBDC has no effects on macroeconomic outcomes, including bank intermediation. Jiang and Zhu (2021) discuss how the interest on a CBDC and the interest on reserves interact as two separate policy tools. Another line of research studies the financial stability implications of a CBDC such as the risk-taking behavior of banks and bank runs. Recent works by Chiu et al. (2020), Fernández-Villaverde et al. (2020), Schilling et al. (2020), Keister and Monnet (2020), Monnet et al. (2020), and Williamson (2020*b*) have made some important progress. Our paper abstracts from these issues and focuses on the effects of a CBDC on bank intermediation in terms of deposit and loan quantities. For research related to the design of a CBDC, see Agur et al. (2020) and Wang (2020). For policy discussions on CBDCs, see Fung and Halaburda (2016); Engert and Fung (2017); Mancini-Griffoli et al. (2018); Chapman and Wilkins (2019); Davoodalhosseini and Rivadenyra (2020); Davoodalhosseini et al. (2020); and Kahn et al. (2020).[4]

[4]Our paper is also related to the literature on private digital currencies and currency competition; see Chiu and Koeppl (2019); Fernández-Villaverde and Sanches (2019); Schilling and Uhlig (2019); Zhu and Hendry (2019); Benigno et al. (2020); Choi and Rocheteau (2020); and Zhou (2020). For a complete introduction to the issues in digital currencies, see Schar and Berentsen (2020).

More broadly, our paper contributes to the monetary theory literature by developing a tractable model with imperfect competition in inside money creation.[5] It is also connected to the literature on how a bank's market power affects monetary policy transmission. Dreschler et al. (2017) provide empirical support for banks' market power in deposit markets and propose a transmission channel accordingly: since a lower nominal interest rate makes cash cheaper to use relative to deposits, banks are compelled to lower the spread between the nominal interest rate and the deposit rate. The effect of a lower nominal interest rate plays a similar role as a higher interest on a CBDC: both policies reduce banks' market power in the deposit market.[6]

The rest of the paper is organized as follows. Section 2 describes the physical environment. Section 3 characterizes the equilibrium. Section 4 calibrates the model and assesses its quantitative implications. Section 5 discusses motivations and implementation of a CBDC. Section 6 concludes and provides some directions for future research. Appendix A provides omitted proofs. Extensions and further discussions are collected in the Online Appendix.

2 Environment

Our model is based on the framework of Lagos and Wright (2005). Time is discrete and continues from zero to infinity. There are four types of agents: a continuum of households with measure 2, a continuum of entrepreneurs with measure 1, a finite number of N bankers (each running a bank), and the government. The discount factor from the current period to the next is $\beta \in (0, 1)$. In each period t, agents interact sequentially in two stages: a frictional decentralized market (DM) and a Walrasian centralized market (CM). There are two perishable goods: y in the DM and x in the CM.

Households are divided into two permanent types, buyers and sellers, each with measure 1. In the DM, a buyer randomly meets a seller. The meeting probability is $\Omega \in (0, 1]$ for both buyers and sellers. The buyer wants to consume y, which is produced by the seller. The buyer's utility from consumption is $u(y)$ with $u'(0) = \infty$, $u' > 0$, and $u'' < 0$. The

[5] Berentsen et al. (2008) models banking in the environment of Lagos and Wright (2005). Gu et al. (2018) demonstrate the inherent instability of banking. Dong et al. (2021) study the effects of competition on bank profits and welfare.

[6] Using a variation of the model of Dreschler et al. (2017), Kurlat (2019) shows that banks' market power raises the cost of inflation. Scharfstein and Sunderam (2016) propose a transmission channel based on banks' market power in the loan market. As the nominal interest rate increases, banks reduce their markup due to lower demand for loans. Wang et al. (2020) estimate a structural banking model and show that the effect of banks' market power in monetary policy transmission is sizable and comparable to that of bank capital regulations.

seller's disutility from production is normalized to y. Let y^* be the socially efficient DM consumption, which solves $u'(y^*) = 1$. Households lack commitment and cannot enforce debt repayment. As a result, the DM trade must be quid pro quo and buyers must use a means of payment to exchange for y. We will discuss available means of payment later. The terms of trade are determined by buyers making take-it-or-leave-it offers. In the CM, both buyers and sellers work and consume x. Their labor h is transformed into x one-for-one. The utility from consumption is $U(x)$ with $U'(0) = \infty$, $U' > 0$, and $U'' < 0$.[7] Buyers' and sellers' preferences can be summarized respectively by the period utilities

$$\begin{aligned} \boldsymbol{U}^B(x, y, h) &= u(y) + U(x) - h, \\ \boldsymbol{U}^S(x, y, h) &= -y + U(x) - h. \end{aligned}$$

Young entrepreneurs are born in the current CM and become old and die in the next CM. Entrepreneurs cannot work in the CM and consume only when old. Young entrepreneurs are endowed with an investment opportunity that transforms x current CM goods to $f(x)$ CM goods in the next period, where $f'(0) = \infty$, $f'(\infty) = 0$, $f' > 0$, and $f'' < 0$. Entrepreneurs would like to borrow from households to invest. However, entrepreneurs and households lack commitment and cannot enforce debt repayment, so no credit arrangement among them is viable.

Like entrepreneurs, young bankers are born in the CM, become old and die in the next CM. Bankers cannot work in the CM and consume only when old.[8] Unlike households and entrepreneurs, bankers can commit to repay their liabilities and enforce the repayment of debt from entrepreneurs. Therefore, banks can act as intermediaries between households and entrepreneurs to finance investment projects. A bank can finance its loans by issuing two liabilities: liquid checkable deposits and illiquid time deposits.[9] Checkable deposits can be used as a medium of exchange to facilitate trading between buyers and sellers in the DM. Banks are subject to the reserve requirement that a bank's reserve holdings must cover at least a fraction $\chi \geq 0$ of its checkable deposits.

The government is a combination of monetary and fiscal authorities. The monetary authority,

[7] For the theoretical analysis, $U(x)$ can be simply linear. The more general functional form $U(x)$ allows us to introduce a parameter that affects total CM output to better match the ratio of M1 to GDP in the quantitative analysis.

[8] Infinitely-lived banks complicate expositions, but have little impact on the results. In this model, banks do not have incentives to retain profits for investment because deposit financing is cheaper. Therefore, they behave as if they live for one period.

[9] Time deposits are not important for our results but help to avoid a technical complication. See Online Appendix I for a detailed discussion.

or the central bank, issues three forms of liabilities: physical currency (or cash), central bank reserves, and a CBDC. Currency is a physical token, pays a zero interest rate, and can be used as a means of payment. The reserves are electronic balances that pay a net nominal interest rate $i_r \geq 0$; they can be held only by banks and cannot be used for retail payments. The CBDC is a digital token or electronic entry that can be used for retail payments. It pays a net nominal interest i_e. We focus on stationary monetary policies, where the total liabilities of the central bank (currency, CBDC, and reserves) grow at a constant gross rate $\mu > \beta$ and the central bank stands ready to exchange its three forms of liabilities at par in the CM. We abstract from government purchases. The government collects revenues from the issuance of new liabilities to pay interest on the CBDC and reserves, and the difference finances lump-sum transfers (T) to buyers (a negative T represents lump-sum taxes).

In the following, we describe how payments flow in the economy. In the DM, buyers use cash, CBDC, and checkable deposits to purchase good y from sellers. We assume that the two electronic payment methods, CBDC and deposits, are perfect substitutes in terms of payment functions. Sellers are distinguished in three types by the payment methods they accept (Lester et al. 2012; Zhu and Hendry 2019). Type 1 sellers (of measure $\omega_1 > 0$) accept only cash and can be interpreted as local cash-only stores that do not accept electronic payments. Type 2 sellers (of measure $\omega_2 > 0$) accept deposits and CBDC, and can be interpreted as online stores. Type 3 sellers (of measure $\omega_3 = 1 - \omega_1 - \omega_2 \geq 0$) accept all three payment methods and can be interpreted as local stores with point-of-sale machines that accept both cash and electronic payment methods.

Since the CM is Walrasian, the equilibrium allocation can be supported by different patterns of payment flows. Some possible types of transactions are as follows. Buyers trade x to rebalance their payment portfolio, spending the deposits issued by old banks, acquiring deposits issued by new banks, and adjusting their cash and CBDC balances. Sellers use their earnings in the previous DM (in cash, CBDC, and old deposits) to buy x. Young entrepreneurs acquire loans from young banks in the form of new deposits to purchase x for investment. During the process, new deposits are transferred from young entrepreneurs to buyers and old entrepreneurs. Old entrepreneurs sell x in exchange for new and old bank deposits to repay their loans from old banks. Between the banks, young banks use their newly issued deposits to exchange for reserves from old banks. Old banks use acquired new deposits to repay remaining liabilities and to purchase x. Note that, as in reality, banks engage in purely financial transactions, except when they use their profits to buy x. Figure 1 summarizes the activities and timeline for all private agents.

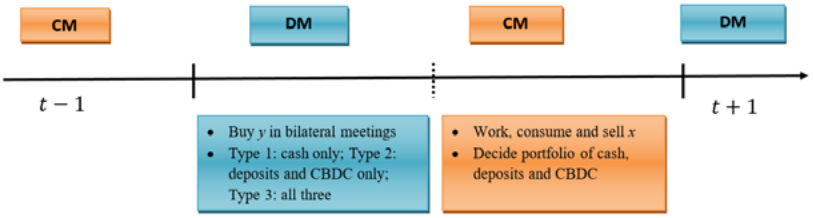

(a) Buyers

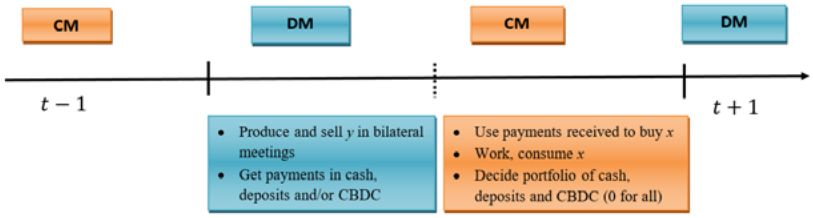

(b) Sellers

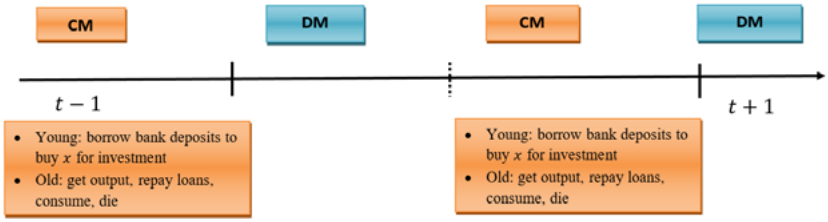

(c) Entrepreneurs

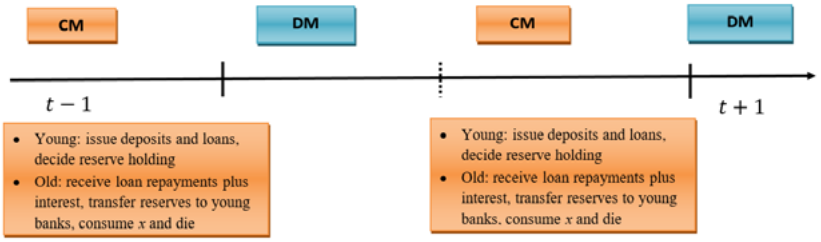

(d) Bankers

Figure 1: Timeline

In the benchmark model analyzed in the next section, we assume that banks cannot hold the CBDC. We also assume that N is fixed, banks engage in Cournot competition in the deposit market, and the lending market is perfectly competitive. This simple environment transparently illustrates our main mechanism. In Online Appendix B.2, we consider the case where banks can use the CBDC as reserves. In Section 4.4 and Online Appendices C and D, we study a model where N is endogenous and the lending market also features imperfect competition. Online Appendix F shows a model with price competition in the deposit market following Burdett and Judd (1983) and Head et al. (2012). Online Appendix G shows an extension of the model that incorporates risk-taking considerations. Our main findings are robust in all of these extensions.

3 Equilibrium Characterization

We focus on stationary monetary policies and stationary equilibria where real allocations are constant over time. It takes four steps to solve for the equilibrium. First, characterize the household's problem to derive the demand for cash, CBDC, and bank deposits as functions of the deposit rate. Second, solve the Cournot game for banks, incorporating the household demand for deposits, to derive the aggregate deposit supply and loan supply as functions of the competitive loan rate. Third, derive the aggregate demand for loans from entrepreneurs. Finally, equate the supply and demand for loans to derive the equilibrium loan rate and loan quantity and plug them into the solutions to private agents' problems to obtain other equilibrium objects, such as the rate and quantity of deposits.

3.1 Households

We first present the buyer's problem, and then the seller's problem. Let W and V be the household's value functions in the CM and DM, respectively. We suppress the time subscript and use prime to denote variables in the next period. Define $\vec{a} = (z, e, d, b)$ as the vector of the real value of cash, CBDC, checkable deposits, and time deposits held by an agent. Let $\vec{i} = (i_z, i_e, i_d, i_b)$ be the vector of net nominal returns, and $\vec{R} = (R_z, R_e, R_d, R_b) = (1 + \vec{i})/\mu$ be the vector of real gross returns. For example, the net nominal interest on cash is $i_z = 0$, and its real gross return is $R_z = 1/\mu$. For brevity, we often refer to R_e as the CBDC rate and R_d the (checkable) deposit rate.

In the CM, a buyer chooses consumption x, labor h, and the real asset portfolio $\vec{a}'$ carried to the next DM and measured at the current price. The value function for a buyer holding an asset portfolio $\vec{a}$ is

$$\begin{aligned} W^B(\vec{a}) &= \max_{x,h,\vec{a}'} \left\{ U(x) - h + \beta V^B(\vec{a}') \right\} \\ \text{subject to } x + \vec{1} \cdot \vec{a}' &= T + h + \vec{R} \cdot \vec{a}, \end{aligned}$$

where $\vec{1}$ is the unit vector $(1,1,1,1)$ and "$\cdot$" denotes the inner product of two vectors. The first-order condition with respect to asset portfolio $\vec{a}'$ is

$$\beta \frac{\partial}{\partial a} V^B(\vec{a}') \leq 1, \text{with equality if } a' > 0 \text{ for } a = z, e, d, b. \tag{1}$$

Note that, since the type of the DM meeting is not revealed until the start of the DM, buyers carry a portfolio of cash, CBDC, and bank deposits to the DM. Three standard results of the Lagos-Wright model are $U'(x) = 1$, all buyers choose the same portfolio $\vec{a}'$, and $\partial W^B(\vec{a})/\partial a = R_a$ for $a = z, e, d, b$.

The buyer's DM value function is

$$V^B(\vec{a}) = \sum_{j=1}^{3} \alpha_j [u(Y(\mathcal{L}_j)) - P(\mathcal{L}_j)] + W^B(\vec{a}), \tag{2}$$

where $\alpha_j = \omega_j \Omega$ is the (unconditional) probability of meeting a seller of type j, and $Y(\mathcal{L})$ and $P(\mathcal{L})$ are the terms of trade and represent the amount of good y being traded and the amount of payment, respectively. The terms of trade in a type j meeting depend on the buyer's usable liquidity $\mathcal{L}_j$, which incorporates the expected return of the asset. Specifically,

$$\mathcal{L}_1 = R_z z, \tag{3}$$

$$\mathcal{L}_2 = R_e e + R_d d, \tag{4}$$

$$\mathcal{L}_3 = R_z z + R_e e + R_d d. \tag{5}$$

Next we turn to the seller's problem. Without loss of generality, we assume that the seller does not take any asset into the DM, or $\vec{a}' = \vec{0}$.[10] Therefore, a type j seller's CM problem is

[10] It can be shown that, if the liquidity premium, defined below, on a liquid asset is positive, then the seller does not take that asset into the DM. The seller is indifferent between holding zero or a positive amount of illiquid time deposits when $R_b = 1/\beta$, which holds in equilibrium as shown below. For simplicity, we assume the seller does not hold time deposits either. Note that a seller enters the CM with positive asset balances ($\vec{a} > 0$) after trading in the previous DM.

$$\begin{aligned} W_j^S(\vec{a}) &= \max_{x,h}\{U(x) - h + \beta V_j^S(\vec{0})\} \\ \text{subject to } x &= h + \vec{R}\cdot\vec{a}. \end{aligned}$$

The type j seller's DM value function is

$$V_j^S(\vec{0}) = \Omega[-Y(\tilde{\mathcal{L}}_j) + P(\tilde{\mathcal{L}}_j)] + W^S(\vec{0}),$$

where $\tilde{\mathcal{L}}$ is usable liquidity held by the seller's trading partner.

The terms of trade in the DM are determined by buyers making take-it-or-leave-it offers and solve

$$\max_{y,p}[u(y) - p] \text{ subject to } p \geq y \text{ and } p \leq \mathcal{L},$$

where the first constraint is the seller's participation constraint and the second is the liquidity constraint. The solution is

$$Y(\mathcal{L}) = P(\mathcal{L}) = \min(y^*, \mathcal{L}). \tag{6}$$

In words, if the buyer has enough payment balances to purchase the optimal amount, then the optimal amount is traded; otherwise, the buyer's liquidity constraint binds and the buyer spends all available payment balances.

Combining (1) to (6), we can characterize the household's solution as follows. First, the demand for time deposits is separable from the demand for liquid assets and is given by $R_b = 1/\beta$. Since time deposits have no liquidity value, their return must compensate for discounting across time. Second, the buyer's demand for payment balances (z, e, d) is determined by

$$\frac{1}{\beta R_z} - 1 = \alpha_1\lambda(\mathcal{L}_1) + \alpha_3\lambda(\mathcal{L}_3), \tag{7}$$

$$\frac{1}{\beta R_a} - 1 \geq \alpha_2\lambda(\mathcal{L}_2) + \alpha_3\lambda(\mathcal{L}_3) \text{ with equality iff } a > 0, \text{ for } a = e, d, \tag{8}$$

where $\mathcal{L}_j$ is defined by (3) to (5), and $\lambda(\mathcal{L}) = \max\{u'(\mathcal{L}) - 1, 0\}$ is the liquidity premium.

Equation (7) states that the marginal cost of holding cash (left-hand side) equals its marginal benefit (right-hand side). The cost is that the buyer must delay consumption and bear the inflation cost to accumulate cash. The benefit is that more cash allows the buyer to consume more in type 1 and type 3 meetings. Equation (8) is for the CBDC and checkable deposits

and has a similar interpretation.

Under the assumptions $u'(0) = \infty$ and $\alpha_1 > 0$, the demand for cash is positive, so (7) holds as an equality. Similarly, if $\alpha_2 > 0$, the demand for total electronic (CBDC plus checkable deposits) balances is also positive. However, because the CBDC and checkable deposits are perfect substitutes, buyers hold only the instrument with the higher rate of return. From (8), if $R_d < R_e$, then the demand for checkable deposits is zero. If $R_d > R_e$, then the demand for the CBDC is zero. If $R_d = R_e$, then the buyer is indifferent between the CBDC and checkable deposits and cares only about the total electronic payment balances.

Equations (7) and (8) define R_d as a function of d, which is the inverse demand function for checkable deposits, denoted as $\mathbf{R}_d(d)$. To derive $\mathbf{R}_d(d)$, it is useful to first obtain the inverse deposit demand without a CBDC. We denote it as $\hat{\mathbf{R}}_d(d)$ (from now on, we will use the accent "ˆ" to denote variables or functions if there is no CBDC). We can solve $\hat{\mathbf{R}}_d(d)$ from (7) and

$$\frac{1}{\beta R_d} - 1 \quad = \quad \alpha_2 \lambda(\mathcal{L}_2) + \alpha_3 \lambda(\mathcal{L}_3), \tag{9}$$

after imposing $e = 0$. For certain values of d, there may exist multiple values of R_d that solve (7) and (9). This is because although (7) and (9) uniquely determine d given R_d, d may not be monotone in R_d. Intuitively, as R_d increases, there are two opposing effects: the substitution effect implies a higher d and the wealth effect implies a lower d. Throughout this paper, we assume that the substitution effect dominates and d is monotonically increasing in R_d.[11] Then, $\hat{\mathbf{R}}_d(d)$ is well-defined and increasing in d, with $\hat{\mathbf{R}}_d(0) = 0$ and $\hat{\mathbf{R}}_d(d) = 1/\beta$ for $d \geq \beta y^*$. With a CBDC, households hold only the electronic payment instrument that bears a higher rate of return. Therefore,

$$\mathbf{R}_d(d) = \begin{cases} [0, R_e) & \text{if } d = 0, \\ R_e & \text{if } d \in (0, \hat{\mathbf{R}}_d^{-1}(R_e)], \\ \hat{\mathbf{R}}_d(d) & \text{if } d > \hat{\mathbf{R}}_d^{-1}(R_e). \end{cases}$$

Figure 2 illustrates the inverse demand for checkable deposits. The solid line represents the demand with a CBDC, and the dashed line represents the demand without a CBDC. The two functions overlap if $R_d > R_e$. Once R_d is below R_e, the demand for checkable deposits drops to zero.

[11] Without this assumption, an equilibrium of the model still exists but may not be unique. A sufficient condition for this assumption is $-xu''(x)/u'(x) \leq 1$.

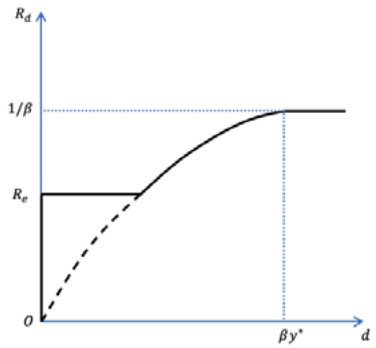

Figure 2: Inverse Demand for Checkable Deposits

Notes. The solid line is the inverse demand for checkable deposits with a CBDC, $\mathbf{R}_d(d)$; and the dashed line is the inverse demand for checkable deposits without a CBDC, $\hat{\mathbf{R}}_d(d)$. The two lines coincide with each other if $R_d > R_e$.

3.2 Banks

Banks issue two types of deposits, checkable deposits (d) and time deposits (b), and invest in two assets, reserves (r), and loans (ℓ). They do not invest in cash under the assumption $i_r \geq 0$. Bankers maximize consumption in the second period of life, which equals the return from loans and reserves, minus interest payments on deposits. They engage in Cournot competition in the deposit market and perfect competition in the loan market. Formally, banker j chooses $\{r_j, \ell_j, d_j, b_j\}$ to maximize its profit, taking as given the gross real rates for time deposits ($R_b = 1/\beta$), reserves (R_r) and loans (R_ℓ), the inverse demand function for checkable deposits ($\mathbf{R}_d(\cdot)$), and other banks' checkable deposit quantities ($D_{-j} = \sum_{i \neq j} d_i$):

$$\max_{r_j, \ell_j, d_j, b_j} \left\{ R_\ell \ell_j + R_r r_j - \mathbf{R}_d(D_{-j} + d_j) d_j - b_j/\beta \right\} \tag{10}$$
$$\text{subject to} \quad \ell_j + r_j = d_j + b_j, \quad r_j \geq \chi d_j.$$

This problem has two constraints. The first is a balance sheet identity at the end of the banker's first CM. The right-hand side is liabilities, which include checkable and time deposits. The left-hand side is assets, which include reserves and loans. The second is the reserve requirement constraint. We also implicitly impose that d_j, b_j, and ℓ_j are non-negative throughout the paper.

If $R_\ell > 1/\beta$, then the bank can make unlimited profits by issuing time deposits and investing

in loans. As a result, $R_\ell \leq 1/\beta$ in equilibrium. From now on, we restrict our attention to $R_\ell \in [0, 1/\beta]$. We also assume $R_r < 1/\beta$. We can separate the bank's problem into two steps. In the first step, the bank chooses funding sources (d_j, b_j):

$$\max_{d_j, b_j} \left\{ [\xi - \mathbf{R}_d(D_{-j} + d_j)]d_j + (\xi_b - 1/\beta)b_j \right\}, \tag{11}$$

where

$$\xi \equiv \max\{R_r, \chi R_r + (1-\chi)R_\ell\}$$

is the gross return on the bank's checkable deposits, and $\xi_b \equiv \max\{R_r, R_\ell\}$ is the gross return on time deposits. The first term in (11) is the profit from issuing checkable deposits, and the second term is the profit from issuing time deposits. Banks can hold their assets in loans or reserves, and therefore the return on assets is the higher of the two. Note that the return on checkable deposits accounts for the cost of satisfying the reserve requirement, while this consideration is absent for time deposits. Additionally, $b_j = 0$ if $R_\ell < 1/\beta$ and $b_j \in [0, \infty)$ if $R_\ell = 1/\beta$: the bank issues time deposits only if the return on loans is sufficient to cover the return of $1/\beta$ required by households.

In the second step, conditional on the choice in the first step, (d_j, b_j), the bank solves an asset allocation problem. If $R_\ell < 1/\beta$, then the bank issues only checkable deposits. It invests only in reserves if loans have a lower return than reserves, and invests only a fraction χ of assets in reserves to satisfy the reserve requirement if loans have a higher return. If the two assets have the same return, then the bank is indifferent between any allocations that satisfy the reserve requirement. If $R_\ell = 1/\beta$, then the bank starts to issue time deposits and ℓ_j can take any value in $[(1-\chi)d_j, \infty)$.

We focus on a symmetric pure strategy equilibrium in which every bank makes the same choice (r, ℓ, d, b). Denote the equilibrium checkable deposits of the Cournot game as $\mathbf{d}(R_\ell)$ to indicate its dependence on the loan rate R_ℓ. Following the discussion in the above paragraph, conditional on $\mathbf{d}(R_\ell)$, we can express the equilibrium loan supply function $\boldsymbol{\ell}(R_\ell)$ as

$$\boldsymbol{\ell}(R_\ell) = \begin{cases} 0 & \text{if} \quad R_\ell < R_r, \\ [0, (1-\chi)\mathbf{d}(R_\ell)] & \text{if} \quad R_\ell = R_r, \\ (1-\chi)\mathbf{d}(R_\ell) & \text{if} \quad R_r < R_\ell < 1/\beta, \\ [(1-\chi)\mathbf{d}(1/\beta), \infty) & \text{if} \quad R_\ell = 1/\beta. \end{cases} \tag{12}$$

To establish the existence and uniqueness of the equilibrium in the Cournot game, the

assumption below, Assumption 1, is maintained throughout the paper.[12] As discussed above, b_j is indeterminate if $R_\ell = 1/\beta$, and ℓ_j is indeterminate for certain values of R_ℓ. We say that the Cournot game has a unique symmetric equilibrium if the symmetric checkable deposit supply is unique.

Assumption 1 *a) For any $D \in [0, \beta y^*)$ and $\zeta \leq 1/\beta$, there exists a unique $d_j \in [0, \beta y^* - D)$ such that $\hat{\mathbf{R}}'_d(D+d)d + \hat{\mathbf{R}}_d(D+d) \lesseqgtr \zeta$ if $d \lesseqgtr d_j$ and $d \in [0, \beta y^* - D)$. b) $\hat{\mathbf{R}}'_d(Nd)d + \hat{\mathbf{R}}_d(Nd)$ increases with d on $[0, \beta y^*/N)$ and is less than R_r if d is sufficiently small.*

In the following, we first characterize the Cournot equilibrium if $R_e = 0$, which is equivalent to the case without a CBDC. It serves as a basis for analyzing the general case where $R_e > 0$. We characterize the Cournot equilibrium by taking the first-order condition of the bank's deposit-issuing problem (11) and imposing symmetry.

Proposition 1 *In the absence of a CBDC, the Cournot game has a generically unique symmetric pure strategy equilibrium, where each bank supplies $\hat{\mathbf{d}}(R_\ell) \in [0, \beta y^*/N)$ checkable deposits. In addition, $\hat{\mathbf{d}}(R_\ell)$ increases with R_ℓ and solves the following equation in d:*[13]

$$\hat{\mathbf{R}}'_d(Nd)d + \hat{\mathbf{R}}_d(Nd) = \xi. \tag{13}$$

Proof. *See Appendix A.* ■

In Figure 3, we plot the aggregate checkable deposit supply curve $\hat{\mathbf{D}}^s(R_\ell) = N\hat{\mathbf{d}}(R_\ell)$ (black curve in the left panel) and the loan supply curve $\hat{\mathbf{L}}^s(R_\ell) = N\hat{\boldsymbol{\ell}}(R_\ell)$ (black curve in the right panel) in the absence of a CBDC. Without a CBDC, banks always issue checkable deposits, and the loan supply is positive if $R_\ell \geq R_r$. If $R_\ell < R_r$, banks hold only reserves as assets, the checkable deposit supply is flat and the loan supply is zero. If $R_\ell = R_r$, the loan supply is vertical. Banks are indifferent between loans and reserves as long as the reserve requirement is satisfied. Both checkable deposits and loans strictly increase with R_ℓ if $R_r < R_\ell < 1/\beta$. If $R_\ell = 1/\beta$, then banks start to issue time deposits to finance loans. They are willing to supply any amount of loans that is no less than $(1-\chi)N\hat{\mathbf{d}}(1/\beta)$.

Now we analyze how a CBDC affects the checkable deposit and loan supply. Since the CBDC is a perfect substitute for checkable deposits regarding payment functions, it alters

[12] Part (a) of Assumption 1 guarantees that the symmetric Cournot equilibrium is generically unique. Part (b) guarantees that the equilibrium deposit supply is increasing in R_ℓ and banks issue checkable deposits for any R_ℓ. When $u(y) = y^{1-\sigma}/(1-\sigma)$, Assumption 1 holds if $\sigma < 1$ and R_z is not too small.

[13] The equilibrium is unique unless $R_\ell = 1/\beta$ and $\chi = 0$. In this case, there is one equilibrium where banks make positive profits, and a continuum of equilibria with $d \geq N\beta y^*/(N-1)$ in which banks make zero profits. We select the single positive-profit equilibrium.

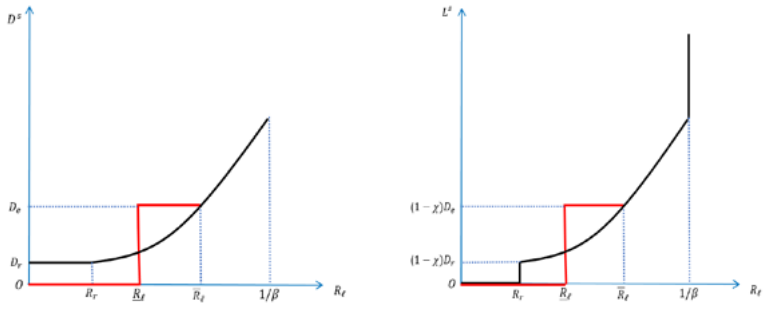

(a) Checkable Deposits (b) Loans

Figure 3: Effects of a CBDC on the Supply of Checkable Deposits and Loans

Notes. (1) $D_r = N\hat{\mathbf{d}}(R_r)$. (2) The red line is the case with a CBDC, and the black line represents the case without a CBDC. The two curves coincide with each other when $R_\ell \geq \bar{R}_\ell$.

the checkable deposit and loan supply only if the CBDC rate, R_e, exceeds the checkable deposit rate in the Cournot equilibrium without a CBDC, which is denoted by $\hat{\mathbf{R}}_d^*(R_\ell) \equiv \hat{\mathbf{R}}_d(N\hat{\mathbf{d}}(R_\ell))$. From (13) and Assumption 1, $\hat{\mathbf{R}}_d^*(R_\ell)$ is constant if $R_\ell \leq R_r$ and strictly increases in R_ℓ if $R_\ell > R_r$. Intuitively, under Cournot competition, a higher return on assets is partly passed on to the checkable deposit rate. Therefore, for a given CBDC rate R_e, the CBDC tends to alter the checkable deposit and loan supply only for low values of R_ℓ. In the following, we discuss in detail how a CBDC alters the checkable deposit and loan supply. To ease presentation, we focus on the case in which $R_e \in (R_r, \hat{\mathbf{R}}_d^*(1/\beta))$. Other cases are analyzed in Online Appendix B.1.

If $R_\ell \geq \bar{R}_\ell$, where $\bar{R}_\ell$ solves $\hat{\mathbf{R}}_d^*(\bar{R}_\ell) = R_e$, the CBDC rate is lower than the deposit rate in the Cournot equilibrium without a CBDC, and a CBDC does not affect the deposit and loan supply.[14] If $R_\ell < \underline{R}_\ell$, where $\underline{R}_\ell$ solves $(1-\chi)\underline{R}_\ell + \chi R_r = R_e$, then a bank's return on assets is insufficient to cover the cost of serving deposits, and it stops operating.

If $\underline{R}_\ell < R_\ell < \bar{R}_\ell$, then a bank matches the CBDC rate and supplies $d_e = D_e/N$ checkable deposits, where

$$D_e = \hat{\mathbf{R}}_d^{-1}(R_e).$$

Intuitively, if a bank reduces its supply of checkable deposits below d_e, then the checkable deposit rate remains equal to the CBDC rate, because the latter sets a floor for the former.

[14]The highest deposit rate in the Cournot equilibrium without a CBDC is $\hat{\mathbf{R}}_d^*(1/\beta)$. If $R_e < \hat{\mathbf{R}}_d^*(1/\beta))$, then $\bar{R}_\ell < 1/\beta$.

The deviating bank has a strictly lower profit because the marginal return of checkable deposits is higher than the marginal cost, that is, $(1-\chi)R_\ell+\chi R_r > R_e$. Therefore, no bank wants to reduce checkable deposits. On the other hand, no bank wants to increase checkable deposits, because that raises the deposit rate and lowers profits. Notice that a CBDC raises deposit quantity compared to the case without a CBDC, i.e., $d_e > \hat{\mathbf{d}}(R_\ell)$. Without a CBDC, banks restrict deposit supply and pay a deposit rate lower than R_e. With a CBDC, this is no longer possible, because R_e becomes a lower bound for the deposit rate. This reduces banks' incentives to restrict the deposit supply and leads to more deposits.

Finally, if $R_\ell = \underline{R}_\ell$, the bank is indifferent between operating and not, and the deposit supply lies in the interval $[0, d_e]$. Proposition 2 summarizes a bank's checkable deposit supply in the Cournot equilibrium with a CBDC.

Proposition 2 *If $R_e \in (R_r, \hat{\mathbf{R}}_d^*(1/\beta))$, a bank's supply of checkable deposits in the symmetric pure strategy equilibrium of the Cournot game is given by*

$$\mathbf{d}(R_\ell) = \begin{cases} 0 & \text{if} \quad R_\ell < \underline{R}_\ell, \\ [0, d_e] & \text{if} \quad R_\ell = \underline{R}_\ell, \\ d_e > \hat{\mathbf{d}}(R_\ell) & \text{if} \quad \underline{R}_\ell < R_\ell < \bar{R}_\ell, \\ \hat{\mathbf{d}}(R_\ell) & \text{if} \quad \bar{R}_\ell \leq R_\ell \leq 1/\beta. \end{cases} \tag{14}$$

Proof. *See Appendix A.* ■

Figure 3 illustrates how a CBDC affects the aggregate checkable deposit and loan supply, $\mathbf{D}^s(R_\ell) = N\mathbf{d}(R_\ell)$ and $\mathbf{L}^s(R_\ell) = N\boldsymbol{\ell}(R_\ell)$, graphed in red. If $R_\ell \geq \bar{R}_\ell$, then the deposit rate offered by banks in the absence of a CBDC is higher than the CBDC rate, and the CBDC does not affect the economy. Therefore, the checkable deposit and loan supply curves with and without a CBDC coincide. If $\underline{R}_\ell < R_\ell < \bar{R}_\ell$, then the supply of checkable deposits and loans is dictated by the CBDC rate R_e. The supply of checkable deposits stays at D_e and the supply of loans stays at $(1-\chi)D_e$. This corresponds to the positive horizontal part of the solid red line. In this interval, the red curve is above the black curve, reflecting that a CBDC can increase the deposit and loan supply. If $R_\ell = \underline{R}_\ell$, banks break even, and the supply of checkable deposits can take any value between zero and D_e, and the supply of loans lies between zero and $(1-\chi)D_e$. This corresponds to the vertical part of the solid red line. If $R_\ell < \underline{R}_\ell$, banks cannot compete with the CBDC and do not operate.[15]

[15] Online Appendix B.1 shows that the introduction of a CBDC with $R_e \in (\hat{\mathbf{R}}_d^*(R_r), (1-\chi)/\beta + \chi R_r)$ expands the supply of checkable deposits and loans for some values of R_ℓ (the second and/or third branch of (14) applies).

3.3 Entrepreneurs

Entrepreneurs take the gross loan rate R_ℓ as given and solve

$$\max_{\ell}\{f(\ell) - R_\ell \ell\}.$$

The inverse loan demand for an entrepreneur is $f'(\ell) = R_\ell$, which defines the aggregate loan demand function,

$$\mathbf{L}^d(R_\ell) = f'^{-1}(R_\ell).$$

The loan demand decreases with R_ℓ. It is always positive and approaches zero (infinity) as R_ℓ approaches infinity (zero). Note that the loan demand function is not affected by a CBDC.

3.4 Effects of a CBDC

We now combine the aggregate loan supply curve, $\mathbf{L}^s(R_\ell)$, with the aggregate loan demand curve, $\mathbf{L}^d(R_\ell)$, to determine the equilibrium loan quantity and rate. The loan market equilibrium is unique because the aggregate loan supply curve is non-decreasing. We can then use the equilibrium loan rate to derive the equilibrium quantity and rate of checkable deposits and loans. In the steady state equilibrium, the government budget constraint is $z + e + r = R_z z + R_e e + R_r r + R_z T$, which determines the equilibrium transfer to buyers. Note that the transfer T does not affect the household's demand for real payment balances, so it does not affect the analysis that determines the equilibrium rates and the quantities of deposits and loans.

Figure 4 shows the equilibrium with and without a CBDC. The blue curve is the aggregate loan demand and the black curve is the aggregate loan supply without a CBDC. They intersect at point a, which corresponds to the equilibrium without a CBDC. Let $\hat{R}^*_\ell$ and $\hat{R}^*_d$ be the rates of loan and checkable deposits, respectively, in this equilibrium.

The red curves illustrate the aggregate loan supply under different values of R_e. They intersect with the loan demand at point b, which corresponds to the equilibrium with a CBDC. We focus on the case in which $\hat{R}^*_\ell > R_r$.[16] If $R_e \leq \hat{R}^*_d$, a CBDC does not affect the equilibrium; otherwise, the effect of a CBDC can be distinguished by three regimes as R_e increases from $\hat{R}^*_d$. These regimes are distinguished along two dimensions: the effect of

[16] If the loan demand is sufficiently low, then the intersection a in Figure 4 would lie on the first vertical part of the loan supply curve and $\hat{R}^*_\ell = R_r$. The reserve requirement constraint would be slack and the equilibrium loan rate would be determined by the interest on reserves R_r. In this case, a CBDC can increase deposits without affecting lending, as in Andolfatto (2020).

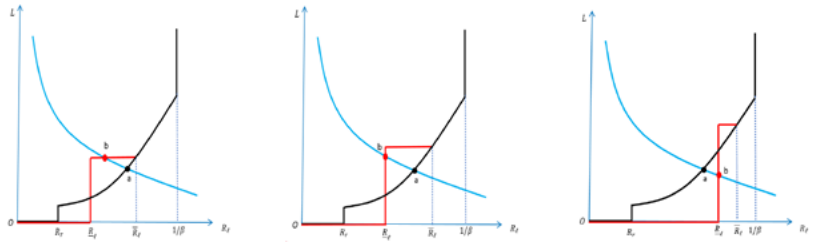

(a) Regime 1: $R_e \in (\hat{R}_d^*, R_{e1})$ (b) Regime 2: $R_e \in (R_{e1}, R_{e2})$ (c) Regime 3: $R_e \in (R_{e2}, 1/\beta)$

Figure 4: Effects of a CBDC

Notes. The blue curve is the aggregate loan demand, the black curve is the aggregate loan supply without a CBDC, and the red curve is the aggregate loan supply with a CBDC. The red curve joins the black curve for $R_\ell > \bar{R}_\ell$.

introducing a CBDC (relative to the case without a CBDC), and the comparative statics of deposits and loans with respect to R_e.

Regime 1 is shown in Figure 4(a). Compared with the case without a CBDC, a CBDC raises the deposit rate and the demand for electronic payment balances. If the CBDC had not been introduced, banks would have restricted their supply of checkable deposits and offered a lower deposit rate. However, the CBDC sets a floor for the rate of checkable deposits. Losing the ability to further reduce deposit rate, banks supply D_e checkable deposits to meet all the demand for electronic payment balances at the CBDC rate, because the marginal profit from checkable deposits is positive. In this regime, the CBDC is not used. A bank invests a fraction $1-\chi$ of its checkable deposits in loans, so the aggregate loan quantity is $L_e = (1-\chi)D_e$. Now we analyze how the economy responds as the CBDC rate increases. As R_e rises, $\underline{R}_\ell$ and $\bar{R}_\ell$ move to the right and the horizontal part of the loan supply curve rises, because D_e increases. Therefore, the rate and quantity of checkable deposits increase. The loan quantity increases and the loan rate decreases. Banks then have a lower profit margin because of the higher deposit rate and the lower loan rate. If $R_e = R_{e1}$, which solves $R_e = (1-\chi)f'(L_e) + \chi R_r$ as an equation in R_e, the profit margin reaches zero and all banks make a zero profit.

As R_e increases beyond R_{e1}, the economy enters into regime 2, illustrated in Figure 4(b). In this regime, a higher R_e increases the rates of checkable deposits and loans. The marginal profit from checkable deposits is zero and banks behave as if they are perfectly competitive. To stay break-even, banks must increase the loan rate to compensate for a higher deposit

rate. This lowers the equilibrium loan quantity. Banks then create fewer checkable deposits to finance loans. However, households increase their electronic payment balances by holding more CBDC. If $R_e < R_{e2} \equiv (1-\chi)\hat{R}_\ell^* + \chi R_r$ (or equivalently, $\underline{R}_\ell = \hat{R}_\ell^*$), a CBDC still leads to more loans and deposits relative to the case without it.

Finally, as R_e increases beyond R_{e2}, regime 3 occurs. It is the same as regime 2, except that the CBDC rate is too high and the quantities of checkable deposits and loans drop below the level without a CBDC. In other words, introducing the CBDC causes disintermediation if and only if $R_e > R_{e2}$. The following proposition summarizes these discussions.

Proposition 3 *There exists a unique steady-state monetary equilibrium with a CBDC. If $R_e \leq \hat{R}_d^*$, then a CBDC does not affect the economy. If $R_e > \hat{R}_d^*$, the effects of a CBDC are as follows.*

1. ***Effects of introducing a CBDC:*** *Introducing a CBDC promotes lending relative to the case without a CBDC if $R_e \in (\hat{R}_d^*, R_{e2})$, and reduces lending if $R_e \in [R_{e2}, 1/\beta]$.*

2. ***Comparative statics with respect to the CBDC rate:*** *Raising R_e promotes lending if $R_e \in (\hat{R}_d^*, R_{e1})$, and reduces lending if $R_e \in [R_{e1}, 1/\beta)$.*

Our analysis delivers two important messages. First, introducing a CBDC does not necessarily cause disintermediation or reduce bank loans and deposits. Indeed, the CBDC expands bank intermediation by introducing more competition to the banking sector if its rate falls between $\hat{R}_d^*$ and R_{e2}. Second, one should not judge the effectiveness of the CBDC based on its usage, but, rather, on how much it affects the deposit and lending rates or quantities. Throughout regime 1, the CBDC is not used, but it increases both deposits and loans. In fact, it maximizes lending at $R_e = R_{e1}$, which is the upper bound of regime 1. Here, the CBDC works as a *potential entrant*. It disciplines the off-equilibrium outcome: if banks reduce their deposit rates below the CBDC rate, then buyers would switch to the CBDC.

We end this section by comparing a CBDC and an interest-rate-floor policy that sets a minimum deposit rate that banks can legally offer. These two policies are similar but, in general, can deliver different outcomes. If the CBDC rate R_e is below R_{e1}, then the CBDC is not used and its effect is identical to a policy that mandates banks to pay a real rate no less than R_e. However, the two policies lead to different outcomes if R_e is larger than R_{e1}. With a CBDC, households use both the CBDC and deposits for transactions, and their total electronic balance increases with R_e despite shrinking deposits. With an interest-rate-floor policy, households end up with lower electronic payment balances and therefore consume less in type 2 and type 3 meetings. The effects of the interest-rate-floor policy (or the CBDC

with $R_e < R_{e1}$) are also related to the effects of a minimum wage on employment in the labor literature. It has been shown that a minimum wage can increase employment if firms have monopsony power (Burdett and Mortensen, 1998; Flinn, 2006; and Ashenfelter et al., 2010).

4 Quantitative Analysis

Theoretically, a CBDC can increase bank lending if its interest rate lies in a certain range. Empirical questions remain as to how large a range this is and how big the effect of a CBDC can be. To answer these questions, we calibrate our model without a CBDC to the United States economy, and conduct a counterfactual analysis to assess the effect of introducing a CBDC. We also use the calibrated model to study the effects of a non–interest-bearing CBDC when the economy trends toward cashless.

4.1 Calibration

We introduce two modifications to the model. First, we assume that banks incur a management cost c per unit of deposits. This simply adds the term $-c(d_j + b_j)$ to the profit function in (10). In our model, this cost is equivalent to a variable asset management cost. Second, we allow sellers in the DM to have some market power. Specifically, the DM terms of trade are determined by the Kalai bargaining, with bargaining power θ to the buyer. These modifications do not affect the qualitative results but capture two features in the data: banks have operational costs and sellers have substantial markups. Both features can be quantitatively important.

Consider an annual model and the functional forms $U(x) = B \log x$, $u(y) = [(y+\epsilon)^{1-\sigma} - \epsilon^{1-\sigma}]/(1-\sigma)$, and $f(k) = Ak^{\eta}$. The parameter ϵ is set to 0.001. This guarantees $u(0) = 0$ so that the Kalai bargaining is well-defined for all σ. It has little effect on our counterfactual analysis. There are 15 parameters to calibrate: $(A, B, N, \Omega, \omega_1, \omega_2, \omega_3, \sigma, c, i_r, \theta, \beta, \eta, \chi, \mu)$. Nine parameters, i_r, c, β, η, μ, χ, and ω_i $(i = 1, 2, 3)$, are set directly. The rest are calibrated internally. We calibrate $\omega_1, \omega_2, \omega_3, c, A, N, i_r, \chi$ and μ using data from 2014 to 2019. The calibration of $(\Omega, B, \sigma, \eta)$ follows the standard approach of matching the money and loan demand curves, which requires the use of longer time series data.

We use four data sets in our calibration exercise: (1) data from the Survey of Consumer Payment Choice (SCPC) and the Diary of Consumer Payment Choice (DCPC) from the Federal Reserve Bank of Atlanta; (2) call report data from the Federal Financial Institutions

Examination Council; (3) new M1 series from Lucas and Nicolini (2015); and (4) several time series on macro variables and reserves from Federal Reserve Economic Data (FRED). In what follows, we briefly discuss the calibration of several key parameters. For more details, see Online Appendix H.

We obtain the payment acceptance parameters, the ωs, from the SCPC (Greene and Stavins 2018) and the DCPC (Premo 2018). The SCPC contains information on the fraction of online transactions, and the DCPC contains information on the perceived fraction of point-of-sale transactions that do not accept cash or debit/credit cards. We use data from the 2016 wave, and the numbers are similar in 2015 and 2017. The SCPC documents that an average household makes 67.8 transactions per month. This includes 6.6 automatic bill payments, 5.9 online bill payments, and 4.7 online or electronic non-bill payments. We count these as online transactions and they represent 25.37% of all transactions. We assume that all online transactions accept only deposits. At the point of sale, the DCPC reports that 15% of transactions do not accept debit/credit cards and 2% of transactions do not accept cash. Then, cash-only transactions are those at points of sale that do not accept cards. This implies $\omega_1 = 15\%(1 - 25.37\%) = 11.19\%$. Deposit-only transactions include online transactions and point-of-sale transactions that do not accept cash. Hence, $\omega_2 = 25.37\% + 2\%(1 - 25.37\%) = 26.86\%$, and $\omega_3 = 1 - \omega_1 - \omega_2 = 61.94\%$.

Next, we calibrate the DM trading probability (Ω), the parameters of the utility functions (σ, B), and the bargaining power (θ) jointly to match the money demand curve and a 20% retail markup. For the monetary aggregate, we use the new M1 series from Lucas and Nicolini (2015), which include cash, checkable deposits, and some interest-bearing liquid accounts. To calculate the money demand in the model, we also need the deposit rates. The call report data, which was also used in Drechsler et al. (2017), contain the interest expenses and balances on the transaction accounts of the United States banks. We take the ratio of these two variables to obtain an average interest rate on transaction deposits. For this calibration, we use data from 1987 to 2008 for the following reasons. First, interest expenses on transaction accounts are not available before 1987. Second, after 2008, demand for M1 rises sharply likely due to non-transactional demand, such as store-of-value or foreign demand. Because our model focuses on the transactional demand, it is more reasonable to exclude the post-crisis data. In our calibration exercise, the ωs are set to their values in 2016. It is possible that the ωs changed during 1987–2008 and differ from their values in 2016. However, since we calibrate the parameters to match the aggregate money demand curve, our approach remains valid if changes in the ωs affect mainly the composition of M1 but not its aggregate. This is likely, because the money demand curve is stable during

Parameters	Notation	Value	Calibration Targets
Calibrated externally			
Discount factor	β	0.96	Standard in literature
Curvature of production	η	0.66	Elasticity of commercial loans
Reserve requirement	χ	5.60%	2014–19 avg. required reserves/trans. balances
Interest rate on reserves	i_r	1.02%	2014–19 avg. IORR
Cost of handling deposits	c	0.02	Avg. operating cost per dollar asset 2.02%
Gross money growth rate	μ	1.0152	2014–19 avg. annual inflation 1.52%
Frac. of type 1 trades	ω_1	11.19%	SCPC 2016
Frac. of type 2 trades	ω_2	26.86%	SCPC 2016
Frac. of type 3 trades	ω_3	61.94%	SCPC 2016
Calibrated internally			
Prob. of DM trading	Ω	0.22	Money demand 1987-2008
Coeff. on CM consumption	B	2.33	Money demand 1987-2008
Curv. of DM consumption	σ	1.66	Money demand 1987-2008
Total factor productivity	A	1.44	Rate on transaction accounts 0.3049%
Number of banks	N	19	Spread b/w transaction accounts and loans 3.39%
Buyer's bargaining power	θ	0.9988	Retailer markup 20%

Table 1: Calibration Results

1987–2008.

We set η, the parameter that governs the curvature of the entrepreneur's production function, to match the elasticity of commercial loans with respect to the prime rate, using the time series from FRED. Choose χ to match the average ratio of required reserves to the total balance in transaction accounts in 2014–2019. Set $i_r = 1.02\%$ to match the average interest rate on required reserves. To calibrate A, c, and N, we use several statistics on the banking sector in 2014-2019 calculated from the call report data. We choose c to match average non-interest expenditures, excluding expenditures on premises or rent, per dollar of assets. Given η, χ, i_r and c, pick A and N jointly to match the average interest rate on transaction accounts and the spread between the loan rate and the rate on transaction accounts.

Table 1 summarizes all the parameter values along with their calibration targets.[17] Figure 5(a) shows the model-predicted money demand curve against the data between 1987 and 2008. Figure 5(b) shows the loan supply and loan demand without a CBDC under the calibrated parameters.

[17]The money demand is rather flat during 1987-2008, so σ is larger than one. This result implies that the DM utility has a high curvature and a θ close to one is needed to match the markup. We have also done a calibration using data from 1987 and 2019. This alternative calibration suggests $\sigma = 0.45$ and $\theta = 0.80$, and the effect of a CBDC is larger than the benchmark calibration.

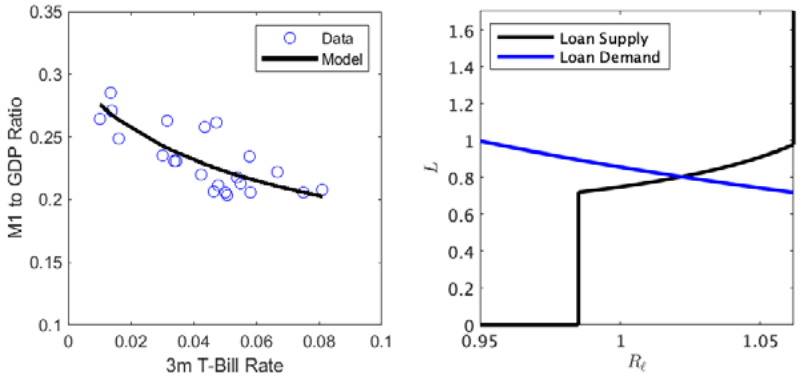

(a) Model Fit (b) Equilibrium

Figure 5: Money Demand and Equilibrium under Calibrated Parameters

4.2 Effects of an Interest-Bearing CBDC

Now we conduct counterfactual analysis and introduce an interest-bearing CBDC as a perfect substitute for checkable deposits. We are particularly interested in how the CBDC affects lending and output with different interest rates. Figure 6 shows the results. In all figures, the horizontal axis is the net nominal interest on CBDC i_e. The first row shows the net nominal deposit and loan interest rates and their difference, that is, the spread. All interest rates are in percentages. The second row displays the percentage changes of deposits, loans and total output relative to the equilibrium without a CBDC.

First note that if the interest rate on the CBDC (i_e) is below 0.30%, which is the deposit rate without a CBDC in our calibration, then the CBDC does not affect the economy; this corresponds to the flat parts of the figures.

Once i_e exceeds 0.30%, the CBDC rate becomes an effective floor of the deposit rate, and from that point on the deposit rate follows the 45^o line. As i_e and the deposit rate increase, total checkable deposits increase as long as banks make positive profits. At $i_e = 0.85\%$, a bank's profit becomes zero, and the checkable deposits reach their maximum. After that, a further increase in i_e leads to a reduction in checkable deposits. To break even with a higher deposit rate, banks must raise their lending rates, which reduces the loan demand and hence the deposit supply. The effect on loan quantity is the same as the effect on deposits. Because a higher loan supply reduces the loan rate, the loan rate first decreases and then increases

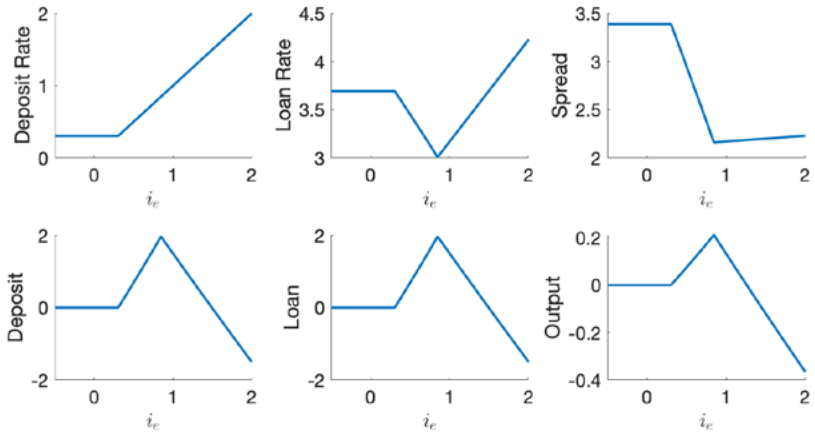

Figure 6: Effects of the Interest Rate on CBDC

with i_e. From Figure 6, if the CBDC rate is between 0.30% and 1.49%, then the CBDC increases both deposits and loans compared to the equilibrium without a CBDC. At the maximum, the CBDC increases checkable deposits and loans by 1.96% and reduces the loan rate to around 3.00% from about 3.70%.

We next focus on the spread. The CBDC competes with checkable deposits, and a higher CBDC rate reduces the spread as long as banks make positive profits. If i_e is sufficiently high, then banks earn zero profits and the spread starts to increase. Intuitively, as the CBDC rate increases, the interest rate on deposits increases. Because of the reserve requirement, a bank can lend only a fraction of its deposits. Therefore, the loan rate must increase even more to compensate for the increase in the deposit rate, explaining the increasing segment of the spread curve.

Lastly, we move to output.[18] The pattern is qualitatively similar to that of loans: as i_e in-

[18]The steady state output aggregates consumption and investment in the DM and the CM:

$$\mathbf{Y} = \sum_{j=1}^{3} \alpha_j P(y_j) + 2B + R_r(D - L) + AL^\eta - R_d D + L.$$

The first two terms are households consumption in the DM and CM, respectively. DM consumption is measured by real payments in terms of CM goods. The third, fourth and fifth terms together measure the CM consumption of old bankers and old entrepreneurs. It equals revenue from reserves (the third term) and production (the fourth term) subtracted by repayment of deposits plus interest (the fifth term). The last term is investment by young entrepreneurs. The net consumption of government is zero.

creases, the total output first increases and then decreases. Quantitatively, the expansionary effect on output is more modest relative to lending, because of the diminishing return in production. Introducing a CBDC increases the total output (relative to the case without a CBDC) if $i_e \in (0.30\%, 1.26\%)$. The highest increase in output is 0.21%, which is achieved at $i_e = 0.85\%$.

In addition to the baseline calibration, we have conducted some robustness checks. First, we extend the model to allow banks to hold the CBDC as reserves. The results are almost identical. Second, we assess the sensitivity of our results with respect to values of χ and i_r. We range χ from zero to 10% and i_r from zero to 1.02%. The results are very close in magnitude. For example, if $\chi = 0$ and $i_r = 1.02\%$, the CBDC increases deposits and loans if its rate is between 0.30% and 1.63%. It increases output if the rate is between 0.30% and 1.37%. At the maximum, deposits and loans increase by 2.13% and output increases by 0.24%.

4.3 Non–Interest-Bearing CBDC in a Cashless Economy

We have so far focused on an interest-bearing CBDC. However, central banks may be wary of paying interest on a CBDC, at least in initial stages of its introduction. If the CBDC does not pay interest, can it still have any effect on intermediation? This section assesses the effect of a zero-interest CBDC as the payment landscape evolves, captured by changes in the ωs.

In particular, we consider the trend of declining cash usage experienced in many countries. Several central banks consider this trend an important motivation for issuing a CBDC. We capture this trend by converting $\Delta\%$ of type 3 sellers to type 2 sellers; that is, the fraction of type 3 sellers changes to $\omega_3 - \Delta\% \times \omega_3$ and that of type 2 sellers changes to $\omega_2 + \Delta\% \times \omega_3$. One interpretation is that some brick-and-mortar sellers have closed their physical stores and sell exclusively online. Therefore, more stores accept only electronic payment methods. However, such a change can occur for other reasons. For example, some consumers (merchants) have recently stopped using (accepting) cash for fear of transmitting and contracting the COVID-19 virus. We evaluate how an economy with and without a CBDC differs as Δ increases.

The solid blue line in Figure 7 illustrates the results without a CBDC. As Δ increases, checkable deposits become a better payment instrument. Banks gain more market power and reduce the deposit rate. Buyers hold more deposits despite the reduced deposit rate, because deposits can be used in more transactions. Banks issue more deposits and make more loans, and the loan rate decreases. More loans lead to a higher output. The spread

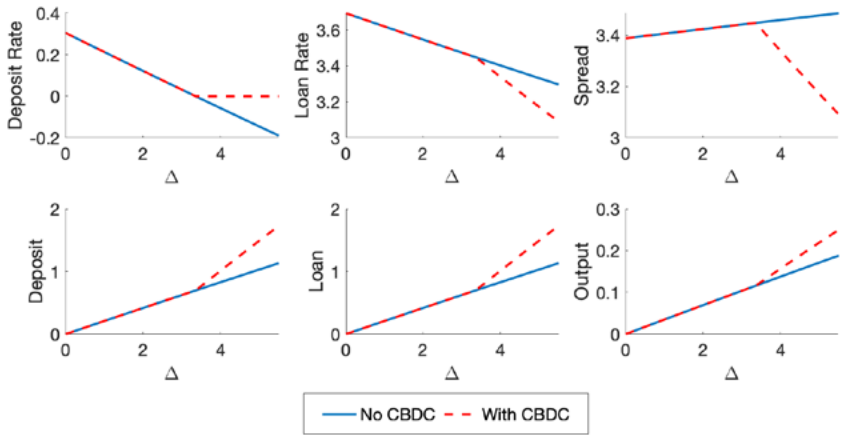

Figure 7: Effects of a Zero–Interest CBDC as Economy Becomes Cashless

Notes. The horizontal axis represents the % of type 3 meetings that change into type 2 meetings.

increases as the reduction in deposit rate exceeds the reduction in the loan rate.

The dashed red curve shows the economy with a zero-interest CBDC. If Δ is low, then the CBDC rate is lower than the deposit rate. Therefore, it does not affect the equilibrium and the dashed red curve overlaps with the solid blue curve. As Δ increases, the CBDC prevents the deposit rate from becoming negative, that is, zero becomes a hard floor of the deposit rate. As a result, banks find it optimal to create more deposits and make more loans. Compared to the case without a CBDC, the deposit rate and output are higher, while the loan rate and spread are lower. As Δ increases, banks gain more market power and the CBDC has larger effects. A zero-interest CBDC starts to affect the economy if 3.40% of type 3 sellers stop accepting cash. Therefore, the Unites States could reach a situation in which a zero-interest CBDC affects the economy with a modest change in the payment landscape.

4.4 Endogenous Bank Entry

We have so far taken the banks' market power as given and assumed that the number of banks is fixed. In this subsection we endogenize it by modelling bank entry through a fixed

operating cost.[19] We study how endogenous entry affects our quantitative results.

With endogenous entry, banks play a two-stage game. In the entry stage, they decide whether to incur the fixed operating cost, denoted by κ, to become active. In the second stage, the N active banks play the Cournot game analyzed in Section 3. We solve the model backward in two steps. First, we solve the N-bank Cournot game and obtain each active bank's profit $\pi(N)$, which decreases with N. Second, we solve the entry game. The equilibrium number of banks, denoted by N^*, satisfies

$$\pi(N^*) \geq \kappa \text{ and } \pi(N^*+1) < \kappa. \tag{15}$$

There exists a unique equilibrium with active banks ($N^* \geq 1$) if κ is sufficiently small.

It turns out that if the loan market is competitive, then endogenizing bank entry does not modify the basic results from the model with a fixed N. The effects of a CBDC on the rate and quantity of deposits and loans remain the same as in the benchmark model as long as R_e is not too big so that at least one bank is active. Intuitively, the banks' market power on the deposit side is disciplined by the CBDC rate, which fully determines the aggregate deposit supply. On the loan side, the market rate is competitive and independent of the number of active banks. As a result, the reduction in bank profits and the subsequent exit of banks do not affect either side. The remaining active banks satisfy the demand for electronic payment balances at the CBDC rate and lend up to the reserve requirement.

Now suppose banks have market power and engage in Cournot competition on the loan side too. In a two-sided Cournot game, an individual bank j considers its price impact on both the deposit and the loan markets and solves

$$\max_{\ell_j, r_j, d_j, b_j} \left\{ \mathbf{R}_\ell(L_{-j}+\ell_j)\ell_j - \mathbf{R}_d(D_{-j}+d_j)d_j + R_r r_j - b_j/\beta \right\}$$
$$\text{subject to } \ell_j + r_j = d_j + b_j \text{ and } \chi d_j \leq r_j,$$

where $\mathbf{R}_\ell(\cdot) = f'(\cdot)$ is the inverse loan demand function and L_{-j} is the aggregate loan supply of banks other than j. After we solve the N-bank two-sided Cournot equilibrium (see Online Appendix D for detailed analysis), we can compute the profit of each bank, $\pi(N)$, and solve

[19] There is indeed evidence that banks have significant fixed operating costs. According to the call report data, expenses on business premises and fixed assets were, on average, 5.9% of a bank's income between 1987 and 2010. Corbae and D'Erasmo (2020) estimate that the fixed cost for all banks in the United States is 0.77% (scaled by loans). Liu (2019) finds substantial fixed costs associated with complying with regulations. In addition, Online Appendix E shows that banks have decreasing average costs, which is consistent with fixed operating costs.

N^* from (15).

In the model with endogenous entry and Cournot competition on both sides, the impact of a CBDC on intermediation is more complicated. Since a CBDC reduces bank profits, it reduces the number of active banks. On the one hand, a CBDC can raise the competitiveness in the deposit market and lead to more deposits and loans. On the other hand, a smaller number of active banks lowers the competitiveness in the loan market and results in fewer deposits and loans. Therefore, the overall effect of a CBDC on intermediation is a quantitative question.

We calibrate the model with two-sided Cournot competition and endogenous entry. The calibration follows two steps. We first calibrate a model with two-sided Cournot competition and a fixed number of banks. From this calibration, we obtain the number of banks and compute $\pi(N)$, the profit of active banks. We then set $\kappa = \pi(N)$ and solve the model with endogenous entry to evaluate the effect of a CBDC.

We find that a CBDC can still increase bank intermediation and output, but the magnitude is about a quarter of that in the benchmark model. The CBDC increases deposits and loans if i_e is between 0.30% and 0.56%. It increases output if i_e is between 0.30% and 0.51%. The maximum increase in loans is 0.42% and in output is 4.53 basis points. The smaller effects are due to two reasons. First, banks have less market power in the deposit market in this calibration, because part of the market power is attributed to loans. Second, a reduction in the number of banks counteracts the effects of the CBDC in the deposit market.[20] If the CBDC does not pay interest, it can have a positive effect on bank lending and output if more than 4.00% of type 3 sellers stop accepting cash.

Before we conclude this section, we briefly discuss the welfare implication of a CBDC in our model. First, different types of agents are affected differently by the CBDC. As i_e increases, buyers bring more electronic payment balances. Based on the calibration, buyers are not constrained in type 3 meetings. More electronic payment balances increase consumption in type 2 meetings and do not affect consumption in other meetings. As a result, the welfare of buyers and type 2 sellers increases with i_e, and the welfare of type 1 and type 3 sellers is not affected. Entrepreneurs benefit from the CBDC if and only if it reduces the loan rate. Banks lose because the CBDC reduces their market power and hence profits. Second, the analysis of aggregate welfare is complex and has many components. As discussed in details by Keister and Sanches (2021), one needs to consider payment efficiency and investment efficiency. In

[20] If the loan market features Cournot competition and the number of banks is exogenous, a CBDC increases deposits and loans if $i_e \in (0.30\%, 0.78\%)$ and output if $i_e \in (0.30\%, 0.69\%)$. The maximum increase in loans is 0.78% and in output is 8.35 basis points. Therefore, bank exits significantly dampen the expansionary effect of a CBDC.

addition, regarding the latter, over-investment is possible due to cheap deposit funding, so a higher level of bank intermediation and output does not necessarily imply higher welfare.

5 Discussion

In this section, we discuss some important issues related to CBDCs, including the motivations for issuing a CBDC and the implementation scheme.

5.1 Motivations for Issuing a CBDC

According to the Bank for International Settlements survey on CBDC (Boar and Wehrli, 2021), domestic payments efficiency is a key motivation for issuing a retail CBDC in both advanced and emerging market economies. Related to this motivation, our findings suggest that a CBDC can discipline the bank's market power in providing transaction deposit balances and improves payment efficiency. In the past, cash has served this disciplinary role. As the economy enters an increasingly "cashless" digital world, the role of cash weakens, and a CBDC offers the general public an outside option for conducting electronic payments.

While introducing a CBDC could promote bank competition, one may ask why the payment market requires special treatment—after all, the government does not enter the supply side of each market that is subject to imperfect competition. In our view, a couple of features make the payment market somewhat special relative to others. Central banks are already and have been intervening in the payment market for a long time by providing payment balances in the form of cash. Given this historical involvement, the public sector has accumulated some tangible and intangible capital (e.g., payment infrastructures and social trust) that are valuable for its entry into the digital payments industry. Furthermore, the government's taxation power implies that the public sector has an advantage, relative to the private sector, in providing safe, liquid balances (Holmström and Tirole, 1998).[21]

Another motivation for issuing a CBDC is related to the implementation of monetary policy. It is commonly argued that, if interest rates are close to zero, monetary policy becomes less effective, because individuals and intermediaries can hold cash to avoid negative interest rates. In our model, cash is an outside option for households as a means of payment, and for banks as reserve balances. The existence of cash limits the ability of a central bank to reduce the interest on reserves and deposits. If cash is replaced with a CBDC, which can bear a

[21]The flight to safe government securities and bank notes in a financial crisis is evidence of the public sector's relative advantage in providing safe assets of which the nominal redemption value is certain.

negative nominal interest rate, then the limit on the interest rates on reserves and deposits can be relaxed. In practice, completely eliminating cash is unlikely in the near future, but whether cash is around or not, our model suggests that the interest on the CBDC becomes a new policy tool. Combining it with traditional monetary policy tools such as the interest rate on reserves, the central bank can implement a larger set of equilibrium allocations. It is straightforward to adapt our framework to discuss such issues; see, for example, Jiang and Zhu (2021).

Finally, there are other motivations, including safety and resiliency of the payment system, financial inclusion, monetary policy sovereignty and data privacy.[22] Regardless of the motivation, our paper helps central banks understand the potential impact of a CBDC on the financial industry.

5.2 CBDC Implementation

We now discuss some issues related to the implementation of a CBDC. In our model, the central bank sets the interest rate on the CBDC. An alternative arrangement is to control its quantity, with the interest rate on the CBDC and deposits endogenously determined by the market (Kumhof and Noone, 2018). Under this quantity rule, competition from the central bank will still induce private banks to raise the deposit rate. However, a key difference is that the CBDC always causes bank disintermediation. Specifically, the CBDC takes a positive market share away from banks, leading to fewer deposits and loans relative to the case without a CBDC.

Another implementation issue is related to the architecture of the CBDC system. Our results on bank intermediation require two conditions: the CBDC pays an interest set by the central bank, and the CBDC is a close substitute for bank deposits in terms of payment functionality. In addition, our theoretical analysis requires that the central bank can offer payment services as efficiently as banks. In practice, what architecture can meet these requirements without returning the market power to commercial banks? One option is that the central bank runs an independent CBDC system by itself. The central bank may be able to utilize new payment technologies to reduce the costs of operating an independent payment system. For example, the Riksbank's e-krona pilot considers running an independent payment system based on

[22]Related to the safety motivation, Chiu et al. (2020) suggest that private agents might not fully internalize the benefit of circulating safe payment balances. Andolfatto (2020) shows that a CBDC can expand deposit funding through greater financial inclusion and desired saving. Our result that a CBDC can expand intermediation suggests a similar impact. Also, Garratt and van Oordt (2021) argue that introducing a CBDC helps to promote privacy which is a public good.

the Distributed Ledger Technology. However, in such a system, it may still be challenging to provide comprehensive customer service and to satisfy anti-money-laundering/know-your-customer requirements.

Alternatively, the central bank can delegate payment and customer services to third parties and take measures to avoid returning market power to banks. One choice is to delegate the operation of the CBDC system to new non-bank players such as private fintech companies. For example, the Central Bank of the Bahamas has built the infrastructure, technology and regulatory framework for its CBDC (the Sand Dollar), and authorises non-banks as Sand Dollar agents to enrol customers. The central bank can also rely on traditional banks to provide payment and customer services and impose sufficient regulations. For example, the central bank could request banks to offer segregated CBDC accounts and impose the following regulations. Banks should not be allowed to charge excessive account fees or delay CBDC transactions, and should pay the CBDC interest in full (for further discussion, see Kahn et al., 2018).

6 Conclusion and Future Research

This paper develops a model with imperfect competition in the deposit market to analyze whether introducing a CBDC would cause disintermediation in the banking sector. We show that, contrary to the common wisdom, a CBDC can promote bank intermediation. Intuitively, if banks have market power, they restrict the deposit supply to lower the deposit rate. An interest-bearing CBDC introduces more competition, which leads to more deposits and lending, and a lower loan rate. However, greater intermediation arises only if the interest rate on the CBDC lies in some intermediate range. If the CBDC rate is too low, then the CBDC does not affect the equilibrium. If the CBDC rate is too high, disintermediation occurs.

Our model is useful for analyzing the effects of CBDCs with various design choices: interest bearing or not, cash like or deposit like, serving as reserves or not, with a fixed quantity or a fixed rate, and so forth. It can also be used to study the role of a CBDC in an increasingly cashless world and the interaction between CBDC-related policies and existing monetary policy instruments, such as interest on reserves.

Our model abstracts from the financial stability issue related to a CBDC. For instance, an interest-bearing CBDC would increase banks' funding costs. As a result of their lower charter value, on the asset side, banks could invest in riskier projects, increasing the total

risk in the financial system. On the liability side, banks could switch to less stable funding sources, such as wholesale funding, increasing the likelihood of runs in the wholesale market. Hence, introducing a CBDC could lead to financial stability concerns. Other policy tools (e.g., capital requirements and emergency lending facilities) could help limit banks' risk taking and alleviate the risk of runs in the banking system. Combining these policies with a well-designed CBDC could potentially promote bank intermediation without increasing risk in the economy. We take some initial steps to explore these issues in Online Appendix G, and leave a full analysis of this important issue for future research.[23]

Finally, our paper does not investigate central bank lending policy. As CBDCs become an alternative to deposit accounts, the central bank may also reconsider its lending policy.[24] Brunnermeier and Niepelt (2019) show that central bank lending to private banks can insulate private lending and investment against the reduction in bank deposits caused by a CBDC.[25] Our analysis suggests that a CBDC can promote bank intermediation without central bank lending. However, central bank lending can still serve as an additional (or complementary) policy tool to affect aggregate investment. For example, with endogenous bank entry and non-competitive loan market, the introduction of a CBDC weakens competition in the loan market, and central bank lending can counteract this negative effect. We leave a full analysis of how to coordinate central bank lending and CBDC policies to a separate paper.

References

[1] I. Agur, A. Ari, and G. Dell'Ariccia (2020) "Designing Central Bank Digital Currencies." *Journal of Monetary Economics*, forthcoming.

[2] D. Andolfatto (2020) "Assessing the Impact of Central Bank Digital Currency on Private Banks." *The Economic Journal*, available at https://doi.org/10.1093/ej/ueaa073

[23] In Online Appendix G, we incorporate the risk-taking channel based on Boyd and De Nicolo (2005), and show that our main result on the level of bank intermediation remains robust, while the introduction of a CBDC can induce more or less risk taking, depending on modeling assumptions.

[24] In the short run, large-scale, direct lending to firms seems difficult due to legal and political considerations, and operational challenges such as managing credit relationships, and screening and monitoring borrowers. Central bank lending through private banks seems to be a more viable option. However, technological developments (e.g., mobile banking) could reduce the operational costs and make direct lending more likely in the long run.

[25] Niepelt (2020) generalizes the equivalence result in Brunnermeier and Niepelt (2019) to a model with central bank reserves. Fernández-Villaverde et al. (2020) study the equivalence regarding maturity transformation between private and central bank intermedaition.

[3] O. C. Ashenfelter, H. Farber, and M. R Ransom (2010) "Labor Market Monopsony." *Journal of Labor Economics* 28(2), 203-210.

[4] J. Barrdear and M. Kumhof (2021) "The Macroeconomics of Central Bank Issued Digital Currencies." *Journal of Economic Dynamics and Control*, 104148,

[5] P. Benigno, L. Schilling, and H. Uhlig (2019) "Cryptocurrencies, Currency Competition and the Impossible Trinity." NBER Working Paper 26214.

[6] A. Berentsen, G. Camera, and C. Waller (2007) "Money, Credit and Banking." *Journal of Economic Theory*, 135(1), 171-195.

[7] C. Boar and A. Wehrli (2021) "Ready, Steady, Go? – Results of the Third BIS Survey on Central Bank Digital Currency."Bank for International Settlements Working Paper, NO. 114.

[8] J. Boyd and G. De Nicolo (2005) "The theory of bank risk taking and competition revisited." *The Journal of Finance* 60(3), 1329-1343.

[9] M. Brunnermeier and D. Niepelt (2019) "On the Equivalence between Private and Public Money." *Journal of Monetary Economics* 106(C), 27-41.

[10] K. Burdett and K. Judd (1983) "Equilibrium Price Dispersion." *Econometrica* 51(4), 955-969.

[11] K. Burdett and D. Mortensen (1998) "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39(2), 257-273.

[12] J. Chapman and C. Wilkins (2019) "Crypto "Money": Perspective of a Couple of Canadian Central Bankers." Bank of Canada Staff Discussion Paper 2019-1.

[13] J. Chiu and T. Koeppl (2019) "The Economics of Cryptocurrencies – Bitcoin and Beyond." Bank of Canada Staff Working Paper 2019-40.

[14] J. Chiu, M. Davoodalhosseini, J. Jiang and Y. Zhu (2020) "Safe Payments." Bank of Canada Staff Working Paper 2020-53.

[15] M. Choi and G. Rocheteau (2020) "Money Mining and Price Dynamics." *American Economic Journal: Macroeconomics* (forthcoming).

[16] Committee on Payments and Market Infrastructures (2018) "Central Bank Digital Currencies." Bank of International Settlements.

[17] D. Corbae and P. D'Erasmo (2020) "Capital Buffers in a Quantitative Model of Banking Industry Dynamics. " Mimeo.

[18] M. Davoodalhosseini (2021) "Central Bank Digital Currency and Monetary Policy," *Journal of Economic Dynamics and Control*, 104150

[19] M. Davoodalhosseini and F. Rivadeneyra (2020) "A Policy Framework for E-Money." *Canadian Public Policy* 46(1), 94-106.

[20] M. Davoodalhosseini, F. Rivadeneyra, and Y. Zhu (2020) "CBDC and Monetary Policy." Bank of Canada Staff Analytical Note 2020-4.

[21] M. Dong, S. Huangfu, H. Sun, and C. Zhou (2016). "A Macroeconomic Theory of Banking Oligopoly." *European Economic Review* 138, 103864.

[22] M. Dong and S. Xiao (2021). "Central Bank Digital Currency: A Corporate Finance Perspective." Mimeo.

[23] I. Dreschler, A. Savov, and P. Schnabl (2017) "The Deposit Channel of Monetary Policy." *Quarterly Journal of Economics* 132, 1819-1976.

[24] W. Engert and B. Fung (2017) "Central Bank Digital Currency: Motivations and Implications." Bank of Canada Staff Discussion Paper 2017-16.

[25] J. Fernández-Villaverde and D. Sanches (2019) "Can Currency Competition Work?" *Journal of Monetary Economics* 106(C), 1-15.

[26] J. Fernández-Villaverde, D. Sanches, L. Schilling, and H. Uhlig (2020a) "Central Bank Digital Currency: Central Banking for All?" *Review of Economic Dynamics* (in print).

[27] C.J. Flinn (2006) "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates." *Econometrica* 74, 1013-1062.

[28] B. Fung and H. Halaburda (2016) "Central Bank Digital Currencies: A Framework for Assessing Why and How." Bank of Canada Staff Discussion Paper 2016-22.

[29] R. Garratt and M. van Oordt (2021) "Privacy as a Public Good: A Case for Electronic Cash." *Journal of Political Economy* (forthcoming).

[30] C. Greene and J. Stavins (2018) "The 2016 and 2017 Surveys of Consumer Payment Choice: Summary Results." Federal Reserve Bank of Atlanta Discussion Paper No. 18-3.

[31] C. Gu, C. Monnet, E. Nosal, and R. Wright (2018) "On the Instability of Banking and Financial Intermediation." Working Papers 1901, Department of Economics, University of Missouri.

[32] A. Head, L. Q. Liu, G. Menzio and R. Wright (2012) "Sticky Prices: A New Monetarist Approach." *Journal of the European Economic Association*, 10: 939-973.

[33] B. Holmström, and J. Tirole (1998) "Private and public supply of liquidity." *Journal of political Economy*, 106.1: 1-40.

[34] J. Jiang and Y. Zhu (2021) "Monetary Policy Pass-Through with Central Bank Digital Currency." Bank of Canada Staff Working Paper (forthcoming).

[35] C. M. Kahn, F. Rivadeneyra, and R. Wong (2018) "Should the Central Bank Issue E-Money." Bank of Canada Staff Working Paper 2018-58.

[36] T. Keister and C. Monnet (2020) "Central Bank Digital Currency: Stability and Information." Mimeo.

[37] T. Keister and D. Sanches (2021) "Should Central Banks Issue Digital Currency?" *Review of Economic Studies*, forthcoming.

[38] M. Kumhof and C. Noone (2018) "Central Bank Digital Currencies—Design Principles and Balance Sheet Implications." Bank of England Staff Working Paper, No. 725.

[39] P. Kurlat (2019) "Deposit Spreads and the Welfare Cost of Inflation." *Journal of Monetary Economics* (106), 78-93.

[40] R. Lagos and R. Wright (2005) "A Unified Framework of Monetary Theory and Policy Analysis." *Journal of Political Economy* 113, 463-484.

[41] R. Lagos and S. Zhang (2019) "The Limits of *onetary* Economics: On Money as a Medium of Exchange in Near-Cashless Credit Economies." NBER Working Paper No. w25803.

[42] R. Lagos and S. Zhang (2021) "The Limits of *onetary* Economics: On Money as a Constraint on Market Power." *Econometrica*, forthcoming.

[43] B. Lester, A. Postlewaite and R. Wright (2012) "Information, Liquidity, Asset Price and Monetary Policy." *Review of Economic Studies* 79, 1209-1238.

[44] K. Liu (2019) "The Impact of the Dodd-Frank Act on Small U.S. Banks. " Mimeo.

[45] R.E. Lucas and J.P. Nicolini (2015) "On the stability of money demand. " *Journal of Monetary Economics* 73(C), 48-65.

[46] T. Mancini-Griffoli, M. S. Martinez, I. A. Peria, A. Ari, J. Kiff, A. Popescu, and C. Rochon (2018) "Casting Light on Central Bank Digital Currency." International Monetary Fund Staff Discussion Notes No. 18-08.

[47] J. Meaning, B. Dyson, J. Barker, and E. Clayton (2018) "Broadening Narrow Money: Monetary Policy with a Central Bank Digital Currency." Bank of England Staff Working Paper No. 724.

[48] C. Monnet, A. Petursdottir, and M. Rojas-Breu (2020). "Central Bank Account for All: Efficiency and Stability." Mimeo.

[49] Niepelt, D. (2020) "Monetary Policy with Reserves and CBDC: Optimality, Equivalence, and Politics." CEPR Discussion Paper DP15457.

[50] Payment Canada (2020) "COVID-19 Pandemic Dramatically Shifts Canadians' Spending Habits."Available at https://www.payments.ca/about-us/news/covid-19-pandemic-dramatically-shifts-canadians%E2%80%99-spending-habits.

[51] J. Premo (2018) "Guide to the 2016 Diary of Consumer Payment Choice."Federal Reserve Bank of Boston: Consumer Payments Research Center.

[52] G. Rocheteau, R. Wright and C. Zhang (2018) "Corporate Finance and Monetary Policy."*American Economic Review* 108, 1147-1186.

[53] F. Schar and A. Berentsen (2020) "Bitcoin, Blockchain, and Cryptoassets: A Comprehensive Introduction." MIT Press.

[54] D. Scharfstein and A. Sunderam (2016) "Market Power in Mortgage Lending and the Transmission of Monetary Policy." Unpublished working paper. Harvard University.

[55] L. Schilling, J. Fernández-Villaverde, and H. Uhlig (2020) "Central Bank Digital Currency: When Price and Bank Stability Collide." Available at SSRN: https://ssrn.com/abstract=3606226.

[56] L. Schilling and H. Uhlig (2019) "Some Simple Bitcoin Economics." *Journal of Monetary Economics* 106(C), 16-26.

[57] Y. Wang, T. Whited, Y. Wu, and K. Xiao (2020) "Bank Market Power and Monetary

Policy Transmission: Evidence from a Structural Estimation." forthcoming in *Journal of Finance*.

[58] Z. Wang (2020) "Tax Compliance, Payment Choice, and Central Bank Digital Currency." Mimeo.

[59] S. Williamson (2020*a*) "Central Bank Digital Currency: Welfare and Policy Implications. " Mimeo.

[60] S. Williamson (2020*b*) "Central Bank Digital Currency and Flight to Safety. " Mimeo.

[61] S. Zhou (2020) "Anonymity, Secondary Demand, and the Velocity of Cryptocurrency." Mimeo.

[62] Y. Zhu and S. Hendry (2019) "A Framework for Analyzing Monetary Policy in an Economy with E-Money." Bank of Canada Staff Working Paper. 2019-1.

Appendix

A Proofs

Proof of Proposition 1. The bank's choice of checkable deposits solves:

$$\max_{d_j} \left[\xi - \hat{\mathbf{R}}_d(D_{-j} + d_j)\right] d_j.$$

First, suppose $\xi < 1/\beta$, which occurs if $\chi > 1$ or $R_\ell < 1/\beta$. Focus the case where $D_{-j} + d_j < \beta y^*$ because banks make negative profit otherwise. By Assumption 1(a), this problem has a unique solution. It satisfies $\hat{\mathbf{R}}'_d(D_{-j} + d_j)\, d_j + \hat{\mathbf{R}}_d(D_{-j} + d_j) = \xi$. Then the symmetric pure strategy Nash equilibrium d must satisfy (13). Because $\hat{\mathbf{R}}'_d$ is positive and $\hat{\mathbf{R}}_d(\beta y^*) = 1/\beta > \xi$, $\hat{\mathbf{R}}'_d(Nd)\, d + \hat{\mathbf{R}}_d(Nd) > \xi$ if d is slightly smaller than $\beta y^*/N$. By Assumption 1, equation (13) has a unique solution, which is increasing in ξ and hence increasing in R_ℓ. Next, we show that there is a solution to (13) on $[0, \beta y^*)$ if $\xi = 1/\beta$. Let ξ_n be an increasing sequence that converges to $1/\beta$ and d_n be the solution to (13) if $\xi = \xi_n$. Then d_n is the Cournot equilibrium supply of checkable deposits if $\xi = \xi_n$. Let $\tilde{d} = \lim_n d_n \leq \beta y^*/N$. We show that $\tilde{d} < \beta y^*/N$ and therefore solves (13) under $\xi = 1/\beta$ by continuity. Suppose towards contradiction $\tilde{d} = \beta y^*/N$. Then a bank's profit under ξ_n is $[\xi_n - \hat{\mathbf{R}}_d(Nd_n)]d_n$, which converges to 0 because d_n converges to $\beta y^*/N$ and ξ_n converges to $1/\beta$. But if a bank unilaterally deviate to $d_n/2$, its profit is $[\xi_n - \hat{\mathbf{R}}((N - 1/2)d_n)]d_n$, which converges to $[1/\beta - \hat{\mathbf{R}}((N-1/2)\beta y^*/N)]y^*/2N > 0$. This implies that for n sufficiently large, a bank can choose $d_n/2$ and gets a higher profit. Therefore, d_n cannot be an equilibrium. This leads to a contradiction. As a result, $\tilde{d} < \beta y^*/N$ and solves (13) if $\xi = 1/\beta$. ■

Proof of Proposition 2. We only prove the third branch of (14), which says $\mathbf{d}(R_\ell) = d_e > \hat{\mathbf{d}}(R_\ell)$ if $\underline{R}_\ell < R_\ell < \bar{R}_\ell$. The other branches are obvious. First, if the total supply of checkable deposits D is lower than $Nd_e = D_e$, then increasing d_j does not change the real gross rate of deposits, which is fixed at R_e. The first-order derivative of (11) with respect to d_j is $\xi - R_e$, which is positive if $R_\ell > \underline{R}_\ell$ by the definition of $\underline{R}_\ell$. Therefore, bank j can always increase its profit by increasing d_j. Second, by the definition of $\bar{R}_\ell$, if $R_\ell = \bar{R}_\ell$, then $\hat{\mathbf{R}}_d(D_e) + \hat{\mathbf{R}}'_d(D_e)\frac{D_e}{N} = \xi$. Therefore, by Assumption 1, the marginal profit of a bank $\xi - \hat{\mathbf{R}}_d(D) - \hat{\mathbf{R}}'_d(D)\frac{D}{N} < 0$ for all $D > D_e$ and $R_\ell < \bar{R}_\ell$. It is profitable for a bank to reduce its supply of deposit if $D > D_e$. Combining both arguments, banks supply D_e checkable deposits in total and $\mathbf{d}(R_\ell) = d_e$ by symmetry. ■

Appendices for Online Publication

B Supplementary Analysis

B.1 Detailed Analysis of Checkable Deposit Supply

In the main text, the discussion assumes $R_e \in (R_r, \hat{\mathbf{R}}_d^*(1/\beta))$, where the checkable deposit curve experiences all four branches in (14). For R_e located out of this range, some branches in (14) disappear. We now describe $d(R_\ell)$ for all $R_e \in [0, 1/\beta]$ in more detail. One main take away is that for $R_e \in (\hat{\mathbf{R}}_d^*(R_r), (1-\chi)/\beta + \chi R_r)$, the introduction of a CBDC could expand the supply of checkable deposits and loans for some values of R_ℓ (the second and/or the third branch of (14) apply).

Case 1. If $R_e \leq \hat{\mathbf{R}}_d^*(R_r)$, then the CBDC rate cannot beat the lowest deposit rate offered by the bank regardless of the level of the lending rate, and the CBDC does not affect the deposit supply. Therefore, $\mathbf{d}(R_\ell) = \hat{\mathbf{d}}(R_\ell)$ for all $R_\ell \in [0, 1/\beta]$ and only the last branch in (14) remains.

Case 2. If $\hat{\mathbf{R}}_d^*(R_r) < R_e < R_r$, then $1/\beta > \bar{R}_\ell \geq R_r$. The bank makes a positive profit for all R_ℓ, and only the last two branches in (14) remain: $\mathbf{d}(R_\ell) = \hat{\mathbf{d}}(R_\ell)$ if $R_\ell \in [\bar{R}_\ell, 1/\beta]$ and $\mathbf{d}(R_\ell) = d_e$ if $R_\ell < \bar{R}_\ell$.

Case 3. If $R_e = R_r$, then $1/\beta > \bar{R}_\ell > R_r$ and $\underline{R}_\ell = R_r$. Note $\mathbf{d}(R_\ell)$ remains the same for all $R_\ell \leq R_r$ (the return of the bank's deposit is bounded below by R_r), $\mathbf{d}(R_\ell) = \mathbf{d}(\underline{R}_\ell) = [0, d_e]$ if $R_\ell \leq \underline{R}_\ell$. In this case, the bank can always break even (though may not make positive profits), and the last three branches in (14) remain: $\mathbf{d}(R_\ell) = \hat{\mathbf{d}}(R_\ell)$ if $R_\ell \in [\bar{R}_\ell, 1/\beta]$; $\mathbf{d}(R_\ell) = d_e$ if $\underline{R}_\ell < R_\ell < \bar{R}_\ell$; and $\mathbf{d}(R_\ell) = [0, d_e]$ if $R_\ell \leq \underline{R}_\ell$.

Case 4. If $R_r < R_e < \hat{\mathbf{R}}_d^*(1/\beta)$, then we are back to the case described in the main text and all four branches in (14) apply.

Case 5. If $\hat{\mathbf{R}}_d^*(1/\beta) \leq R_e < (1-\chi)/\beta + \chi R_r$, then the CBDC rate exceeds the highest deposit rate offered by the bank (without a CBDC) and the CBDC affects the economy for all values of $R_\ell \in [0, 1/\beta]$. At the same time, we have $1/\beta > \underline{R}_\ell > R_r$. The first three branches in (14) remain: $\mathbf{d}(R_\ell) = 0$ if $R_\ell < \underline{R}_\ell$; $\mathbf{d}(R_\ell) = d_e$ if $R_\ell \in (\underline{R}_\ell, 1/\beta]$; and $\mathbf{d}(R_\ell) = [0, d_e]$ if $R_\ell = \underline{R}_\ell$.

Case 6. If $R_e = (1-\chi)/\beta + \chi R_r$, then the CBDC affects the Cournot equilibrium for all values of $R_\ell \in [0, 1/\beta]$, and $\underline{R}_\ell = 1/\beta$ (or the bank cannot make positive profits and can

break even only when $R_\ell = 1/\beta$). Only the first two branches of (14) remain.

Case 7. If $R_e > (1-\chi)/\beta + \chi R_r$, then the required rate for checkable deposits is higher than the highest possible return on the bank's assets. As a result, $\mathbf{d}(R_\ell) = 0$ for all $R_\ell \in (0, 1/\beta]$. In this case, only the first branch in (14) remains. If $R_\ell = 1/\beta$, the bank could still offer time deposits. This case occurs only if $\chi > 0$.

B.2 CBDC as Reserves

Now we allow banks to hold the CBDC to satisfy the reserve requirement.[26] The CBDC then plays two roles. First, it is a means of payment that competes with checkable deposits. Second, it can lower the cost for the banks to hold reserves if it has a higher return than reserves.

The household's and entrepreneur's problems remain the same as in the main text. The bank's problem changes to

$$\max_{b_j, e_j, r_j, \ell_j, d_j} \left\{ R_\ell \ell_j + R_r r_j + R_e e_j - \mathbf{R}_d(D_{-j} + d_j) d_j - b_j/\beta \right\}$$
$$\text{s.t.} \quad \ell_j + e_j + r_j = b_j + d_j, \; e_j + r_j \geq \chi d_j,$$

where e_j is bank j's CBDC balance. As before, we solve the Cournot game among banks for each value of R_ℓ and trace out the aggregate loan supply. The only difference is that now banks hold the CBDC to satisfy reserve requirement if its rate is higher than the rate on reserves.

The red curves in Figure 8 illustrate the resulting aggregate loan supply curve. Again, we focus on the case with $R_e > R_r$. Same as in Figure 4, the aggregate loan supply curve with the CBDC coincides with the horizontal axis if R_ℓ is low. If R_ℓ is intermediate, the curve is flat. The deposit rate matches the CBDC rate, and the quantity of loans is fully determined by the CBDC rate and equals D_e. If R_ℓ is above a cut-off $\bar{R}_\ell^R$, the return of the checkable deposits is higher than that of the CBDC, and the aggregate loan supply curve is upward sloping.

[26] If banks can hold the CBDC but not as reserves, then the effect of a CBDC remains the same as in the benchmark model and banks do not hold any CBDC. The intuition is as follows. If $R_e \leq R_r$, then banks prefer reserves to the CBDC. If $R_e > R_r$, then the marginal benefit of investing in the CBDC is negative. To invest in the CBDC, banks need to issue deposits with at least the same return as the CBDC. In addition, banks must hold reserves and invest only a fraction of deposits on the CBDC, which implies that the total cost of holding the CBDC exceeds the return.

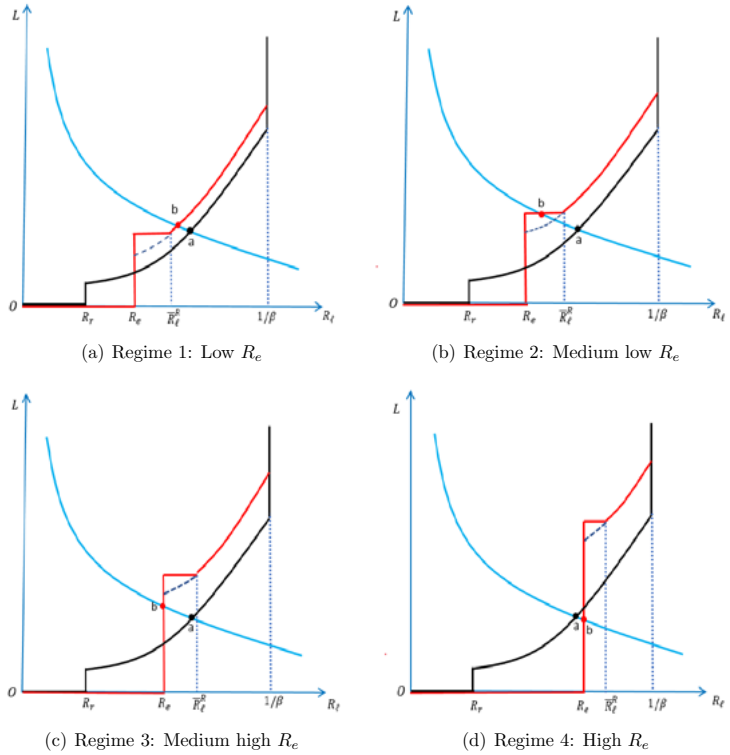

(a) Regime 1: Low R_e

(b) Regime 2: Medium low R_e

(c) Regime 3: Medium high R_e

(d) Regime 4: High R_e

Figure 8: CBDC as Reserves

Notes. (1) The blue curve is the loan demand, the black curve is the loan supply without a CBDC, the dashed black line is the loan supply when a CBDC is used as reserves but cannot be used for payments, and the red curve is the final new loan supply with a CBDC that can be used for both reserves and payments. The dashed line coincides with the red curve except for $R_\ell \in \left(R_e, \bar{R}_\ell^R\right]$. All three curves join each other at $R_\ell = 1/\beta$.

Compared with the baseline design, besides being an alternative payment method (*payment competition effect*), the CBDC can also reduce the bank's cost to hold reserves (*cost-saving effect*). The two effects together shift the loan supply curve without a CBDC, shown by the solid black curve, upward to the one with the CBDC. To decompose these two effects, we also plot the aggregate loan supply curve in an auxiliary model, where we shut down the payment competition effect. It is the dashed blue curve on $(R_e, \bar{R}_\ell^R)$ but overlaps with the red curve otherwise. The shift from the solid black curve to the dashed curve captures the cost-saving effect, and the shift from the dashed curve to the red curve captures the payment competition effect.[27]

We can analyze Figure 8 in the same way as Figure 4. It is easy to see that a CBDC that serves as reserves can also promote bank intermediation. Compared with the benchmark model, the CBDC has the additional cost-saving effect. This effect can be active and increase lending even if the CBDC has a lower rate of return than checkable deposits, as shown in Figure 8(a). This happens if the rate on reserves is low. Therefore, this design can increase bank intermediation more compared to the benchmark design. Moreover, it can promote bank intermediation for a wider range of i_e than in the benchmark model.

C Imperfectly Competitive Loan Market

This section considers the setup where banks engage in Cournot competition in both deposit and loan markets. We show that the equilibrium is the same as in a setup where banks have deposit and loan departments and these departments allocate funds among themselves in a perfectly competitive market. Therefore, we only need to analyze the equilibrium in the latter setup. The latter setup is attractive because we can use the loan supply and loan demand curves in the interbank market to study the equilibrium. Moreover, the loan supply curve is similar to the one obtained in the main text. A similar analysis shows that the CBDC can increase both deposits and loans. Throughout, we impose the following assumption.

Assumption 2 *The production function satisfies* $2f''(L) + f'''(L)L < 0$.

[27] In the increasing part of new (red) loan supply curve ($R_\ell \in [\bar{R}_\ell^R, 1/\beta)$), the payment competition effect is muted, and the higher loan supply relative to the old (black) supply curve is due to the cost-saving effect. In contrast, with the baseline CBDC design, the cost-saving effect is absent and the loan supply curves with and without the CBDC coincide if $R_\ell \in (\bar{R}_\ell, 1/\beta)$).

First, we consider the original setup without the interbank market. Bank j solves

$$\max_{\ell_j, r_j, d_j, b_j} \left\{ \mathbf{R}_\ell(L_{-j} + \ell_j)\ell_j - \mathbf{R}_d(D_{-j} + d_j)d_j + R_r r_j - b_j/\beta \right\} \tag{16}$$
$$\text{st } \ell_j + r_j = b_j + d_j \text{ and } \chi d_j \leq r_j,$$

where $\mathbf{R}_\ell(\cdot) = f'(\cdot)$ is the inverse loan demand function and L_j is the aggregate loan supply of banks other than j. Now bank j internalizes its impact on the loan rate. To solve (16), form the Lagrangian

$$\max_{\ell_j, r_j, d_j, b_j} \Big\{ \mathbf{R}_\ell(L_{-j} + \ell_j)\ell_j - \mathbf{R}_d(D_{-j} + d_j)d_j + R_r r_j - b_j/\beta$$
$$+ \Lambda_1(b_j + d_j - \ell_j - r_j) + \Lambda_2(r_j - \chi d_j) \Big\}.$$

One can obtain the first-order conditions and then impose symmetry to obtain the following set of equilibrium conditions.

$$0 = \mathbf{R}'_\ell(L)\frac{L}{N} + \mathbf{R}_\ell(L) - \Lambda_1, \tag{17}$$
$$0 \geq [(1-\chi)\Lambda_1 + \chi R_r] - \frac{\partial^+}{\partial D}\mathbf{R}_d(D)\frac{D}{N} - \mathbf{R}_d(D), \tag{18}$$
$$0 \leq [(1-\chi)\Lambda_1 + \chi R_r] - \frac{\partial^-}{\partial D}\mathbf{R}_d(D)\frac{D}{N} - \mathbf{R}_d(D), \tag{19}$$
$$0 = R_r - \Lambda_1 + \Lambda_2, \tag{20}$$
$$0 = \Lambda_2 \text{ if } L < (1-\chi)D \text{ and } 0 < \Lambda_2 \text{ if } L = (1-\chi)D, \tag{21}$$
$$1 = \beta\Lambda_1 \text{ if } b_j > 0 \text{ and } 1 > \beta\Lambda_1 \text{ if } b_j = 0, \tag{22}$$

where $\frac{\partial^+}{\partial D}\mathbf{R}_d(D)$ and $\frac{\partial^-}{\partial D}\mathbf{R}_d(D)$ denote the right and the left derivatives of $\mathbf{R}_d$, respectively. They are introduced because $\mathbf{R}_d(D)$ has a kink at $D = D_e$. The above conditions determine the equilibrium with Cournot competition in both the deposit and loan markets. Assumption 2 and the assumptions on $\mathbf{R}_d$ imply that the above conditions are necessary and sufficient, and the equilibrium is unique.

Now consider the setup with an interbank market. Each bank has a deposit department and a loan department. Each of the department cares only about their own profit. The deposit departments create deposits and lend them in a competitive interbank market. The loan departments borrow deposits in the interbank market and lend to entreprenuers. The equilibrium interbank market rate R_I equates the demand and the supply of funds in the interbank market. Deposit departments engage in Cournot competition with each other in

deposit creation while the loan departments engage in Cournot competition with each other in lending.

Then bank j's deposit department solves

$$\max_{\ell_j, r_j, d_j, b_j} \left\{ R_I \ell_j - \mathbf{R}_d(D_{-j} + d_j) d_j + R_r r_j - b_j/\beta \right\} \tag{23}$$
$$\text{st } \ell_j + r_j = b_j + d_j \text{ and } \chi d_j \leq r_j,$$

which is the same as the problem in our benchmark model except that now the lending rate in the interbank market appears in the problem.

The Cournot equilibrium among the deposit departments determines the aggregate loan supply in the interbank market. Without loss of generality, assume that $R_I \geq R_r$. Then we can form the Lagrangian

$$\max_{\ell_j, r_j, d_j, b_j} \left\{ R_I \ell_j - \mathbf{R}_d(D_{-j} + d_j) d_j + R_r r_j - b_j/\beta + \Lambda_1(b_j + d_j - \ell_j - r_j) + \Lambda_2(r_j - \chi d_j) \right\}.$$

Again, take the first-order conditions and impose symmetry to obtain

$$0 \geq [(1-\chi)R_I + \chi R_r] - \frac{\partial^+}{\partial D}\mathbf{R}_d(D)\frac{D}{N} - \mathbf{R}_d(D), \tag{24}$$
$$0 \leq [(1-\chi)R_I + \chi R_r] - \frac{\partial^-}{\partial D}\mathbf{R}_d(D)\frac{D}{N} - \mathbf{R}_d(D), \tag{25}$$
$$0 = R_r - R_I + \Lambda_2, \tag{26}$$
$$0 = \Lambda_2 \text{ if } L^s < (1-\chi)D \text{ and } 0 < \Lambda_2 \text{ if } L^s = (1-\chi)D, \tag{27}$$
$$1 = \beta R_I \text{ if } b_j > 0 \text{ and } 1 > \beta R_I \text{ if } b_j = 0, \tag{28}$$

where $L^s = \sum_{j=1}^{N} \ell_j$ is the aggregate loan supply. Notice that the loan supply curve obtained here is identical to the one in the benchmark model except it depends on R_I instead of R_ℓ.

The lending department j solves

$$\max_{\ell_j} \left\{ \mathbf{R}_\ell(L_{-j} + \ell_j)\ell_j - R_I \ell_j \right\}. \tag{29}$$

From the Cournot game among the loan departments, we can obtain the aggregate loan

demand in the interbank market. Under symmetry, the loan demand L^d satisfies

$$\mathbf{R}'_\ell(L^d)\frac{L^d}{N} + \mathbf{R}_\ell(L^d) = R_I. \tag{30}$$

Notice that if we combine (24)-(30) and use the market clearing condition in the interbank market, $L^d = L^s = L$, we obtain the same set of equilibrium conditions as in (17)-(22) with $R_I = \Lambda_1$. In the original problem, Λ_1 is the shadow value of loanable funds. It is equal to the interbank market rate R_I if there is a perfectly competitive interbank market. This proves that the model with and without a perfect competitive interbank market yield the same equilibrium.[28]

D Fixed Operating Costs and Endogenous N

Assuming that the the number of banks is fixed in the benchmark model, we found that a CBDC can increase bank lending by reducing their market power. However, if banks have fixed operating costs, a CBDC could cause banks to exit because it reduces the profit of operating banks. This in turn can strengthen the market power of remaining banks and offsets the positive effect of the CBDC. In this section, we show theoretically that taking this into account, the CBDC can still improve bank intermediation. In the following, we first study the case with a perfectly competitive loan market and then the case with a Cournot loan market.

Each period, $\bar{N}$ potential bankers decide whether to be active. If a potential banker decides to be active, the banker needs to pay a fixed operating cost κ. After the decisions are made, active bankers engage in Cournot competition in the deposit market and perfect competition in the loan market. This model can be solved in two steps. We first solve our benchmark model for different values of N and calculate the profit of a bank without the fixed operating cost, denoted by $\pi(N)$. Then the equilibrium number of banks N^* satisfies $\pi(N^*) \geq \kappa$ and $\pi(N^*+1) < \kappa$. We can again use the loan supply and demand curve to show that the CBDC can increase bank lending. The only difference is that we need to take into account the effect on the number of banks.

Figure 9(a) shows the equilibria with and without a CBDC. Again, the blue curve is the loan demand curve and the black curve is the loan supply curve if $N = \hat{N}^*$, where $\hat{N}^*$ is the equilibrium number of banks without a CBDC and point a corresponds to the equilibrium.

[28] If banks are heterogeneous, the equilibrium with and without the interbank market can differ.

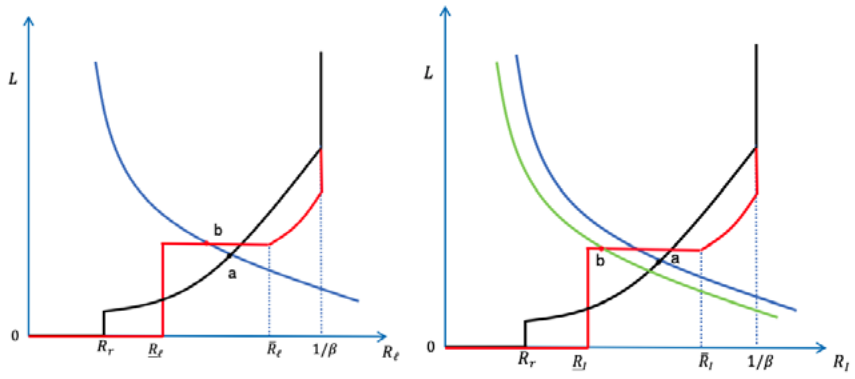

(a) Cournot in Deposit Market (b) Cournot in Deposit and Loan Markets

Figure 9: Fixed Operating Costs

Now consider introducing a CBDC with R_e higher than the gross real deposit rate without a CBDC but not too high. At the new equilibrium, $N^* < \hat{N}^*$ because the CBDC reduces bank profits. The red curve plot the loan supply if $N = N^*$. The same as in the benchmark model, the loan supply is 0 if $R_\ell < \underline{R}_\ell$. If $R_\ell \in (\underline{R}_\ell, \bar{R}_\ell)$, the aggregate loan supply is $(1-\chi)D_e$, which is also the same as in the benchmark model. The intuition is also the same: in this region, banks cannot reduce the deposit rate by reducing the supply because the CBDC rate is an effective lower bound. As they are making positive profits per unit of deposits, they satisfy all the demand for electronic payment balances. As R_ℓ moves beyond $\bar{R}_\ell$, the CBDC is not effective and the loan supply equals to the one without the CBDC under $N = N^*$. Notice that $\bar{R}_\ell$ is higher than its value if the number of active banks is fixed at $\hat{N}^*$. Because there are fewer banks and more market power, a higher lending rate is needed to make a CBDC ineffective. Moreover, the red curve lies below the black curve if $\bar{R}_\ell < R_\ell < 1/\beta$, because the aggregate loan supply decreases with N if the CBDC is not effective.

The equilibrium with the CBDC is point b. It features more deposits, loans and lower lending rate compared to the equilibrium without the CBDC. Moreover, it coincides with the equilibrium with the number of banks fixed at N^*. Intuitively, the CBDC sets the deposit rate and force active banks to compete, which completely off-sets the effect of fewer active banks. Indeed, as long as there is one active bank, the equilibrium remains at point b.

The above result depends crucially on the assumption of perfect competition in the loan

market. If the banks also have market power in the loan market, they would reduce aggregate loan supply if fewer banks are active. As shown in Appendix C, the equilibrium with a fixed N is the same as the equilibrium in a model where each bank is split into a deposit and a loan department and they trade in a competitive interbank market. Therefore, we can focus on the latter to simplify the analysis. In particular, we can plot the loan supply and loan demand in the interbank market to obtain the equilibrium just as in the previous analysis.

Figure 9(b) shows the effect of a CBDC if the loan market is imperfectly competitive. Same as the above, suppose that $N = \hat{N}^*$ in the equilibrium without the CBDC. The black curve is the loan supply in the interbank market if $N = \hat{N}^*$. The blue curve is the loan demand curve. Their intersection, point a, is the equilibrium without the CBDC. With the CBDC, N decreases to N^*. The the red curve is the loan supply in the interbank market with the CBDC if $N = N^*$. The loan supply is the same as in the case with a perfectly competitive loan market, except that it depends on the lending rate in the interbank market, R_I. But there is an additional effect on the loan demand. Because fewer banks are active, they have more monopoly power in the final loan market. Active banks then demand higher lending rate from entrepreneurs and lend out less for any rate in the interbank market. As a results, the loan demand in the interbank market declines. This shifts the loan demand to the left to the green curve. If the shift is not too big, the equilibrium changes to point b. The amount of loans in the interbank market is higher after introducing the CBDC. Because the lending departments lend out all the borrowed funds in the interbank market, entrepreneurs get more loans and face a lower rate. As a result, the CBDC promotes bank intermediation. However, if the loan demand shifts too much, it could intersect on the vertical part of the loan supply curve, as a result, the loan quantity may decrease and a CBDC reduces bank intermediation.

E Empirical Evidence of Fixed Operating Costs

We now show that banks have increasing return to scales, which is majorly caused by decreasing average costs. This provide an indirect evidence of significant fixed operating costs. We first calculate a bank's profit and revenue per unit of assets from the call report data between 1987 and 2010, which we call profit rate and revenue rate, respectively. We then regress it on the bank's assets (in trillion dollars). To eliminate the outliers, we drop the observations with the lowest 1% of asset or the lowest 1% profit rate.

Table 2 shows the regression results. The first three columns are results on revenue rates. The first column is from a simple OLS. The second column adds bank fixed effects and

Table 2: Evidence of Fixed Operating Costs

	Revenue	Revenue	Revenue	Profit	Profit	Profit
Assets	-0.0027***	-0.0182***	-0.0069***	0.0037***	0.0003	0.0012***
	(0.0004)	(0.0054)	(0.0026)	(0.0004)	(0.0004)	(0.0004)
Bank FE	No	Yes	Yes	No	Yes	Yes
Time FE	No	No	Yes	No	No	Yes
N	885313	885313	885313	885313	885313	885313

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

the last column further adds time fixed effects. The coefficients on assets are negative and significant at 1% level in all specifications. This suggests that average revenue drops with total assets. The last three columns are the same regressions but with the dependant variable being the profit rate. The coefficients on assets are positive and highly significant except one specification. This means that the profit rate is increasing in total assets. Since the revenue rate is decreasing in assets, this suggests the average cost must be decreasing in assets. This is consistent with fixed operating costs.

F Price Competition in the Deposit Market

Consider an alternative deposit market structure, where banks set the real interest rate on deposits. They have market power due to information frictions following Burdett and Judd (1983) and Head et al. (2012). There are a continuum of banks. Each of them quote rates for its checkable deposits and time deposits. Due to information frictions, a buyer does not see all the quotes. Instead, he or she sees one quote with probability b_1 and two quotes with probability b_2. For simplicity, assume that $b_1 + b_2 = 1$. If the buyer sees two quotes, he or she chooses to stay with the bank that quotes a higher rate. After choosing the bank, the buyer works and makes portfolio choices. There is heterogeneity in the portfolio choices because buyers face different interest rates on their deposits. Nevertheless, the inverse demand function for a buyer remains to be $\mathbf{R}_d$. Throughout this section, we require that Assumption 1 holds with $N = 1$. It is convenient to work with demand function instead of inverse demand. Therefore, define $\mathbf{D}(R_d) = \mathbf{R}_d^{-1}(R_d)$ and $\hat{\mathbf{D}}(R_d) = \hat{\mathbf{R}}_d^{-1}(R_d)$. Notice that $\mathbf{D}(R_d) = \hat{\mathbf{D}}(R_d)$ if $R_e = 0$.

Banks engage in perfect competition in the loan market. Given the loan rate, they choose

R_d and R_b to maximize the expected profit. Same as before, banks offer time deposit only if $R_\ell \geq 1/\beta$. If $R_\ell = 1/\beta$, they are indifferent between any amount of time deposits and $R_b = 1/\beta$. Therefore, we can focus only on the choice of R_d. As in the main paper, we derive the loan supply as a function of R_ℓ from the bank's problem and then intersect it with the loan demand curve to determine the equilibrium. The latter remains unchanged, so we focus on the former.

Following Head et al. (2012), there is a continuum of R_d quoted in the equilibrium. They all lead to the same expected profit. Banks trade off the profit from a customer and the probability of getting a customer. Given R_ℓ, the distribution of R_d is $F(\cdot; R_\ell)$. One can show that $F(\cdot; R_\ell)$ is non-atomic and has an interval support. Given R_ℓ, the lowest quoted rate solves

$$\begin{aligned} & \max_{r,\ell,R_d} b_1[R_\ell \ell + R_r r - R_d \mathbf{D}(R_d)] \\ \text{st} \quad & r + \ell = \mathbf{D}\left(R_d\right) \text{ and } r \geq \chi \mathbf{D}\left(R_d\right). \end{aligned}$$

A bank with the lowest rate gets a customer only if the customer sees only one quote. This happens with probability b_1. Conditional on having a customer, the problem is the same as before with $D_{-j} = 0$, i.e., the bank is a local monopoly. This problem can be rewritten as

$$\max_{R_d} b_1 \left(\xi - R_d\right) \mathbf{D}\left(R_d\right). \tag{31}$$

where $\xi = \max\left\{R_\ell \left(1-\chi\right) + \chi R_r, R_r\right\}$. Let $\hat{\underline{\mathbf{R}}}_d(R_\ell)$ solve the following equation in R_d:

$$(\xi - R_d)\hat{\mathbf{D}}'(R_d) - \xi \hat{\mathbf{D}}(R_d) = 0.$$

Then the solution to (31) is $\underline{\mathbf{R}}_d(R_\ell) = \max(R_e, \hat{\underline{\mathbf{R}}}_d(R_\ell))$. The CBDC rate sets a floor for the lowest deposit rate. Then $F(\cdot; R_\ell)$ satisfies the equal-profit condition

$$b_1 \left[\xi - \underline{\mathbf{R}}_d(R_\ell)\right] \mathbf{D}\left(\underline{\mathbf{R}}_d(R_\ell)\right) = \left\{b_1 + 2b_2 F\left(R_d; R_\ell\right)\right\} \left(\xi - R_d\right) \mathbf{D}\left(R_d\right).$$

The left hand side is the profit from quoting the lowest rate. The right hand side is the profit from quoting the higher rate R_d. A bank with rate R_d gets a customer if either the customer has a quote only from this bank or the customer's other quote has a lower rate. This equal-profit condition gives a closed-form solution

$$F\left(R_d; R_\ell\right) = \frac{b_1}{2b_2} \left\{ \frac{\left[\xi - \underline{\mathbf{R}}_d(R_\ell)\right] \mathbf{D}\left(\underline{\mathbf{R}}_d(R_\ell)\right)}{\left[\xi - R_d\right] \mathbf{D}\left(R_d\right)} - 1 \right\}.$$

Denote the highest rate in the support of F as $\bar{\mathbf{R}}_d(R_\ell)$. It solves $F(R_d; R_\ell) = 1$. If $R_\ell \leq 1/\beta$, then $\bar{\mathbf{R}}_d(R_\ell) < 1/\beta$.

The accepted quotes have the price distribution

$$G(R_d; R_\ell) = b_1 F(R_d; R_\ell) + b_2 F(R_d; R_\ell)^2.$$

Use $\hat{F}$ and $\hat{G}$ to denote the equilibrium distributions of quoted and accepted rates if there is no CBDC (i.e., if $R_e = 0$). By Assumption 1, $\bar{\mathbf{R}}_d(R_\ell)$ is increasing in R_ℓ. This implies that as R_ℓ increases, F, $\hat{F}$, G and $\hat{G}$ increases in the sense of first-order stochastic dominance.

Similar as in the benchmark model, the aggregate loan supply without the CBDC is

$$\hat{\mathbf{L}}^s(R_\ell) = \begin{cases} 0 & \text{if } R_\ell < R_r \\ \left[0, (1-\chi)\int \hat{\mathbf{D}}(x)\, d\hat{G}(x; R_\ell)\right] & \text{if } R_\ell = R_r \\ (1-\chi)\int \hat{\mathbf{D}}(x)\, d\hat{G}(x; R_\ell) & \text{if } 1/\beta > R_\ell > R_r \\ \left[(1-\chi)\int \hat{\mathbf{D}}(x)\, d\hat{G}(x; R_\ell), \infty\right] & \text{if } R_\ell = 1/\beta \end{cases}.$$

It takes the same form as the case with Cournot competition. It increases with R_ℓ because $\hat{\mathbf{D}}$ is increasing with R_ℓ and $\hat{G}$ increases with R_ℓ in the sense of first-order stochastic dominance.

Now introduce a CBDC with $R_r < R_e < 1/\beta$ as in the main text. Again, define $\underline{R}_\ell$ as the solution to

$$(1-\chi) R_\ell + \chi R_r = R_e.$$

Let $\bar{R}_\ell$ satisfy $\underline{\hat{\mathbf{R}}}_d(\bar{R}_\ell) = R_e$. If $R_\ell < \underline{R}_\ell$, banks do not operate. If $R_\ell \geq \bar{R}_\ell$, $\underline{\hat{\mathbf{R}}}_d(R_\ell) \geq R_e$ and the CBDC does not change the distribution of R_d, i.e., $F(\cdot; R_\ell) = \hat{F}(\cdot; R_\ell)$ and $G(\cdot; R_\ell) = \hat{G}(\cdot; R_\ell)$. If $\underline{R}_\ell < R_\ell < \bar{R}_\ell$, then $R_e > \underline{\hat{\mathbf{R}}}_d(R_\ell)$. This implies that $F(\cdot; R_\ell)$ and $G(\cdot; R_\ell)$ first-order stochastically dominate $\hat{F}(\cdot; R_\ell)$ and $\hat{G}(\cdot; R_\ell)$, respectively. If $R_\ell = \underline{R}_\ell$, then $R_d = R_e$ is degenerate, so banks are indifferent between any amount of deposits and lend up to the reserve requirement. On the other hand, households are indifferent between the CBDC and checkable deposits. As a result, the deposit quantity can be anything between

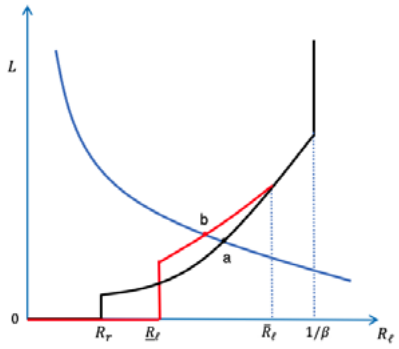

Figure 10: Equilibrium with Burdett-Judd Deposit Market

0 and $\hat{\mathbf{D}}(R_e)$. The aggregate loan supply is

$$\mathbf{L}^s(R_\ell) = \begin{cases} 0 & \text{if } R_\ell < \underline{R}_\ell \\ \left[0, (1-\chi)\hat{\mathbf{D}}(R_e)\right] & \text{if } R_\ell = \underline{R}_\ell \\ (1-\chi)\int \hat{\mathbf{D}}(x)\, dG(x; R_\ell) & \text{if } \bar{R}_\ell > R_\ell > \underline{R}_\ell \\ (1-\chi)\int \hat{\mathbf{D}}(x)\, dG(x; R_\ell) & \text{if } \frac{1}{\beta} > R_\ell \geq \bar{R}_\ell \\ \left[(1-\chi)\int \hat{\mathbf{D}}(x)\, dG(x; R_\ell), \infty\right] & \text{if } R_\ell = 1/\beta \end{cases}.$$

Figure 10 shows the loan demand and loan supply curves. The black curve is the aggregate loan supply curve without the CBDC and the red curve is the curve with the CBDC. Similar to the Cournot model, the red curve is above the black curve if $R_\ell \in \left(\underline{R}_\ell, \bar{R}_\ell\right)$ and overlaps with the black curve if $R_\ell > \bar{R}_\ell$. Different from the Cournot model, it is increasing on $\left(\underline{R}_\ell, \bar{R}_\ell\right)$. The blue curve is the loan demand curve. Its intersections with the loan supply curves correspond to equilibria with and without the CBDC. In this figure, we plot the case where i_e is intermediate. The equilibrium with the CBDC (point b) has a higher loan quantity and a lower loan rate than the equilibrium without the CBDC (point a). Because the equilibrium is between $\underline{R}_\ell$ and $\bar{R}_\ell$, households do not use the CBDC in equilibrium despite that it has a positive effect on bank intermediation. Same as in the Cournot model, a higher i_e increases both $\underline{R}_\ell$ and $\bar{R}_\ell$. As R_e increases from R_r, the CBDC first increases bank intermediation and then decreases bank intermediation.

Proposition 4 *There exists a unique steady-state monetary equilibrium. With a proper interest rate, the CBDC can increase bank intermediation.*

G Bank Profit and Risk-Taking Behaviors

So far, we show that the CBDC can increase bank lending by reducing the market power of banks. A reduction in market power lowers bank profits, which may have implications on the risk of the economy if investment is risky. Although this is not the main focus of this paper, our framework is flexible enough for studying the effect of the CBDC on the risk. It turns out the CBDC can increase or decrease risk in the economy depending on whether it is the entrepreneurs or the banks who decide the risk level of the investment projects. If entrepreneurs make the decision, the CBDC can increase loans and decrease risk taking. If banks make the decision, the CBDC can increase both loans and risk taking. In the latter case, the government can use contingent transfers to reduce the risk taking and use the CBDC to reduce banks' market power. This opens the door to the general question of coordinations between the CBDC policy and policies targeting directly at the risk-taking behaviors. We leave a careful study of this policy coordination problem to future research.

G.1 Risk-Taking Behaviors of Entrepreneurs

Following Boyd and De Nicolo (2005), each entrepreneur has access to a risky project that requires 1 unit of investment. Each project gives s CM good with probability $p(s)$ and 0 with probability $1-p(s)$, where p is strictly decreasing and concave with $p(\overline{s})=0$ and $1 \geq p(0) > 0$. The projects with the same risk level are perfectly correlated. Each entrepreneur has a heterogeneous cost κ to operate the investment. The measure of entrepreneurs with κ less than x is $G(x)$.

Entrepreneurs have limited liability. They pay back the loans only if the project is successful. Notice that entrepreneurs care only about the expected profit and do not have incentives to diversify investments across projects with different risk levels. Therefore, their problem is

$$\max_{s} p(s)(s - R_\ell),$$

which yields a decision for risk level

$$R_\ell = s + \frac{p(s)}{p'(s)}.$$

This defines risk taking as an increasing function of R_ℓ denoted as $\mathbf{s}(R_\ell)$. Higher lending rate induces entrepreneurs to take more risk and utilize their limited liabilities. Then the

profit of a entrepreneur is

$$\pi(\mathbf{s}(R_\ell)) = p\left[\mathbf{s}(R_\ell)\right]\left[\mathbf{s}(R_\ell) - R_\ell\right],$$

which is decreasing in R_ℓ by the envelope condition. An entrepreneur invests if and only if $\pi(\mathbf{s}(R_\ell)) \geq \kappa$. Therefore, the total loan demand is $G(\pi(\mathbf{s}(R_\ell)))$, which is decreasing in R_ℓ. We can invert it to obtain the inverse loan demand function $\mathbf{R}_\ell$.

Now we consider the bank's problem. Banks engage in Cournot competition in both the deposit and the loan markets. To simplify presentation, we assume that all bank deposits are insured by the government at zero cost. The government finances deposit insurance by lump-sum tax in case of a bank default. We assume that banks do not have enough assets to cover liabilities if entrepreneurs default on the loan, which requires χ to be small. If a bank fails, it exhausts its assets to pay back deposits but do not have further obligation. Banks engage in Cournot competition in both deposit and loan markets. Therefore, bank j solves

$$\begin{aligned}\max_{\ell_j, r_j, d_j} \; & \mathbf{q}(L_{-j} + \ell_j)\left[\mathbf{R}_\ell(L_{-j} + \ell_j)\ell_j - \mathbf{R}_d(D_{-j} + d_j)d_j + R_r r_j\right] \\ \text{st } & \ell_j + r_j = d_j \text{ and } \chi d_j \leq r_j,\end{aligned} \tag{32}$$

where $\mathbf{q}(\cdot) = p[\mathbf{s}(\mathbf{R}_\ell(\cdot))]$.

We form the Lagrangian

$$\begin{aligned}\max_{\ell_j, r_j, d_j} \Big\{ & \mathbf{q}(L_{-j} + \ell_j)\left[\mathbf{R}_\ell(L_{-j} + \ell_j)\ell_j - \mathbf{R}_d(D_{-j} + d_j)d_j + R_r r_j\right] \\ & + \Lambda_1(d_j - \ell_j - r_j) + \Lambda_2(r_j - \chi d_j)\Big\}.\end{aligned}$$

Imposing symmetry, we can obtain a set of equilibrium conditions

$$0 = \mathbf{q}(L)\left[\mathbf{R}'_\ell(L)\frac{L}{N} + \mathbf{R}_\ell(L)\right] + \mathbf{q}'(L)\left[\mathbf{R}_\ell(L)\frac{L}{N} - \mathbf{R}_d(D)\frac{D}{N} + R_r r_j\right] - \Lambda_1, \tag{33}$$

$$0 \geq [(1-\chi)\Lambda_1 + \mathbf{q}(L)\chi R_r] - \mathbf{q}(L)\left[\frac{\partial^+}{\partial D}\mathbf{R}_d(D)\frac{D}{N} + \mathbf{R}_d(D)\right], \tag{34}$$

$$0 \leq [(1-\chi)\Lambda_1 + \mathbf{q}(L)\chi R_r] - \mathbf{q}(L)\left[\frac{\partial^-}{\partial D}\mathbf{R}_d(D)\frac{D}{N} + \mathbf{R}_d(D)\right], \tag{35}$$

$$0 = \mathbf{q}(L)R_r - \Lambda_1 + \Lambda_2, \tag{36}$$

$$0 = \Lambda_2 \text{ if } L < (1-\chi)D \text{ and } 0 < \Lambda_2 \text{ if } L = (1-\chi)D. \tag{37}$$

Again Λ_1 is the marginal value of loanable funds. Different from the case without risk, the bank profit when it does not fail enters into (33). If this profit is lower, Λ_1 is lower. Intuitively, banks have an incentive to issue more loans to reduce entrepreneur's risk, which increases their probability of being profitable. However, if the bank profit drops, they have less incentive to reduce the risk and as a result the value of loanable funds drops. This channel could reduce bank lending and increase risk if the CBDC cut bank profits.

Proposition 5 *If entrepreneurs decide the level of risk, a CBDC can increase bank lending and decrease risk if R_e is slightly higher than the gross real rate of deposits without the CBDC.*

Proof. If there is no CBDC, $\mathbf{R}_d = \hat{\mathbf{R}}_d$ is a smooth function. Therefore, (34) and (35) hold as equality. Denote the equilibrium quantities without the CBDC as $\hat{D}^*, \hat{L}^*, \hat{\Lambda}_1^*, \hat{\Lambda}_2^*$. Consider the case where $\hat{\Lambda}_2^* > 0$ and the reserve requirement binds. We show the if i_e is set such that R_e is only slightly higher than $\hat{\mathbf{R}}(\hat{D}^*)$, then $D^* = D_e$, $L^* = (1-\chi)D^*$ constitute an equilibrium with a CBDC. To achieve this, we only need to show that we can find $\Lambda_1^* > 0$ and $\Lambda_2^* > 0$ such that (33)-(36) hold with $(D^*, L^*, \Lambda_1^*, \Lambda_2^*)$. We first define Λ_1^* to be the solution of (33) if we replace (D, L) by (D^*, L^*). Notice that Λ_1^* is continuous in D_e and hence continuous in R_e. Now we evaluate (35) at (D^*, L^*, Λ_1^*). Because $\frac{\partial^-}{\partial D}\mathbf{R}_d(D_e) = 0$ and $\frac{\partial^-}{\partial D}\hat{\mathbf{R}}_d(\hat{D}^*) > 0$, the right-hand side of (35) is strictly greater than 0 if R_e is slightly higher than $\hat{\mathbf{R}}_d(\hat{D}^*)$. Similarly, the right-hand side of (34) evaluated at (D^*, L^*, Λ_1^*) is less than or equal to 0. Otherwise, it suggests that the equilibrium without the CBDC should feature more than D^* deposits and L^* loans, which contradicts the assumption. Lastly, because Λ_1^* and L^* are continuous in R_e, (36) defines a positive Λ_2^* if R_e is not too far from $\hat{\mathbf{R}}_d(\hat{D}^*)$. Now we have constructed Λ_1^* and Λ_2^* that satisfy the requirement. Therefore, D^* and L^* are the equilibrium quantities if R_e is higher but not too much higher than $\hat{\mathbf{R}}_d(\hat{D}^*)$. Because $L^* > \hat{L}^*$ and risk is decreasing in total loan supply, a CBDC can increase bank lending and reduce risk. ■

G.2 Risk-Taking Behaviors of Banks

We next study a variation of the model where banks choose the risky investment. This model is also studied in Boyd and Di Nicolo (2005). It is the same as the above except that now the banks decide both the risk of the investment s and the quantity of the investment ℓ. The investment returns $s\ell$ with probability $p(s)$ and 0 with probability $1-p(s)$. Returns are perfectly correlated across projects.

Bank j solves

$$\max_{s_j,\ell_j,r_j,d_j} p(s_j)\left[s_j\ell_j - \mathbf{R}_d(D_{-j}+d_j)d_j + R_r r_j\right] \quad \text{st } \ell_j + r_j = d_j \text{ and } \chi d_j \le r_j. \tag{38}$$

We form the Lagrangian:

$$\max_{s_j,\ell_j,r_j,d_j} \left\{ p(s_j)\left[s_j\ell_j - \mathbf{R}_d(D_{-j}+d_j)d_j + R_r r_j\right] + \Lambda_1(d_j - \ell_j - r_j) + \Lambda_2(r_j - \chi d_j) \right\}.$$

Impose symmetry and eliminate Λ_1 and Λ_2 to obtain

$$0 \ge [(1-\chi)s + \chi R_r] - \left[\frac{\partial^+}{\partial D}\mathbf{R}_d(D)\frac{D}{N} + \mathbf{R}_d(D)\right], \tag{39}$$

$$0 \le [(1-\chi)s + \chi R_r] - \left[\frac{\partial^-}{\partial D}\mathbf{R}_d(D)\frac{D}{N} + \mathbf{R}_d(D)\right], \tag{40}$$

$$0 = p(s)\frac{L}{N} + p'(s)\left[s\frac{L}{N} - \mathbf{R}_d(D)\frac{D}{N} + R_r\frac{D-L}{N}\right], \tag{41}$$

$$s = R_r \text{ if } L < (1-\chi)D \text{ and } s > R_r \text{ if } L = (1-\chi)D. \tag{42}$$

Proposition 6 *If banks decide the risk level directly, a CBDC can increase bank lending and risk taking if R_e is slightly higher than the gross real rate of deposits without the CBDC.*

Proof. Let $\hat{s}^*, \hat{D}^*, \hat{L}^*$ be the equilibrium without the CBDC. And assume that $\hat{s}^* > R_r$. Then $\hat{L}^* = (1-\chi)\hat{D}^*$. Therefore, $\hat{s}^*, \hat{D}^*$ solve

$$0 = [(1-\chi)s + \chi R_r] - \left[\hat{\mathbf{R}}'_d(D)\frac{D}{N} + \hat{\mathbf{R}}_d(D)\right], \tag{43}$$

$$0 = [s + \frac{p(s)}{p'(s)}](1-\chi) + \chi R_r - \hat{\mathbf{R}}_d(D). \tag{44}$$

Now we show that if R_e is higher than but close to $\hat{\mathbf{R}}_d(\hat{D}^*)$, then $D^* = D_e$ and $L^* = (1-\chi)D^*$ constitute an equilibrium. Let s^* be the solution to (44) when $D = D^*$. If R_e is close to $\hat{\mathbf{R}}_d(\hat{D}^*)$, D^* is close to $\hat{D}^*$. Therefore, s^* is close to $\hat{s}^*$ and strictly larger than R_r. Moreover, because $D^* > \hat{D}^*$ and $s + p(s)/p'(s)$ is increasing, $s^* > \hat{s}^*$. If s^*, D^* and L^* satisfies (39)-(41), they are the equilibrium with the CBDC. First, (41) is satisfied by the definition of s^* and the fact that $\hat{\mathbf{R}}_d(D^*) = \mathbf{R}_d(D^*) = R_e$. Second, (40) is satisfied by a similar argument as in the previous section. Lastly, one can show that (39) is satisfied because otherwise, the

equilibrium without the CBDC should have a deposit quantity larger than D^*. This shows that (D^*, L^*, s^*) is the equilibrium if R_e is slightly higher than $\hat{\mathbf{R}}_d(\hat{D}^*)$. Therefore, both bank lending and risk level increase after introducing the CBDC. ■

The government can keep the risk level unchanged by giving banks a lump-sum transfer if they do not fail. Define

$$\tau = \left\{ R_e - [\hat{s}^* + \frac{p(\hat{s}^*)}{p'(\hat{s}^*)}](1-\chi) - \chi R_r \right\} \frac{D^*}{N} = [R_e - \hat{\mathbf{R}}_d(\hat{D}^*)] \frac{D^*}{N}.$$

One can check that $(D^*, L^*, \hat{s}^*)$ constitutes an equilibrium. Notice that τ and the CBDC rate are two orthogonal policy tools. Although the government can choose τ to maximize the expected productivity $p(s)s$, the quantity of investment can be too low due to the market power of banks. The CBDC can help to raise investment quantity. Although such transfers may not be possible in practice, the message is that the government may have other policy tools that target directly at the risk-taking behaviors. These tools can be used in conjuncture with the CBDC rate to reduce banks' market power without making them more risky. An example of such a tool is the contingent transfer studied above. There are more practical tools like capital requirements. A careful study of the interaction between these policies and the CBDC rate can be very interesting and is beyond the scope of this paper. Therefore, we leave it for future research.

H Calibration Method and Data

In the calibration, we use Kalai bargaining as the DM trading mechanism. It is more flexible and allows sellers in the DM to have positive markups. The bargaining power to the buyer is $\theta \in [0, 1]$. The solution maximizes the buyer's surplus given that he or she gets θ fraction of the total surplus and the liquidity constraint, i.e., if the buyer has a real balance $\mathcal{L}$, the Kalai solution solves

$$\max_{y,p} [u(y) - p] \text{ s.t. } u(y) - p = \theta [u(y) - y] \text{ and } p \leq \mathcal{L}.$$

All the analysis in the main text stays unchanged except that

$$\lambda(\mathcal{L}) = \max \left\{ \frac{u'[Y(\mathcal{L})]}{(1-\theta)u'[Y(\mathcal{L})] + \theta} - 1, 0 \right\},$$

where $Y(\mathcal{L})$ satisfies $(1-\theta)u[Y(\mathcal{L})] + \theta Y(\mathcal{L}) = \mathcal{L}$. If $\theta = 1$, Kalai bargaining reduces to the buyer's take-it-or-leave-it offer studied in the main text.

Data

From FRED, we obtain the time series for inflation, 3-month t-bill rates, prime rates, GDP and total commercial loans. The SCPC data contain the number of transactions by type. Table 9 in Greene and Stavins (2018) contains these numbers for 2015-2017. DCPC asks consumers to record whether they think a transaction accepts cash or cards. Page 13 and 14 in Premo (2018) contain summary statistics of answers to these questions.

For calibration, we need times series of interest rates on transaction deposits and loan, and information on the operation costs of banks. We obtain them from call report data from 1987-2019. This data contain quarterly information on balance sheet and income statement of banks in the US. We obtain this data from WRDS by using the SARS code by Drechsler et al. (2017). To obtain the rates on transaction deposits, we first divide interest expenses on transaction accounts (item code: RIAD 4508) by total transaction deposits (RCON2215) to obtain the rates for each bank in a given quarter. Then we take the average across all banks weighted by their transaction deposits to obtain a quarterly industry average. Lastly, we aggregate to the annual level. Notice that RIAD4508 starts only from 1987.

To obtain the loan rates, we first divide the interest income from loans (RIAD4010) by the loan quantity (RCON3360) to obtain bank-level loan rates. These rates are very heterogeneous across banks. This could be because loans of different banks have different risk. Since we do not model risky investment, we focus only on the safe loans. We define our loan rate to be the first percentile of the loan rate distribution in the data. This leads to rates comparable to the FRED on loans of minimal risk, which is reported between 1998 and 2008. In this period, our average loan rate is about 4.4%, while the FRED data have an average rate of 4.7%. This gives a 3.69% loan rate. We have also done a calibration with the average loan rate 5.19%. The results are similar but the positive effect of the CBDC is larger because the higher loan rate implies a higher market power in the banking sector.

Lastly, we compute the operational cost per dollar of asset between 2014 and 2019 to calibrate c. Unfortunately, assets data are missing for many banks during this time frame, but we observe total deposits (RCON2200+RCFN2200) for this period. We also observe both assets and deposits between 1987 and 2010. In this period, assets is about 1.505 times the total deposits. We then assume that the this ratio is stable over time. Therefore, we can divide the operational cost per dollar of deposits by 1.505 to obtain operational cost per dollar of assets. To this end, we first calculate the average operational cost for each bank by subtracting expenses on premises or rent (RIAD4217) from the non-interest expenses (RIAD4903). Then take an average across banks weighted by total deposits to get the industry average. Finally,

we aggregate to annual level and set c to be the time average divided by 1.505.

Computation

One straightforward way to calibrate the model is to solve the equilibrium given each parameter value and choose one that best fits the money demand curve, the deposit rates and the spread. This method, however, can be computationally cumbersome because one needs to solve the model for each data point used for the money demand and then optimize over a six-dimensional parameter. One key insight is that the money demand can be solved independent of the banking sector. This leads to the following algorithm that greatly simplifies the calibration.

1. Match the money demand between 1987 and 2008 to obtain $B, \sigma, \theta, \Omega$. In this step, we set $i_r = 0$ and $\chi = 2.4\%$ to match the average interest on reserves and the average ratio of required reserves excluding vault cash to transaction balances during this period.

   (a) Fix the value of Ω and θ. Fit the money demand curve by choosing B, σ. More specifically, for each interest rate, calculate the steady state equilibrium using the nominal interest rate and the deposit rate for each year. Then choose (B, σ) to minimize the distance between the model predicted M1 to GDP ratio and the data. The M1 to GDP ratio in the model can be calculated by $\left(R_z Z + R_d D\right)/\mathbf{Y}$, where

   $$\mathbf{Y} = \sum_{j=1}^{3} \alpha_j P(y_j) + 2B + \frac{D - L}{1 + \pi} + AL^{\eta} - R_d D + L$$

   is the output and π is the inflation rate. It is the sum of the consumption of households in DM and CM (the first two terms), the consumption of old bankers $(R_r(D-L)+R_l L-R_d D)$, the consumption of old entrepreneurs $(Af(L)-R_l L)$, and the investment by young entrepreneurs (L). The DM consumption is measured by the amount of payment. Because the reserve requirement binds during this period, $L = (1-\chi)D$. The first-order condition of the entrepreneurs implies $A\eta L^{\eta-1} = R_\ell$. Therefore,

   $$\mathbf{Y} = \sum_{j=1}^{3} \alpha_j P(y_j) + 2B + \frac{\chi D}{1 + \pi} + \frac{\left(R_\ell + \eta\right)\left(1 - \chi\right) D}{\eta} - R_d D.$$

   We plug in the time series for R_ℓ and R_d. The above formula does not involve A. This insight allows us to calibrate B and σ independent of the bank's problem.

(b) Calculate the markup. If it is less than 20%, then decrease θ, otherwise increase θ. Repeat 1-2 until the markup in the model matches 20%.

(c) Calculate the model fit at different values of Ω. And find the value that gives the best fit.

2. Match banking data from 2014 to 2019 to obtain A, N

(a) Set N such that the solution of the Cournot competition leads to a spread of 3.39% between loans and transaction account.

(b) Set A to match a 0.3049% interest rate on transaction accounts.

I The Role of Time Deposits

The analysis of this paper does not depend on the existence of time deposits. But they help to avoid a technical complication: if checkable deposits are the only funding source, the loan supply curve is not continuous and an equilibrium where banks adopt pure strategies may not exist.

To see this, suppose checkable deposits are the only funding source and $\chi = 0$. When $R_\ell = 1/\beta$, there are two types of equilibria in the Cournot game among banks. In the first type, the gross real deposit rate is $1/\beta$ and banks make 0 profits. This occurs if for any bank j, its competitors supply $D_{-j} \geq D^* = \beta y^*$. Buyers are unconstrained even if bank j supplies 0 checkable deposits and the liquidity premium for checkable deposits disappears. Therefore, $R_d = R_\ell = 1/\beta$ and bank j is indifferent between any amount of checkable deposits. Because the lending rate is positive, the deposit quantity is equal to the loan quantity. Therefore, the total loan supply equals the total deposits supply and is at least $ND^*/(N-1)$ by symmetry.

In the second type of equilibrium, the gross real return on deposits is less than $1/\beta$ and banks earn positive profits. In this case, for any bank j, its competitors supply $D_{-j} < D^*$ deposits. Therefore, it is optimal for bank j to supply less than $D^* - D_{-j}$ and obtain a positive profit. By symmetry, the total deposits and loans are $N\hat{\mathbf{d}}(1/\beta) < D^*$.

The loan supply cannot take any value between $N\hat{\mathbf{d}}(1/\beta)$ and $ND^*/(N-1)$ in a symmetric pure-strategy equilibrium, generating a hole in the loan supply curve. In fact, one can show that the loan supply cannot be between D^* and $ND^*/(N-1)$ even if we allow for asymmetric strategies. Figure 11 shows the implication on the equilibrium. The black curve is the loan supply without the CBDC and the red curve is the loan supply with the CBDC. If the loan

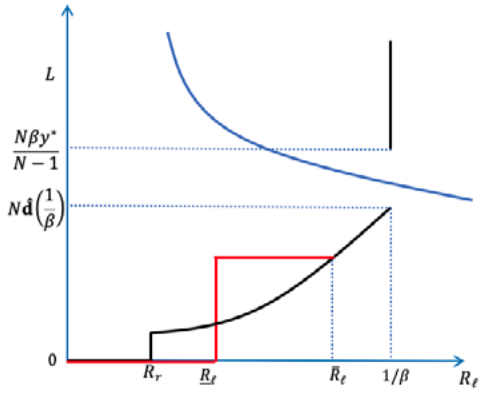

Figure 11: No Time Deposits

demand goes through the hole of the loan supply curve as in Figure 11, there does not exist an equilibrium. If time deposits are available, banks can use funding raised by time deposits to fill the gap in the loan supply curve. An equilibrium always exists.

Digital Money and Central Bank Operations

INTERNATIONAL MONETARY FUND

Digital Money and Central Bank Operations

Charles Kahn, Manmohan Singh, and Jihad Alwazir

WP/22/85

WORKING PAPER

 WP/22/85

IMF Working Paper
Monetary and Capital Markets Department

Digital Money and Central Bank Operations

Charles Kahn, Manmohan Singh, and Jihad Alwazir

May 2022

ABSTRACT: The rise of new and proposed monetary vehicles, including CBDC, stablecoins, payment service providers etc., are unprecedented. An important question for central banks is the extent to which these innovations upend the role of and implementation of monetary policy. The paper focuses on the interest rate channel and if digital money (especially CBDC) will change monetary policy and central bank operations. We argue that new policy instruments make sense only to the extent that there is limited substitutability between the various payment sectors. We analyze trends in currency-in-circulation, and how it may impact central bank's seigniorage, monetary base, and transactional velocity of digital money if money demand declines. Liquidity outside the monetary base will also be important to understand.

JEL Classification Numbers:	E5; E41; G15; F30
Keywords:	Base money; CBDC; central banking operations; currency in circulation; digital money; mobile phone operators; seigniorage.
Author's E-Mail Address:	cmkahn@illinois.edu; msingh@imf.org; jalwazir@imf.org

WORKING PAPERS

Digital Money and Central Bank Operations

Prepared by Charles Kahn, Manmohan Singh, and Jihad Alwazir[1]

[1] The authors would like to thank Peter Stella, Karl Habermeier, Simon Gray, and Adrian Armas for comments. Authors are responsible for any errors.

Contents

BOXES

FIGURES

Introduction

The rise of new and proposed monetary vehicles, including central bank digital currency (CBDC), stablecoins, payment service providers (e.g., mobile network operators), etc., means an unprecedented change in the retail and wholesale payments system. At the same time, central banks and infrastructure providers are examining new ways to facilitate transfer of value across wholesale payments platforms. An important question for central banks is the extent to which these innovations upend the role of and implementation of monetary policy.

On the one hand, it can be argued that while the changes the monetary systems are undergoing are rapid, they are no more extensive than changes seen in the past: the rise of card-based payments at the retail level, netting systems and new cross-border systems at the wholesale level—all of which were taken in stride by the departments in the central banks responsible for monetary policy implementation. On the other hand, several observers and commentators have argued that the speed of the change, the ability of technology to effect rapid, possibly uncontrolled transmission from one payments arrangement to another, and the increased attractiveness of new non-bank, and non-regulated structures mean that central banks must take account the extra stress on their systems,[1] and indeed the threat of irrelevance. Coupled with this threat is a potential opportunity, noted by other commentators: the addition of new monetary instruments may enable the central bank to operate on additional dimensions, opening up new avenues for more targeted policy responses, through adjustment of multiple interest rates or monetary aggregates.

This paper attempts to start to untangle this puzzle, examining the effect on monetary policy implementation of the introduction of CBDCs and of bank and nonbank stablecoins into an economy (Quarles, 2021; Carney, 2021). The significance of the advent of alternative payment arrangements differs depending on whether the monetary authority is conducting a traditional policy intended to affect overall real interest rates, or a quantitative easing or tightening, focused on altering the differential between returns on two different categories of assets, for example maturity premia, risk premia, or liquidity premia. It also differs depending on whether the authority implements its policy by targeting nominal interest rates or by targeting monetary aggregates. We take the perspective in this paper that controlling relevant interest rates is the macroeconomic goal of the monetary authority, and that monetary aggregates are used as a target by an authority that finds interest rates are difficult to observe or directly control on a timely basis.

In addition, a central bank has other goals besides reaching a macroeconomic target. New payment systems can be sources of financial instability and so it may be appropriate to have policies which focus on these risks. Different forms of new money will yield different seigniorage revenues; if these revenues are an important part of the central bank's mandate, it may need to establish policies that tilt the mixture of payment arrangements chosen in the economy. There are many different implications for central bank operations from the introduction of digital money; the paper will focus on three questions: (i) Can digital money affect the interest rate channel?; (ii) Would it require a new instrument for the central bank?; and (iii) What are the implications for currency-in-circulation, monetary base and seigniorage?

The paper is outlined as follows: Section I describes key features of the existing and new types of payments arrangements we consider. Section II focuses on the interest rate channel and whether digital money will change monetary policy and central bank operations (with a focus on CBDC). It also suggests the need to

[1] For example, Gorton and Zhang (2021) argue that, while there is nothing new about privately provided money, stablecoins create systemic risks that should be addressed by regulation and CBDC issuance.

redefine monetary aggregates to take into account new payment assets. Section III focuses on other central bank goals. It argues that new instruments make sense to the extent that there is limited substitutability between the various payment sectors, both on the demand side of customers for the instruments, and on the provision side. Section IV centers on currency-in-circulation (CiC) trends globally, and how it may impact central bank operations such as seigniorage, transactional velocity of digital money, and sterilization, including parallels between CiC and CBDC.

Many of the issues that arise from the introduction of non-bank providers of payments assets are already present in countries with large penetration of e-money. Section V focuses on a sub-set of digital money: e-money and payment service providers (PSP) such as mobile network operators (MNOs), and how liquidity outside the monetary base may be important to understand, including decline in money demand and other money metrics such as M2. Section VI concludes that we will need a better understanding how the elasticities of their demand of the new digital technology compare with the elasticities of the demand for the existing methods of payment.

I. Payments Assets, Old and New

Proposed and actual innovations in payments are arriving at a mind-boggling rate. To understand the different implications, we need to provide our own typology (for different typologies, see Adrian and Mancini-Griffoli, 2021; Carstens, 2021, McLaughin, 2022, Bech and Garratt, 2017). For the most part, existing money comes in two varieties: debt of a central bank, used by individuals in the form of physical currency, and bank money—that is, debt of commercial banks in a form acceptable for payments. Central bank money is fiat currency, issued by the central bank essentially without cost. Commercial bank money promises redemption in central bank debt. Commercial bank money is backed by reserves of central bank debt, by regulatory structures assuring the safety and soundness of the issuing commercial bank, and by deposit insurance and central bank lender-of-last-resort facilities.

Since money is useful for payments it commands a liquidity premium. Thus, issuers of it can reap profits by providing it to agents in the economy who desire it for payments purposes. When provided by the central bank, the profits are "seigniorage": the central bank trades its monetary asset for non-monetary assets (for example Treasury bonds) and profits from the difference between the interest payable on the bonds and the interest cost it pays (typically zero) on the monetary asset. [2] Similarly, a commercial bank profits from the spread between the interest it receives on the loans it makes, and the lower interest it pays to individuals who hold transactions deposits for payments purposes, net of the costs of any reserves it is necessary carry as backing for the deposits. The mix of cash and demand deposits that the public prefers to hold depends on the convenience and relative cost of each.

We consider two basic types of monetary innovations: (i) central bank digital currency (CBDC); and (ii) bank electronic money and fintech-issued electronic money. In each case we only consider moneys whose rate of exchange with existing money is intended to be "fixed."[3] The primary difference in types of innovations we consider is in the nature of the guarantee of fixed redemption. The guarantee for CBDC comes from the fact

[2] Seigniorage is (and has been) an integral part of non-AE. Central banks send a check to their Ministry of Finance as part of the MoF budget. See Reserve Bank of India speech: https://www.rbi.org.in/Scripts/BS_SpeechesView.aspx?Id=1111.
[3] Realistically speaking only assets denominated in the prevailing unit of account are serious contenders for use as mainstream payments assets.

that the central bank can issue whatever currency it needs to redeem CBDC. The redemption guarantee for the other forms of money depends on the nature and extent of the backing the issuing entity holds: electronic money may be backed wholly or fractionally, and the backing asset may be central bank reserves or short-term government assets. The liabilities may be liabilities of the payment institution, backed by the assets of the institution as a whole, or the backing assets may be ring fenced; for example, held by the issuing authority in custody for the holders of the electronic money, protected from any bankruptcy risk of the institution (see Box 1 and Box 2).

Box 1. Types of Central Bank Digital Currency

A variety of arguments have been made as to why central banks might wish to issue CBDC, among them, financial inclusion, spurring retail payments innovation by slow moving financial institutions, simplifying wholesale and international payments, breaking the zero-nominal interest rate lower bound in monetary policy, and protecting central banks from irrelevancy. Corresponding to this variety of justifications, there has been a variety of design proposals, suggesting CBDCs at the retail level and at the wholesale level; CBDCs implemented through accounts at the central bank, or through wallets maintained by intermediate institutions, CBDCs which are non-interest bearing or interest bearing (or possibly negative interest bearing), and CBDCs with restrictions on the amounts a user can accumulate or with additional functionality for use in smart contracts.[1/] Each variation would potentially put its own wrinkle on monetary policy (Sanches and Keister, 2021; Kahn et al., 2020; Adrain and Mancini-Grifolli, 2019).

In this paper we focus on a set of features we regard as the most relevant case, a non-interest bearing asset issued by the central bank, useful for payments purposes and acting as a substitute for CiC, and freely redeemable in CiC (although authorized firms will probably handle customer service for administering the electronic wallets in which CBDC is stored).[2/] To the extent that the CBDC is an improved means of payment over CiC, interest would be unnecessary for its acceptance. At the retail level, an interest-bearing CBDC is most easily thought of as a savings vehicle substituting for time deposits and the like, and its economic effect would be comparable, for example, to an expansion of a program for issuing retail government savings bonds. At the wholesale level, an interest bearing CBDC would quickly replace other forms of central bank reserves. For central banks, moving from being a net recipient of interest from the rest of the economy to a net payer of interest to the rest of the economy would be a dramatic shift, one that no central bank would willingly contemplate.

1/ See for example the recent speech by Brainard (2022) in which she highlights "*design features that could be introduced to limit such risks, such as offering a non-interest bearing CBDC and limiting the amount of CBDC an end user could hold or transfer."* Brainard also notes that *"such that financial intermediaries rather than the Federal Reserve interface directly with consumer."* (See also Box 2 on Stablecoins.)

2/ Some have considered the possibility of an interest bearing CBDC; for example, CBDC may remunerate retail in the same vein as deposits at a bank. Also, inter-bank market settle in central bank money or "reserves"; however, excess reserves and paying interest on excess reserves (IOER) is only a recent post-Lehman phenomenon. This paper is not about QE related excess reserves, which is a paper in itself (e.g., should wholesale CBDC be remunerated at policy rate, or at excess reserve rate, or at zero rate like CiC, etc.)

3/ If policy rate in a country is 12 percent, remunerating the *stock* of CBDC at 12 percent (and not just the *flow)*, may likely result in negative seigniorage—unless required reserves from the banking system are relatively large. This will be a significant change from CiC that generates positive seigniorage.

Box 2. Stablecoins are Money if Backed by Central Bank Reserves

Stablecoins are sometimes dismissed as the poor relations of the cryptocurrency family. Think of them like government money market funds that are anchored at a par value of $1, versus stocks whose prices can swing around wildly. Stablecoins are designed to be a medium of exchange, rather than a speculative asset. And, despite their apparent stability, they may pose bigger challenge for policymakers than their freewheeling crypto-cousins. The market for stablecoins backed by high quality liquid assets is around $180 billion and sizable growth is expected for US dollar backed coins. To maintain a stable value, issuers need to back the coins with a riskless asset, such as short-term US Treasury obligations. This introduces into the economy a privately established dollar-denominated currency that is not backed by reserves at the US central bank. Their appearance on the scene requires a rethink of the basics of monetary policy—a rethink so fundamental that it is useful to go back to the foundations of monetary policy, and rework from there.

The traditional understanding of monetary policy was based on the idea that the money supply was influenced by the central bank's open market operations (OMOs). Most money was bank account money—deposits at banks—after all. Central banks require that deposits be partially backed by central bank reserves, they tended to be a multiple of those reserves; in this model, US Treasury bonds, while a safe and interest-bearing, cannot be used to back bank accounts. So, when the Fed carried out OMOs, trading its reserves for Treasuries, it changed the amount of reserves available to banks and thus altered the money supply. One consequence of stablecoins becoming available as money is that the money supply no longer needs to be backed by reserves only—Treasuries work just as well. Thus, the central bank's ability to influence the money supply through OMOs will be reduced as stablecoins grow.

Reserves will continue to be in demand in the banking system. The crucial difference between reserves and safe collateral is not their safety (both are safe) but their liquidity. For some purposes—specifically for instantaneous transmission over large value payments systems in order to meet obligations—reserves are useful, and Treasuries are not. However, this need for reserves does potentially have knock-on effects as the bank must be able to provide payment services to any customer with a demand deposit. To the extent that the rest of the economy depends on banks to make payments for them, the banking system will need reserve balances (Singh, Kahn, Long, 2021).

But, as the stablecoin business grows sizably, the demand for Treasuries, or Bunds or JGBs will grow as well. Central banks have no means of directly meeting this demand through standard monetary policy. An alternative would be to allow—or even encourage—stablecoin issuers to use reserves as backing rather than Treasuries (Singh et al, 2021). Nonbank stablecoin issuers would likely favor direct access to reserves through Fed master account and access to central bank payment rails, as this would be preferable to siloing caches of Treasuries or obtaining reserves through a correspondent bank (i.e., very unlike Tether). Should circumstances warrant, reserves are more plentiful in the post QE and post COVID era than good collateral (e.g., Eurozone).

At one extreme are stablecoins whose backing is central bank reserves or ring-fenced deposits of a banking system; at the other extreme are some "e-moneys" issued by companies and backed primarily by the reputation of the company itself (as is the case with PSP or MNOs in some jurisdictions—see Section V).

II. Interest Rate Channel

The most basic route for macroeconomic effects of monetary policy is the interest rate channel.[4] By altering, for example, the interest rate on government bonds, macroeconomic policy alters the level of investment in the economy, thereby adjusting the levels of economic activity and inflation. While the introduction of new means of payment can themselves have effects on economic activity, these are unlikely to be of first-order importance or sufficiently rapid to pose a significant effect on monetary policy in these circumstances. Thus target interest rates for real investment in the country are unlikely to change as a result of payments innovations.

Typically, the central bank targets the interest rates in the interbank market. Changing the cost of funding for banks, the monetary authority attempts to affect general economic activity, as banks expand or contract their lending activity in response. However new means of payment can also have direct effects on the financial institutions' behavior, either encouraging or discouraging its activity. Thus, one possible consequence of a payments innovation arises through the potential for new payments arrangements to change spreads between bank funding costs and lending rates, either by increasing the cost of funds through introduction of payments arrangements which compete with bank payment systems, or by enhancing the efficiency of those systems. Not only may changes in the payment system alter the average level of the differential between the rates set by the monetary authority and the rates at which banks lend to the public, they may also affect the variability of the differential.

The effects are even more significant if the monetary authority targets monetary aggregates. Money multipliers for various types of new payments media will be different from existing money multipliers and are likely to vary based on different external shocks. The question of which aggregates to target and how strongly to respond to changes in those aggregates will depend on the substitutability between the various payment mechanisms.

Consider for example, the effects of policy of quantitative easing whereby a central bank purchases relatively illiquid, non-payments assets in return for money. Changing the relative availability of the two types of assets changes their relative price. The effects on the interest rate premium will generally be more dramatic the less substitutable and more segmented the markets for two assets. If on the other hand, intermediate assets are readily available which serve as acceptable substitutes for each, the effect of the policy on the targeted asset's return is likely to be diluted.[5]

The introduction of a new payment asset increases the options available for making payments, and thereby in general increases economic efficiency. However, to the extent that the new payment asset substitutes for existing payments assets, it reduces the effectiveness of attempts to change the supplies of those existing assets. For example, it is more difficult for the central bank to attempt to reduce the liquidity of an economy if other agents can provide assets which are ready sources of liquidity.

Introducing a CBDC does not lead to this difficulty, since the CBDC is simply another asset issued by the central bank useable for payments. Since the CBDC is denominated one-for-one in units of existing currency

[4] We focus on this channel as most relevant for the introduction of new moneys; but other channels are also important.
[5] The argument is really about the difference between the power of a monetary authority to engage in monetary policy by causing an economy wide change in interest rates, or by a quantitative easing or tightening, targeting one sector of interest rates relative to another. The question is really how tight the link remains between HQLA and reserves, once stablecoins become a major component of the demand for HQLA.

and acquired in exchange for other currency or assets it makes no fundamental change in the conduct of monetary policy; with a CBDC the outside money supply is simply the total of CBDC and cash in circulation. Indeed, to the extent that the CBDC increases the usefulness of central bank money, drawing demand away from the monetary assets of other agents, the CBDC has the potential to increase the central bank's control over monetary policy and the ability to reap seigniorage.

More subtly, however, the use of CBDCs might increases the velocity of money; some transactions are not only more convenient electronically, they are also quicker to achieve. The amount of time that money needs to stay in a person's possession between one transaction and the next falls. In other words, less money is needed in aggregate to achieve the same value of transactions.

When a nonbank private institution (e.g., fintech) issues electronic money, it acts as a substitute for the central bank's money, increasing the elasticity of demand for central bank money. However, to the extent that the issuing institution uses central bank reserves as backing for electronic money, it restores some of the power of the central bank to affect the liquidity premium—see Section V for liquidity outside the money base. An important factor that affects substitution between CBDC and privately issued electronic money is perceived safety. Increases in perceived risk of private money would be expected to cause significant swings in the mixture of payments assets held by the public.[6]

In the case of electronic (or digital) money issued by banks the considerations noted in the previous paragraph continue to apply; however, there are two additional distinctions. First the regulatory structure makes the money a closer substitute for central bank digital currency. Second, because these institutions are also lenders, changes in the costs they face can have a direct impact on their willingness to lend, and thus conceivably a stronger and more immediate effect than similar changes in costs of institutions which are solely in the business of providing payments.

Macroeconomic policy is intended to set an interest rate on investment that is consistent with the optimal level of aggregate economic activity in the economy. If a monetary aggregate is targeted, it should serve as a useful proxy of the size of the economy’s liquidity premium or the tightness of the supply of payments assets. It therefore becomes important to understand how monetary aggregates should be adjusted to take into account the introduction of new payments assets. The ideal aggregate would measure the total of all forms of payments assets, public and private which act as close substitutes. However, it may not be possible to observe all of these magnitudes. A central bank will have readily available information on CiC, as well as information on the size of transactions accounts. To the extent that it cannot measure the magnitude of transactions assets in unregulated institutions, the measure is imperfect. A partial remedy would arise from measuring the assets used as reserves by these institutions and adjusting by a money multiplier to move from base to aggregate.

III. New Instruments for Other Goals

Central Banks have multiple objectives; it is important to consider whether the introduction of new forms of payment arrangement require new instruments to achieve these multiple goals. Consider some objectives generally attributed to central banks: price stability; financial stability, and generation of government revenue

[6] In LICs or EMs, the issue of people wanting to transact outside the domestic banking system is more urgent and more legitimate and thus not restrict wallet holdings of CBDC; else, people will just hold other coins (not bank accounts). Thus, those licensed will issue CBDC freely/widely and collect seigniorage.

through seigniorage. Different forms of new money will yield different seigniorage revenues. Lender of last resort functions and associated interest rates and liquidity policies are instruments designed to improve financial sector stability. To the extent that the new payment facilities are separate sources of instability, it may be appropriate to have separate instruments targeted toward them. In either case, however, new instruments, only make sense to the extent that there is limited substitutability between the various payment sectors, both on the demand side of customers for the instruments, and on the provision side. If arguments for interest bearing (e.g., wholesale) CBDC are compelling, this new instrument will need to justify (and align with) the new objective.[7]

Highly liquid assets with associated risk are a source of financial instability. Historically, monetary policy handled the dual goals of financial stability and macroeconomic control by separate instruments: in the US for many years, macroeconomic stability was the province of open market operations, while financial stability was encouraged with discount window lending and deposit insurance, as well as safety and soundness regulation. As long as the institutions providing new forms of payment instruments are regulated and insured, this same division of work can continue without significant alteration. Different types of payments institutions might require different risk premia for deposit insurance, depending for example on the degree to which payments liabilities are backed by central bank reserves or liquid assets. The use of electronic payments might increase the speed of any bank run, necessitating more generous deposit insurance.

The main financial stability concern will arise from payments institutions which are unregulated or underregulated. As in the shadow banking crisis, the problem arises because the costs of regulation (whether small or large) lead institutions to engage in regulatory arbitrage, taking on systemic risks not fully borne by the institutions or the users of the institutions' services. Such problems could easily arise again if fintech payments institutions become widely popular. The most important protection against this, is to ensure that the benefits of joining the regulated sector are sufficiently great to offset the costs of the regulation. In the case of payments systems, for example, access to the payments backbone is an extremely large carrot, provided that the costs of regulations are reasonably adjusted to the risks imposed by the payments arrangement (for example, by allowing payments institutions which do not act as lenders to have regulation that is tailored to the payments function only).

When different payments arrangements arise, it becomes worthwhile to consider whether there is value in using central bank regulation to encourage a particular mix among the arrangements. For instance, seigniorage revenues will be greater from a CBDC than from the same amount of payment activity in a private payment institution, and the elasticity of demand for central bank reserves are likely to vary with the type of institution—see Box 2. However, the power of such fine-tuning will be limited both by the ease of substitution by customers between the different payment methods based on expense, by any swings in preferences among the monetary assets, based on perceived risk.

IV. Central Bank Operational Issues

The introduction of new payments platform by has great potential but may lead to new risks and in the long run, including possible changes in monetary policy transmission. The new services have the potential to change the

[7] New objectives include liquidity management via CBDC, nonbank access to wholesale CBDC, etc.

relationship between base money in the economy and capacity to carry out transactions, with consequences for central bank seigniorage and monetary policy transmission.

The benefits include extending services to the large unbanked cash-based segment of the population and the potential for rapid growth of these services is dramatic**.** There is anecdotal evidence that in the short time the payment service providers (PSP) have been available, mobile phones provide the foundation technology to exchange mobile money. While the systems are of convenience to individuals who already are sophisticated in the use of electronic arrangements, the truly transformational effect on the economy could come from the spread of these services throughout the countryside and to the unbanked segments of the population. This is especially true where fraction of individuals with bank accounts is very small, and the economy is cash-based. An elaborate system of payments organizations has arisen as a way for individuals to make remote payments to, for example, utility companies. On the other hand, the vast majority of unbanked population has mobile phones and the mobile business is concentrated in large operators. Thus the phone-based system could rapidly expand beyond the initial use case for such phone-based systems as Kenya's M-pesa, which initially were used primarily for transmission of cash from workers in the cities to relatives in the countryside.[8]

A. Decline in Monetary Base and Seigniorage

A higher money base allows for higher income from seigniorage; cross-country data shows that in countries with low levels of bank penetration, the ratio of money base to GDP is higher. Countries with high money demand (i.e., money base to GDP is above average) can sterilize relatively more than average.

However, the monetary base could potentially be reduced by a more widespread use of mobile operator payments systems. Mobile payments accounts are beneficial as they extend services to the large unbanked segment of the population. The introduction of mobile operator payments system reduces the size of the monetary base; it swaps part of cash under circulation to demand deposits, through the payment to the mobile operator. This first effect reduces monetary base as only the fraction of reserve requirement on the demand deposit is now part of the monetary base. Second, as a larger number of customers use the mobile operator to make payments, there is less incentive for the mobile operator to keep the whole amount as demand deposits, so there is a further reduction in monetary base as demand deposits and consequently bank reserves also decrease. This translates into an increase in the velocity of money, where more transactions in the economy can be paid by using a smaller amount of base money. This will entail adjustment of operational calculations for those CBs using monetary aggregates for regulation.

We present a illustrative scenario to show the potential costs of a reduction in the monetary base ratio if mobile payment systems substitute the use of cash—Figure 1. We present three scenarios where base money shrinks from 16 to 10 percent of GDP: (i) a severe scenario, where the reduction takes place in the next six years where MB/GDP falls by 1 percent per year; (ii) a baseline scenario in the next 12 years, where MB/GDP falls by ½ percent per year; and (iii) a mild scenario in the next 24 years where MB/GDP falls by ¼ percent per year.

[8] MNOs are generally required to maintain liquid assets equal to the amount of money issued electronically. The funds are usually pooled and held by a bank in the name of the MNO. This arrangement ensures a customer's money will be available on demand. Often, the only regulation of the mobile phone operators is by the communications authorities; and they regulate for technical standards of the communications, not for any financial or liquidity standards.

Figure 1. Seigniorage and Base Money Changes

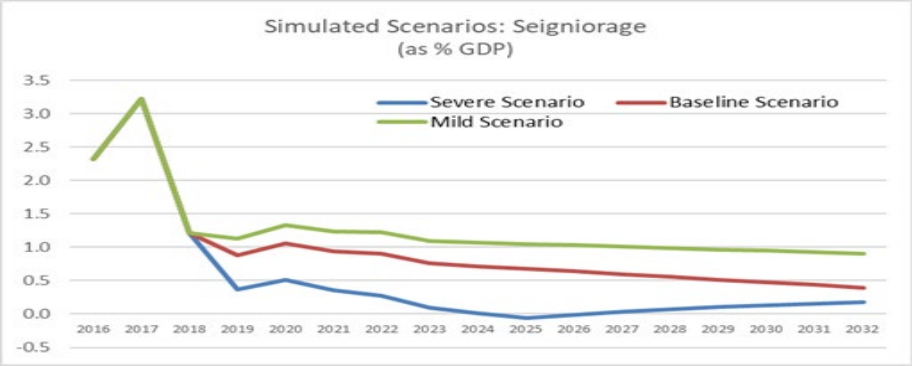

Source: IMF staff estimates; simulation.

Some central banks have adequate data to have a preliminary understanding of the use of e-wallets and associated velocity. They know the average holdings in e-wallets on particular dates (e.g., quarter-end or month-end). A recent study (technical assistance mission) finds volume using e-wallets for payment of goods and services (end-2018) have been 4.3 billion pesos; the average holdings or balance in e-wallets was about 265 million pesos (end-2018). This results in a "*transactions velocity*" of about 16. This provides a useful angle to understand how the interest rate sensitivity of these holdings compares with that for other money aggregates. The comparison, GDP/ M0, a standard metric and is roughly 3.0 as per monetary data files of this country.

A reduction of the monetary base in the future may somewhat constrain the ability of a CB to mop up excess liquidity and conduct monetary policy. A lower demand for cash reduces the rate of growth of the monetary base; this in turn reduces seigniorage revenues from money creation; reverse would be the case in a CBDC world (see Section II), if CiC increases base money (see Box 3).[9]

[9] See Reserve Bank of India speech by deputy-governor Rabi Sankar, "*Central Bank Digital Currency – Is This the Future of Money*" (July 2021), paragraph 31 on CBDC and seigniorage. https://www.rbi.org.in/Scripts/BS_SpeechesView.aspx?Id=1111

Box 3. Demand for Money and Seigniorage

Developed economies have a currency demand between 2 and 4 percent of GDP, while for most other countries it ranges between 2 and 11 percent of GDP. Many central banks remunerate 70-80 percent of central bank profits must be transferred to the Treasury. If the digital/fintech growth is sizable, base money/GDP will decline; so will seigniorage. If a country adopts CBDC, and there is disintermediation of the banking system deposits, base money may go up; so will seigniorage.

The traditional definition of seigniorage depends on both inflation ("tax inflation") and the level of demand for reserve money. In the short run, seigniorage also depends on changes in reserve money. An illustrative example where inflation is 6 percent, and Reserve Money as percentage of GDP is 16 percent would result in seigniorage revenue of 0.9 percent of GDP. The illustration presented below does not include the short-term effect where base money/GDP is constant between the two periods (i.e., it is zero in the equation). If base money declines, zero in the equation will become negative and seigniorage will be lower. If base money increases, seigniorage will increase.

$$s_t \equiv \frac{(H_t - H_{t-1})}{P_t Y_t} = (h_t - h_{t-1}) + \left(\frac{\pi_t}{1+\pi_t}\right) h_{t-1} = 0 + \left(\frac{6percent}{1+6percent}\right) * 16 \text{ percent} = 0.9 \text{ percent}$$

Where:

$H_t - H_{t-1}$: Flow of Reserve Money

$P_t Y_t$: Nominal GDP

$h_t = \frac{H_t}{Y_t P_t}$ Reserve Money as percentage of GDP

B. Currency in Circulation and Demand for Money

The new payments systems may represent a leakage in the transmission channels for monetary policy**.** Demand for cash depends on the alternatives available to cash. For instance, in economies where individuals are rapidly moving away from the cash economy into banking services, we expect to see the demand for cash falling relative to the demand for bank accounts. In economies where nonbank alternatives to cash are increasing, we expect a decrease in the ratio of cash outstanding to GDP, while the effect on broader monetary aggregates will depend on the degree to which reserves are held against the new money substitutes. For instance, in jurisdictions where regulations require holding reserves one for one against e-moneys, movement from cash to e-money will have no effect on broader aggregates (e.g., M2 in Kenya), while movement from bank deposits to e-money will reduce broader aggregates (e.g., Kyrgyz Republic) if there are no requirements to reserves against e-moneys.

The graphs in Figure 2 show the reduction trend in the ratio of CiC to GDP in a group of countries, where the fall in demand for cash is potentially related to moving away from the cash economy to nonbank alternatives to cash.[10] More importantly, the country teams acknowledge digital forays that may/maynot map fully into the trend lines (e.g., Finland). These countries are very heterogeneous, including developed economies such as

[10] We removed all dollarization cases (Nigeria, Angola, Argentina, etc.) that would reduce CiC. Regressions were not a way out as dollarization, CiC, M0, etc. are not easy to control simultaneously. For example, "dummy" to "control for dollarization" is not a good route, as some cases dollarization doubles; or is flat, or increases marginally.

Sweden and Norway, major emerging economies such as China, India, and Russia, and developing economies as Mongolia, Kenya, Armenia, Kyrgyz Republic, Kazakhstan, Tanzania, and Rwanda.[11]

Generally, CiC decline usually pulls M0 with it; however the behavior of M0/GDP may not have declined due to financial deepening, e.g., Russia, Mongolia, Armenia, and Rawanda (i.e., larger banking sector, and thus more required reserves that contributes to M0 and thus higher seigniorage).

Figure 2. Currency in Circulation and Monetary Base—Trends (2007-2020)

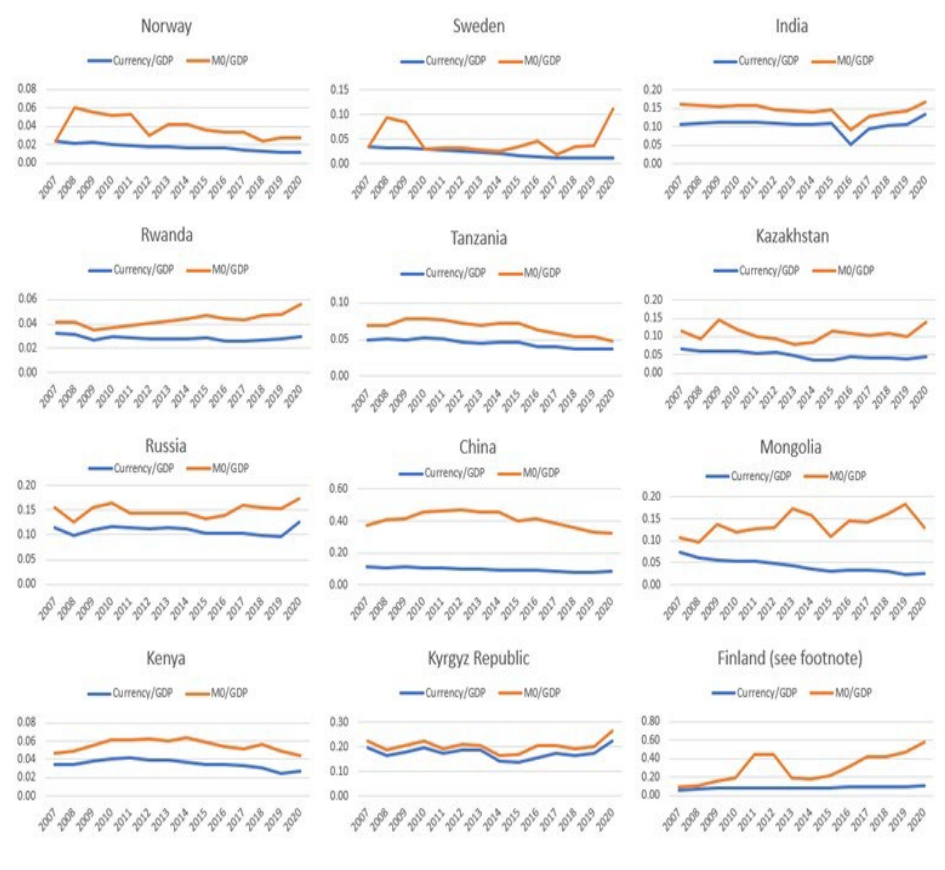

Source: IMF staff estimates.
F/N (Finland): *Beginning January 2002, the reporting of currency in circulation is determined by the accounting provisions of the ECB on the issue of euro banknotes' so "banknotes in circulation" on the NCB's balance sheets are not critical for analytical purposes.*

[11] As per discussions with country teams at the IMF. There are more countries that exhibit digital money forays (e.g., Korea etc.); however related issues, example hoarding of "yellow notes" in Korea masks the fintech progress viewed from the lens of CiC and M0.

Recent research using "cash usage" metric also suggests declining demand for cash (see Khiaonarong and Humphrey, 2022). The metric is developed using BIS's granular data on 25 countries that includes cash, credit cards and e-money: the metric is (cash/cash+cards+e-money). Although the metric cannot be used for seigniorage calculation, it is a useful harbinger of where CiC may be trending. Although seemingly a detour, one of the preconditions for implementing a full-fledged inflation-targeting regime is the absence of fiscal dominance. This means that the government should take central bank's profit as an exogenous variable and let the central bank to run an independent monetary policy consistent with its legal mandate to preserve price stability and extent of sterilization (see Box 4). Under a fiscal dominance situation, the government could induce the central bank to increase its transfers to the Treasury beyond a level consistent with its macroeconomic goals in order to meet budgetary needs. Under such circumstances, the central bank may be unable to secure a stable and permanent low inflation. As a result, society does not trust the inter-temporal purchasing power of the domestic currency, and it is not possible to anchor inflation expectations.

Box 4. Demand for Money and Sterilization

If the demand for money as percentage of GDP is high and the return of the net foreign assets (NFA) is high, sterilization can be absorbed within the central bank balance sheet more easily. In order to illustrate this simple arithmetic, let us assume a steady state situation where the balance sheet of the central bank does not growth as percentage of nominal GDP. We further assume that the return on net foreign assets is r_t^*. For simplicity we use the following notation for the analytical derivation: $\frac{NFA}{PY} = \gamma$, $\frac{NDA}{PY} = \chi$, $\frac{MB}{PY} = \lambda$ and $\frac{N}{PY} = \eta$; NFA: central bank net foreign assets, NDA: central bank net domestic assets, MB: monetary base, N: Central bank net worth, and PY: nominal GDP. In steady state, the return of central bank net worth has to be zero to ensure its balance sheet does not keep growing or shrinking. This implies that $(1 + r^*)\gamma - (1 + i)\chi - \lambda = \eta$ and using the identity, $\gamma - \chi - \lambda = \eta$, we can express the central bank sterilization cost in steady-state, $i\chi$, in the terms of the following simple equation:

$$i\chi = (\lambda + \eta)r^*$$

This simple equation shows that the sterilization costs that the central bank balance sheet can absorb in the long run. Using the average for emerging market economies, where λ *or* money base to nominal GDP $= 10$ percent, and further assuming a return on NFA of 3 percent in the long run, the steady-state, in this example sterilization cost is estimated in 0.30 percent of GDP. (Note η is zero in the equation, as we assume central bank's net worth does not change; level variable].[1/] Restricting sterilization due to demands for budgetary needs to be below 0.30 (in this example) will adversely impact the conduct of monetary policy. Thus, the inroads of fintech/digital money and if demand for money (or λ) decreases, constraints to sterilization are possible.

[1/] A constant monetary base in the long run converges to a value of λr^*.

V. Payment Service Providers and Demand for Money

Electronic money (or e-money) may be regulated as discussed in earlier sections (e.g., CBDC, stablecoins backed by reserves, quasi-CBDC issued by banks or nonbank within regulatory perimeter etc.) and will be part

of the central bank balance sheet.[12] However some e-money may be unregulated which is the focus of this section (IMF, 2021).[13]

The willingness of individuals and companies in the economy to absorb the monetary base issued by the central bank depends on the degree to which individuals wish to hold cash and the degree to which customers wish to hold demand deposits in banks (plus the rules by which banks hold monetary reserves against customers' bank deposits). Since banks hold reserves which are a fraction of demand deposits, while cash holdings are one-for-one central bank money, a change in consumer preferences for cash relative to bank accounts will change underlying demand for central bank reserves. If new payments systems are effectively subject to lower reserve requirements than traditional banks, then demand for central bank monetary base (especially CiC) further deteriorates.

In theory, availability of non-bank private PSPs like mobile network operators will both reduce the level of seigniorage and the effectiveness of transmission of monetary policy.[14] However, the empirical evidence on this issue is tentative. In part this follows because the most dramatic innovations are relatively recent, and it is difficult to interpret difference in money demand equations over long periods of time.

As customers make payments to other companies, the mobile operator will in effect use its bank accounts to make the payments, so balances in customer accounts and in the mobile operator's bank account will decrease one for one. But over time different individuals are topping off their accounts at the same time that others are utilizing theirs, so that these totals will remain fixed on average. For this reason, the mobile operator has no need to hold all its balances in a low interest demand deposit. It could, instead choose to invest in long term financial assets outside the banking system entirely. In this case, the central bank will be concerned, not only with the loss of seigniorage but also with the possibility of financial instability as in the classic Diamond Dybvig (1983) bank run model, since the mobile operator lacks the liquidity needed to honor the total demands of all customers should they decide simultaneously to use their funds for payments.

The key point is the asset side of non-bank private issuers of new money. If they keep 100 percent as banking deposits, no major change in the transmission mechanism and supply of banking loans should be expected. On the contrary, if they keep other assets such as treasury bonds or other financial assets, then it weakens the supply of credit from the banking system, assuming the banking current system follows a structure of both financial intermediation and maturities transformation. The impact on transmission will be more if non-bank payment institutions follow a more narrow-banking approach with assets different from the banking system.

Liquidity outside the monetary base is important to recognize.[15] The importance of these interactions will determine the extent to which the new arrangements reduce the effectiveness of monetary policy, as shown in sub-section A. Recent experience in some African countries requires phone companies to hold liquid reserves against the funds (though possibly not bank balances) that are in customers' accounts. In other countries, there is no such requirement.

[12] As the title of the paper suggests, digital money is a broader concept, and we use e-money as a sub-set of digital money.
[13] Regulation does not necessarily imply regulation as banks, but some regulation to bring them within the regulatory perimeter
[14] In some countries, where PSPs are mandated to hold central bank reserves on their "float," liquidity outside monetary base will not change.
[15] CiC (currency in circulation) is generally close to M0 in EMs and LICs but where financial deepening has been fast and significant (but not quantitative easing), then M0 is a more complete metric for seigniorage calculation as contribution from required reserves may be sizable along with CiC.

A. Liquidity Outside Monetary Base via Mobile Payments—An Illustrative Example

Assume a mobile operators' customer prepay for services. Assume the customers pay 200 to receive services. The mobile phone company uses 100 of this to invest in its infrastructure and puts the rest in bank deposits. Over time as customers receive services from the phone company, the balance in the customer accounts reduces and the net worth of the phone company increases correspondingly as its liabilities decrease, but assets are unaffected.

Mobile Operator Phone Company

Assets	Liabilities
100 Infrastructure	200 Customer Accounts
100 Demand Deposits at Commercial Bank	

In the country as a whole, there are a large number of unbanked individuals. They hold a total of 1,000 in currency in circulation, and 100 in the form of accounts with mobile phone operators for services. There are also a number of banked individuals; they hold 3,000 in bank deposits, 1,000 in currency in circulation, and 100 in mobile phone accounts. Thus, the commercial bank has deposits equal to 3,100. Assume the reserve requirement is 50 percent. The commercial bank's balance sheet is as follows:

Commercial Bank

Assets	Liabilities
1550 Commercial Loans	3000 Individuals' Demand Deposits
1550 Reserves at Central Bank	100 Phone Company Demand Deposits

Thus, the central bank has an outstanding monetary base of 1,550 reserves of commercial banks plus 2000 in currency in circulation. This 3,550 is the source of the central bank's seigniorage.

Central Bank

Assets	Liabilities
3550 Interest Bearing Financial Assets	2000 Currency in Circulation
	1550 reserves

Now suppose that the phone company introduces a facility which allows unbanked individuals to make payments through their accounts. Because of the convenience of this arrangement, individuals increase their holdings of balances with the phone company from 100 to 200. For the phone company, this increase does not represent an increase in demand for phone company services, so it makes no sense to make further investment in infrastructure. These balances could be held as additional balances in bank deposits

Mobile Operator Phone Company

Assets	Liabilities
100 Infrastructures	300 Customer Accounts
200 Deposits with Commercial Bank	

The commercial bank finds its deposits increasing—it gains 100 from the mobile phone company. On net the central bank, however, finds there is a lower demand for its monetary base, as, in aggregate, usage has switched from CiC to reserve backed transactions. Total demand for these assets is the same, but it is more

concentrated in the fractionally backed component. Outstanding monetary base reduces to 3,500 in total; 1,600 reserves in commercial banks, plus 1,900 of currency in circulation. As we have described it so far, the effect of introducing the phone company payments accounts is the same as the effect in developing countries of an increase in banking penetration. (Note: If the mobile phone company holds 100 percent of treasury bonds, then the amount of demand deposits is reduced; the demand for collateral and/or central bank reserves would go up.)[16]

B. Payment Service Providers (and Mobile Payments) Interface with Banking

Careful examination is required of the potential use of the new services for bypassing existing channels for international remittances, and whether this is desirable. International remittances are 20-35 percent of GDP for many countries (El Salvador, Tajikistan, Serbia, Armenia, Philippines, etc.). Remittances are currently made through the banking system and the transmission process is apparently efficient. However, many remittances are via phone account payments systems and the cross-border flows are sizable and increasing.[17] It therefore becomes important to understand the extent to which these services can currently (or have the potential) to be used to make international remittances outside of the banking system; M2 metrics are incomplete if sizable payments are outside the banking system; base money decline is also being observed If LICs target money aggregates (and maybe on way to inflation-targeting), monetary aggregates continue remain important.

Payment services are a fundamental portion of the financial industry and are highly regulated because of their potential risks. Nonetheless, e-money and the mobile phone accounts in the new arrangement are effectively the equivalent of demand deposits, and as such subject to the same concerns. It is imperative that the central bank move quickly to bring a regulatory umbrella over these services. So far, phone company accounts have not grown to a level significant enough to have any economy-wide effect, but this could rapidly change especially in remittances receiving countries where payment activities maybe outside the banking system (and result in incomplete M2).

The new services put tremendous competitive pressure on the existing payment and banking systems; it will be necessary to reconsider where unnecessary regulatory burdens can be relaxed, while encouraging them to develop their own innovations. The mobile phone operators have enormous customer bases compared to the banks. They have expertise in customer service and platform design. They have low regulatory burdens. The banks will find it extremely difficult to induce unbanked customers away from a phone company account. On the other hand, for the time being the banks have natural advantages over the phone companies in offering banking services to customers who already have bank accounts with them: deposit insurance, links to existing savings accounts, and the inertia of moving to a different payments' platform. These natural advantages can reinforce their links with their existing customer base, provided they develop improved banking services.

VI. Conclusion

The rapid development of new methods of payments has made for enormous benefits throughout the world, and in particular has radically changed the situation for individuals who were previously disconnected from the

[16] This would be the case if nonbank stablecoin issuers would be allowed to back their coins with central bank reserves; demand for central bank reserves would go up—see Box 2.

[17] For example, in Central Asia and Caucasus region. Russia in particular exercises stronger controls on access to SIM cards than countries that receive remittances (e.g., Armenia, Tajikistan and Kyrgyz Republic) from Russia.

modern financial system. Further encouragement of the development and expansion of the role of these new systems is an imperative for all central banks. But this development will pose new challenges for those in charge of implementing monetary policy.

The focus will be on the interest rate channel and if digital money (especially CBDC) will change monetary policy and central bank operations. New policy instruments make sense only to the extent that there is limited substitutability between the various payment sectors. Trends in currency-in-circulation, and their impact on central bank's seigniorage, monetary base, liquidity outside the monetary base, and transactional velocity will need to be understand better. As in the case of problems that have arisen from dollarization of deposits, or from new liquidity provision through shadow banking, effective regulation and policy making require understanding and readjustment. If anything, the new arrangements in payments are likely to be adopted even more rapidly than those earlier examples. Effective responses in the new environment will require careful monitoring of the demand for these new technologies and the factors that affect that demand.[18] Data gathering should begin now, before the changes in payments practices become overwhelming.

[18] It will be interesting to see how retail CBDC will compete with digital money use at household level. The MNOs are ambitious, with goals of extending their reach much more broadly into payments services, into microlending, and most significantly into foreign remittances. Like unidentified e-wallets (issued by banks), mobile account payments are subject to a variety of restrictions primarily designed for AML protection, including limits on individual payments, use for foreign transactions, and cash withdrawal.

References

Adrian, Tobias and Tommaso Mancini-Griffoli (2021), *A New Era of Digital Money*, Finance and Development, https://www.imf.org/external/pubs/ft/fandd/2021/06/online/digital-money-new-era-adrian-mancini-griffoli.htm.

_____(2019), The Rise of Digital Money, IMF Fintech Notes, June 15.

Bech, Morten and Rod Garratt (2017), Central Bank Cryptocurrencies, BIS Quarterly Review, September.

Brainard, Lael (2022), *Preparing for the Financial System of the Future, Speech* at the 2022 U.S. Monetary Policy Forum, New York, New York, February 18 at Federal Reserve Board website.

Carney, Mark (2021), *The Art of Central Banking in a Centrifugal World,* Andrew Crockett Memorial Lecture at the BIS, June 28, https://www.bis.org/events/acrockett_2021_speech.pdf.

Carstens, Agustin (2021), *Digital Currencies and The Future of The Monetary System*, Speech at the Hoover Institution policy seminar, Basel, January 27 at BIS website.

Diamond, Douglas W., and Philip H. Dybvig (1983), *Bank Runs, Deposit Insurance, and Liquidity*, Journal of Political Economy, Vol. 91, No. 3, pp. 401–19.

Gorton, Gary and Jeffrey Zhang (2021), *Taming Wildcat Stablecoins,* University of Chicago Law Review, Forthcoming.

International Monetary Fund, 2021, E-Money- Prudential Supervision, Oversight,and User Protection, DP/2021/027, December.

Kahn, Charles, Francisco Rivadeneyra and Tsz-Nga Wong (2020), *Should the Central Bank Issue E-Money?,* Journal of Financial Market Infrastructures, Vol. 8, No. 4, pp. 1–22.

_____and Manmohan Singh (2021), *If Stablecoins are Money, They Need to be Backed by Reserves,* Comment at the RISK.net, February 10, https://www.risk.net/comment/7744611/if-stablecoins-are-money-they-should-be-backed-by-reserves.

Khiaonarong, Tanai and David Humphrey (2022,) *Falling Use of Cash and Demand for Retail Central Bank Digital Currency*, IMF Working Paper 2022/027.

McLaughlin, Toni (2020), *The Regulated Liability Network*, Citibank whitepaper.

Quarles, Randal K. (2021), *Parachute Pants and Central Bank Money*, Speech at the 113th Annual Utah Bankers Association Convention, Sun Valley, Idaho, June 28.

Sanches, Daniel and Todd Keister (2021), *Should Central Banks Issue Digital Currency?*, Federal Reserve Bank of Philadelphia Working Paper 21-37.

Sankar, T. Rabi (2021), *Central Bank Digital Currency – Is This the Future of Money*, Keynote address at the webinar by the Vidhi Centre for Legal Policy, New Delhi, July 22 at Reserve Bank of India.

Singh, Manmohan, Apoorv Bhargava and Peter Stella, 2021, The New Era of Money Supply and Its Impact on Policy. CentralBanking.com (May 17). The New Era of Money Supply and Its Impact on Policy - Central Banking.

_____, Charles Kahn and Caitlin Long, (2021), Interoperability of Stablecoins - Central Banking.

PUBLICATIONS

Agent-Based Simulation of Central Bank Digital Currencies

Agent-Based Simulation of Central Bank Digital Currencies*

Amanah Ramadiah[1†], Marco Galbiati[2], and Kimmo Soramäki[1]

[1] Financial Network Analytics Ltd

[2] S&P Global

November 9, 2021

Abstract

This paper presents a multi-period agent-based model for the study of macro-financial effects related to the introduction of a retail Central Bank Digital Currency (CBDC). Calibrating it with aggregate statistics of the German retail payment market, we exemplify how the model can be used to quantify the impact of a CBDC on i) the usage of alternative means of payments, ii) the composition of consumer's wealth, and iii) the banking sector disintermediation. We find that CBDC can be configured without largely impacting the banking sector balance sheet. However, we also find that card companies may suffer a substantial decline in their transaction revenues. We see this model as a framework that can be enriched and tuned to answer a myriad of questions relevant to different jurisdictions from a macro-financial angle. The model is publicly available in the FNA simulation platform for running other policy experiments i.e., testing the efficacy of alternative configurations of CBDCs.

Keywords: CBDC, payment systems, agent-based model, economic impact, disintermediation

JEL Classification: C63, E41, G21, G28

*We thank participants of the Central Bank Research Association 2021 Annual Meeting and the 25th Workshop of Economics with Heterogeneous Interacting Agents. We are grateful for valuable comments from Mauro Gallegati, Rod Garrat, Jonas Gross, Serafin Jaramillo and Zijian Wang.

†Corresponding author. Email: amanah@fna.fi

1 Introduction

Central banks around the world are actively exploring central bank digital currencies (CBDCs). A recent survey from Boar and Wehrli (2021) finds that more than 80% of the 65 central banks surveyed are undertaking extensive work on CBDCs. The survey also points out that, in the next three years, central banks representing about 20% of the world's population are expected to issue a retail (general purpose) CBDC. The implementation of a retail CBDC has the potential to fundamentally reshape our current monetary and financial systems, with implications not yet fully understood and currently actively researched. The Federal Reserve Bank of Boston, for example, is undertaking work on retail CBDC research with the MIT Digital Currency Initiative (Federal Reserve Bank of Boston, 2020). Additionally, as reported in Reuters (2021), European Central Bank (ECB) policymaker Ignazio Visco has stated that ECB is exploring ways to launch a digital euro. At the same time, a live CBDC has been fully deployed. The Central Bank of The Bahamas launches its Sand Dollar to the general public in October 2020 (Central Bank of the Bahamas, 2019).

Despite the rising interest, many central banks appear not yet convinced that the benefits of CBDC issuance will predominate the costs (Barontini and Holden, 2019). One of the major roadblocks is the fact that the implications of CBDC on the financial and economic system, as well as the behavior of players within a CBDC ecosystem, remain unknown. Accordingly, the literature on CBDC has been growing rapidly for the past few years (see e.g., Kiff et al. (2020) for the recent survey). Some, for example, have studied the possibility of different CBDC designs and discussed various approaches to cope with the banking disintermediation. Bindseil (2021) and Panetta (2018) proposes imposing holding limits, while Kumhof and Noone (2018) suggest a more progressive proposal that would restrict on-demand convertibility of deposits into CBDC. Gross and Schiller (2021) shows that, unless central banks decide to provide additional funds, CBDC will crowd out bank deposits. Others focus on the CBDC implication on systemic risk (Fernández-Villaverde et al., 2020) and its impact on monetary policy (Keister and Monnet, 2020).

This papers offers, to the best of our knowledge, the first Agent-Based model of an economy with a CBDC. While a number of features are still simplistic, our model already includes several key ingredients needed to discuss the main issues related to a CBDC. Its modular structure makes it amenable to extensions and modifications, making it a ***tool for enquiry & a modelling platform*** for a "user" to run policy experiments. Some of the issues that can already be looked into in this stylized framework are: CBDC adoption, disintermediation, leverage of the banking sector (and thus, by extension, financial stability), conduct of monetary policy. At the moment, agents and their interactions are highly simplified, while complex concepts are reduced to simple variables or ratios. However, as we rely on simulations as opposed to numerical solutions, nothing will stop us from adding parts to it, to make it more realistic and capable to address additional questions related to a CBDC. For now, we show what we can obtain with this minimal structure; future expansions will be relatively easy, exploiting a number of "plugs" that we embedded in the model in a long-term development view.

The remainder of the paper is organised as follows: in Section 2 we provide an overview of prior research in payment simulation and agent-based simulation. We present the details of the agent based model in section 3, and describe the parameter calibration in section 4. We then present and discuss the results in section 5. Finally, we discuss our conclusions and further works in section 6.

2 Relevant Literature

The literature on CBDC is in constant and rapid evolution. So, instead of attempting a review that would necessarily be incomplete and soon become outdated, we give a short overview of two families of studies this paper draws on: the literature on payments and the one in agent-based modeling.

2.1 Payment Simulation

As discussed in Leinonen and Soramäki (2004), simulation techniques allow one to build models that closely imitate the real world. This approach is particularly useful when generic econometric models are limited to deal in complex economic or financial systems. In this respect, simulation techniques have long enjoyed success in modelling payment and settlement systems. Koponen and Soramäki (1998) present the first Payment Simulator to assess liquidity impact of Finnish banks when joining the Pan-European interbank payment system TARGET. Simulation approaches are well established for understanding large-value payments systems probably since the regulatory approval process of the CLS system for settling FX trades, as simulations were extensively used to understand behaviour of the system both in normal and abnormal circumstances. Today, the majority of systemically important FMIs employ simulators to stress test and validate changes to their systems as well as to guide them on design choices.

In a study published by the Bank of Finland, for example, Hellqvist and Koskinen (2005) apply a stress testing simulation to the Finnish Bond clearing and settlement system. Furthermore, in Soramäki et al. (2007) and Beyeler et al. (2007), the approach was joined with network theory to understand the topology of the Fedwire system and to simulate failures in critical financial infrastructures. The Eurosystem has also recently embraced payment system simulations as an ongoing oversight tool by specifying how the transaction level data may be used (EU, 2010) and developing a TARGET2 simulation platform. Another recent example is the project to develop new features for CHAPS interbank payment system in the UK. McLafferty and Edward (2013) use real payment data to quantify the liquidity efficiency that could be obtained in CHAPS, the UK's large-value payment system, by the implementation of a liquidity saving mechanism. A more recent example is Byck and Heijmans (2021) who study liquidity in the Canadian large value payment system. The simulation of retail payment systems or the cash cycle have not been studied to the same extent. However, agent-based modelling (ABM) approaches have been deployed to understand these more granular systems.

2.2 Agent-Based Modelling

Agent-based modelling is a computational modelling paradigm that enables the description of how agents behave and interact in a system. The methodology of ABM encodes the behavior of individual agents in simple rules so that we can observe the emergent results of these agents' interactions (Wilensky and Rand, 2015). Figure 1 shows the schematic representation of the typical elements of ABM. There are a number of agents in an environment who are interacting with each other and also with the environment.

Fig. 1: Schematic representation of an agent-based model (ABM).

An ABM is suited to modelling complex systems such as the economy, where complexity, heterogeneity, networks and heuristics combine to produce emergent behaviour. ABMs have already delivered strong results on different applications, such as financial system stress-test (Farmer et al., 2020), funding risk (Hałaj, 2018), housing market (Geanakoplos et al., 2012), payment system (Galbiati and Soramäki, 2011) and financial market simulation (Markose et al., 2005). As described in Turrell (2016), the characteristics of ABM may make it a well-suited approach for exploring the impact of different possible CBDC specifications.

3 An Agent-Based Model of CBDC

This section lays out a parsimonious set of building blocks to investigate recurring issues in the CBDC debate: i) digital currency adoption, ii) credit and disintermediation, iii) bank leverage, iv) monetary policy. We aim at a minumum sufficient structure i.e. we try to introduce only what is strictly needed, in terms of agents, assets, prices, choices, rules, to model these issues. We take inspiration from existing models, from which we borrow the gist of the argument rather than assumptions and formalism, and assemble these modelling blocks into a simple, modular model, with a view to scale it up in the future.

The model is cast in discrete time, with each t ($t = 1, 2...$) representing a "day". The model features 4 classes of agents: consumers, merchants, a commercial bank, a central bank. There is also an underlying 'economy', represented by risky projects undertaken by the commercial bank. Agents interact in time according to given rules, generating dynamics for the model's key outputs: i) composition of consumer's wealth (allocated between cash, CBDC, financial assets), ii) diffusion of means of payments and thus percentage of transactions settled via cash, cards and CBDC, iii) bank deposits, iv) interest rates, iv) the size of the economy.

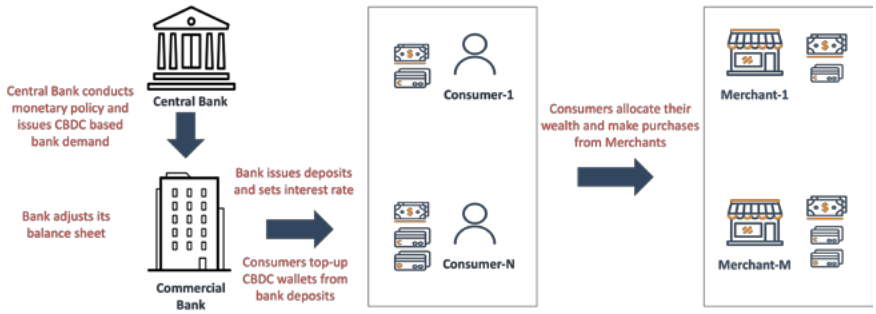

Fig. 2: Illustration of the model dynamics.

As summarizes in Figure 2, the consumers make daily purchases from the merchants using either cash, deposit (card) or CBDC. The decision regarding the means of payment is based on both the consumers' wealth allocation and the merchant's acceptance.[1] Moreover, we assume that the commercial bank adjust its balance sheet by borrowing CBDC from the central bank to face consumers' withdrawals and to fund risky projects. Meanwhile, the central bank conducts monetary policy (e.g., sets the maximum allowed CBDC balance and determines the rate of CBDC borrowing) and issues CBDC based bank demand. The following subsections describe in more detail the 4 classes of agents and their actions.

3.1 Consumers

There is a large, fixed number of consumers, indexed by $i = 1...N_C$. At any given day t, each consumer i has overall wealth $W_i(t)$ allocated between i) cash, ii) CBDC, iii) an interest-bearing current account, or deposit, used for payments (say via debit or credit card), and iv) an otherwise unspecified asset. The latter is not liquid i.e. cannot be used for payments before being liquidated into the current account, and is be used as store of value. In sum,

$$W_i(t) = C_i(t) + K_i(t) + B_i(t) + A_i(t) \tag{1}$$

[1] The commercial bank does not charge any usage fees for both CBDC and card transactions. When transaction fees exist, there would be some bargaining powers between consumers and merchants (Bolt and Soramäki, 2008).

where W is wealth, C is cash, K is CBDC, B is bank deposit and A is illiquid asset. At each t, every i does two things: she decides on the allocation of her wealth and she makes purchases.

Purchases & payments Daily purchases are represented by a weighted bipartite network of N_c consumers and N_m retailers. Randomly drawn at each t, this network is denoted by:

$$\Pi(t) = \left\{ \{p^{\alpha}_{i,j}(t)\}_{j \in M^i} \right\}_{i=1..N_c} \tag{2}$$

where $p^{\alpha}_{i,j}$ is the value of the α-th purchase that i makes from retailer j and M^i is the set of retailers from which i makes purchases. We build $\Pi(.)$ in steps[2]:

1. Draw a total number of purchases in the network: $\eta(t) \sim Poisson$
2. For each purchase, draw a pair of consumer i (from the set of all consumers) and merchant j (from the set of all merchants)
3. Randomly draw the purchases' size as $log(p^{\alpha}_{i,j}) \sim Normal$ [truncated] for all i, j and α

Each of the η purchases is either 'online' or 'offline', according to independent random draws with probabilities p_{online} and $1 - p_{online}$. Only 'offline' purchases can be settled in cash. The settlement of each purchase takes place according to the following rule (where overlap means a set of liquid assets of which the consumer has enough for the purchase and are accepted by the merchant):

$$\text{Online} \begin{cases} \text{overlap is } \emptyset \begin{cases} \text{merchant accepts CBDC \&} \\ \text{consumer has CBDC wallet:} & \textbf{CBDC} \\ \text{otherwise:} & \textbf{no purchase} \end{cases} \\ \text{otherwise : } \textbf{random draw} \end{cases}$$

$$\text{Offline} \begin{cases} \text{overlap is } \emptyset : \textbf{no purchase} \\ \text{otherwise : } \textbf{random draw} \end{cases}$$

Let first consider an online purchase. If the consumer has not enough B nor CBDC to pay for the purchase, but CBDC is accepted by the merchant, the consumer immediately 'tops up' its CBDC wallet and uses that. Otherwise the purchase fails. Topups, however, are possible only if the consumer holds a CBDC wallet. Initially, no consumer holds one but, at each t, those who do not have a wallet will acquire one with probability:

$$prob_i(openWallet(t)) = g(w_i(t-1), y_j(t-1)),$$

[2]The network is exogenous. However, we could make it part of the adoption 'story' assuming a different network formation rule. For example, retailers may decide to offer a number of payment options, in order to attract consumers. And vice-versa, consumers would decide purchases (also) on the basis of the payment means accepted by retailers.

where $w_i(t)$ and $y_j(t)$ (resp. $y_j(t)$) is the percentage of consumers (merchants) who hold (accept) CBDC.[3]

Let's now look at the offline purchase case. If no payment option is available, the purchase fails. Otherwise, any available liquid asset (C, K or B) is used according to a random draw. Online and offline purchases thus differ because: i) offline purchases may be settled with any liquid asset, while online ones purchases admit B or K but not C; ii) a consumer may top up her CBDC wallet 'on the spot' for online purchases, but not for offline ones.

The total amounts paid with the three liquid assets in a day are indicated respectively by $\check{K}, \check{B}, \check{C}$, whose sum $\check{P}$ may turn out to be less than $\sum_j p^{\alpha}_{i,j}$ as purchases will fail if settlement is not possible (when e.g., the consumer runs out of cash or CBDC cannot be used).

Wealth allocation At $t = 0$, no consumer holds CBDC. C, B and A are attributed to consumers according to an observed empirical distribution:

$$W_{.}(0) = C_{.}(0) + B_{.}(0) + A_{.}(0), \quad [C, B, A] \sim f(.) \tag{3}$$

In this section, we omit the consumer's index as this is not needed and all the equations hold in the aggregate too.

At each $t > 0$, before purchases are made, consumers move part of their B into CBDC ('top-up'), provided that both: i) they hold a wallet (as they have encountered merchants using CBDC)[4]

At each $t > 0$, before purchases are made, consumers move part of their B into CBDC ('top-up'), provided that both: i) they hold a wallet (as they have encountered merchants using CBDC)[5], and ii) at $t-1$ they found themselves constrained in CBDC. The top-up is set to τ times the daily average CBDC payments *that the consumer could have made,* up to t. The CBDC top-up is also limited by design to a maximum β so the consumer's CBDC balance evolves as

$$K(t+1) - K(t) = \Delta K(t) = \begin{cases} \min\left[\tau \mathring{K}, \beta\right] - \check{K}(t) & \text{if there is a top-up} \\ -\check{K}(t) & \text{otherwise} \end{cases} \tag{4}$$

where $\mathring{K}$ is average daily amount of CBDC payments the consumer could have made. We have $\mathring{K} \geqslant mean(\check{K})$ but, if the consumer holds enough CBDC as to never be constrained, the two number come close as time goes by.

The consumer buys and sells A using the bank account, thus shifting wealth between A and B. This latter can be used for payments, but A yields a fixed r_A (paid into the bank account) possibly higher than the rate paid on deposits r_B. We do not explicitly model consumer preferences but assume that the consumer targets a certain

[3] This contributes to the CBDC adoption story. For example, the adoption rate is higher the larger the proportion of consumers and merchants in the system using CBDC.

[4] Note that the first time top up depends on a random draw.

[5] Note that the first time top up depends on a random draw.

ratio $\frac{A}{W}$, which depends on the spread in the returns offered by the illiquid asset and the bank account. We assume this target to be:

$$\left(\frac{A}{W}\right)^{*} = 1 - \frac{1}{r_A/r_B(t)} \tag{5}$$

and we posit a 'gradual & approximate' adjustment towards it. That is, the consumer chooses $A(t+1)$ so that:

$$\frac{A(t+1)}{W(t)} - \frac{A(t)}{W(t)} = -\frac{1}{\nu}\left(\frac{A(t)}{W(t)} - \left(\frac{A}{W}\right)^{*}\right). \tag{6}$$

The term 'approximate' refers to the fact that $\frac{A(t+1)}{W(t)}$ is an approximation to the ratio to be targeted $\frac{A(t+1)}{W(t+1)}$, and $\nu \geqslant 1$ makes the adjustment gradual. For example, if $\nu = 2$, the distance between the (approximate) current value and the target is halved. The target converges to 1 (i.e., all wealth is held in A) when $r_A/r_B(t) \to \infty$, and instead it converges to 0 as $r_B(t) \to r_A$, as indeed there is no reason to hold the non-liquid asset, if this does not offer superior returns).

The bank account B goes up as: i) returns from A are received, and ii) A is liquidated. It instead falls when: i) A is invested into, ii) CBDC is topped up, iii) cash is withdrawn, and iv) card payments are made. We also assume that each consumer receives an exogenous amount of salary, denoted by κ. So B changes as follows:

$$\begin{aligned} &B(t+1) - B(t) = \Delta B(t) = \\ &\quad [A(t)r_A + B(t)r_B] - \left[\Delta K(t) + \check{K}(t)\right] - \left[\Delta C(t) + \check{C}(t)\right] - \Delta A(t) - \check{B}(t) + \kappa \end{aligned} \tag{7}$$

where the terms in square brackets are respectively interest income, CBDC top-ups (see above Equation (4)), cash withdrawals (Equation (8) below), while $\Delta A(t)$ is the additional investment into A, and $\check{B}$ are card payments.

The last variable to be specified is C, for which we imagine similar choices as for CBDC: when she runs out of it, the consumer withdraw τ_c days worth of expected cash payments. That is:

$$C(t+1) - C(t) = \Delta C(t) = \begin{cases} \min\left[\tau_c \mathring{C}, \beta_c\right] - \check{C}(t) & \text{if there is a cash withdrawal} \\ -\check{C}(t) & \text{otherwise} \end{cases} \tag{8}$$

where the cash withdrawal, $\min\left[\tau_c \mathring{C}, \beta_c\right]$, equals $\Delta C(t) + \check{C}(t)$ when a withdrawal is made.

In summary, wealth grows due to the return on A and the interest paid on B. This, along with salary κ, allows a stream of purchases, modeled as a random networks. Payments are made by using liquid assets B, C, or K. Bank account B is also the source of cash (via withdrawals) and CBDC (via top-ups). It is also used to buy A, which can then be liquidated back into B. The demand for C and K (withdrawals and

top-ups) is a function of previous usage. The demand for A depends on the spread $r_A - r_B(t)$ and on $\frac{A}{W}$.

3.2 Merchants

All merchants accept cash and at $t = 0$ a percentage of them accepts card ($\mathscr{P}_B$) and CBDC payments ($\mathscr{P}_K$) too. As time goes by, each merchant start accepting CBDC at a random time, which depends on their customers' previous demands:

All merchants accept cash and at $t = 0$ a fraction $\mathscr{P}_B$ of them accepts card. A subset of these merchants, amounting to a fraction $\mathscr{P}_K < \mathscr{P}_B$ of the total, accepts CBDC payments too. As time goes by, each merchant start accepting CBDC at a random time, which depends on their customers' previous demands:

$$prob_j(accepts(t)) = h(z_j(t-1)) \tag{9}$$

where $z_j(t)$ is the percentage of consumers that visited merchant j and were able to pay via CBDC up to t, and h is some function such that $h(0) = 0$ and $h(1) \leq 1$. There is no feedback from the merchants' sales into the economy: from a modelling perspective, their purpose is only to contribute to a CBDC adoption story.

3.3 Commercial Bank

The bank does the following:

- Issues B to consumers. In particular, the bank sets rate r_B and the amount of deposits is determined by the consumers' demand (Equation (7))
- Borrows CBDC from the central bank (K_{bank}), paying rate r
- Uses the above funds to:
  - face consumers' withdrawals
  - fund risky projects X, which yield a daily random return $r_X(t)$ distributed according to a distribution Λ

We imagine the following **timing**: given $X(t), \sum_i B_i(t) = B(t)$ and $K_{bank}(t)$, the return $r_X(t)$ is drawn, the bank pays the due interest and time-t profit is computed. Then, the bank sets $r_B(t+1)$ to which consumers react choosing $B(t+1)$, causing the deposits inflow $\Delta B(t)$. To this, the bank responds setting $X(t+1)$ and $K_{bank}(t+1)$. And so on.

The bank's time-t profit is:

$$X(t)r_X(t) - B(t)r_B(t) - K_{bank}(t)r(t), \tag{10}$$

which we assume to be distributed to (un-modelled) shareholders. The bank faces two constraints. First, it cannot exceed a leverage ratio of $\bar{\Gamma}$ i.e., at all t, it must be:

$$\frac{debt}{assets - debt} = \frac{B(t) + K_{bank}(t)}{X(t) - [B(t) + K_{bank}(t)]} = \Gamma(t) \leqslant \bar{\Gamma} \tag{11}$$

where Γ denotes the bank's actual leverage, defined here as debt to equity. We assume that breaching the constraint entails a 'default', so the bank stops operating. We do not model what happens in this event, as this model is supposed to describe a normal functioning of the economy. We can, however, imagine that a breach of the constraint brings about a bank run, or special intervention from the central bank, and leave this for future investigation.

Second, the bank must continuously meet its obligations i.e., pay interest and satisfy bank withdrawals.[6] That is, at each time it needs at each time liquidity for an amount equal to $[B(t)r_B(t) + K_{bank}(t)r(t)] - \Delta B(t)$. To do so, the bank borrows new funds from the central bank (ΔK_{bank}) and/or liquidates X (ΔX) which, we assume, is illiquid in the sense that part of it is lost when liquidating it.[7] The liquidity constraint of the bank then is as follows:

$$[B(t)r_B(t) + K_{bank}(t)r(t)] - \Delta B(t) = \Delta K_{bank}(t) + \begin{cases} \Delta X(t)\xi & \text{if } \Delta X(t) < 0 \\ \Delta X(t) & \text{otherwise} \end{cases} \tag{12}$$

where $\xi \in (0, 1)$ indicates that, in order to obtain £1 of liquidity, the bank has to sell $£1/\xi > 1$ worth of X.

In the above equation, on the l.h.s is the bank's *liquidity need* and on the r.h.s is the amount of new funds. If the need is positive, funds are obtained by borrowing from the central bank and/or by changing X. If the liquidity need is negative, the bank can instead pay back to the central bank and/or invest into X.[8] As the liquidity need is fixed at t, the bank chooses $K(t+1)$ or $X(t+1)$, and the other quantity is determined by difference, according to Equation (12).

When choosing between X and K_{bank} in response to changes in B, the bank faces a trade off between profitability and leverage. When the bank chose to liquidate X, the difference between assets and debt (denominator in Equation (11)) is unchanged but debt (numerator) falls, thus the bank de-leverages. If the bank instead prefers to borrow K_{bank}, the bank substitutes central bank debt for consumer debt, so debt is unchanged. As assets are unchanged too, leverage thus remains constant. From a leverage perspective, liquidating X is thus preferable. However, from a profitability perspective, borrowing from the central bank is preferable: by substituting central bank debt for consumer deposits, the bank can preserve the lucrative asset X.[9]

[6] These include cash withdrawals and CBDC top-ups, as per Equation (7).

[7] We assume that the commercial bank holds no cash. It can either get it instantly from the central bank using CBDC reserves, or sell part of X against CBDC, which is then instantly converted into cash.

[8] So far, the bank can also increase X when facing a *positive* liquidity need, in which case it would borrow from the central bank to both meet the liquidity need and invest in X. And vice versa: it could decrease X even when not needed i.e when facing a negative liquidity need.

[9] We assume that r is lower than the return offered by X. This is true in the current low rates environment, and is probably also true historically, as X represents the overall portfolio of the financial

Given a path of central bank interest rates $r(t)$, the bank could maximize the expected discounted future stream of profits (Equation (10)) under constraints (Equation (11) and Equation (12)), choosing optimal paths of r_B and X or, alternatively, K. However, in the spirit of the ABM method, we assume here that the bank follows simple rules instead of precisely maximizing profits. As for the consumer, however, we imagine that the bank slowly adjusts its leverage ratio towards a target that depends on σ, the spread between the expected return on X and the average cost of debt: $\sigma(t) = \frac{\Lambda}{\frac{B(t)r_B + K_{bank}(t)r}{B(t)+K_{bank}(t)}}$. We assume this target to be:

$$\Gamma^*(t) = \bar{\Gamma}\left(1 - \frac{1}{\sigma(t)}\right),$$

so if σ drops below 1 (the expected return on X falls below the cost of financing it) the bank targets 0 leverage. If instead σ goes to infinity, the bank asymptotically targets the maximum allowed leverage $\bar{\Gamma}$ (beyond which it would 'default'). The target is approached again 'slowly and approximately' according to the following equation (where, to reduce clutter, we write K for K_{bank} and we omit time (t) except for $(t+1)$ variables so, e.g., B alone stands for $B(t)$):

$$\frac{B + K(t+1)}{X(t+1) - [B + K(t+1)]} - \frac{B+K}{X - [B+K]} = -\frac{1}{\mu}\left(\Gamma - \Gamma^*\right). \tag{13}$$

As for the consumers' equation in Equation (6), 'approximately' refers to the fact that the first ratio is an approximation of the *true* leverage to be targeted (which would have $t+1$ for B too): the bank myopically acts as if deposits at $t+1$ will be the same as at t. Parameter $\mu \geqslant 1$ again determines the speed of adjustment towards target.

Given history up to t, either $X(t+1)$ can be freely determined, or $K_{bank}(t+1)$, but not both, because of the cash flow constraint in Equation (12). Using then Equation (13), we write $K_{bank}(t+1)$ as a function of $X(t+1)$ (writing again K for K_{bank} and omitting time (t) except for $(t+1)$ variables):

$$K(t+1) = \frac{D}{1+D}X(t+1) - B \tag{14}$$

where $D = \frac{B+K}{X-[B+K]} - \frac{1}{\mu}\left(\Gamma - \Gamma^*\right)$. Plugging this into Equation (12), we obtain $X(t+1)$:

$$X(t+1) = \left[B(1+r_B) + K(1+r) + X\xi\right]\frac{1+D}{D + \xi(1+D)} \tag{15}$$

In reality, the above Equation (15) should have $B(t+1)$ in the place of B. However, we assume that at t the bank does not know $B(t+1)$[10], so $X(t+1)$ is set using $B(t)$ as an approximation.

Finally, regarding the choice of r_B, we assume that the bank applies a mark-up on

sector.

[10] After all, our single bank is a metaphor for the whole financial sector, so it is reasonable to assume that this does not have perfect knowledge of the demand for deposits.

the cost of borrowing CBDC, r, set by the central bank:

$$r_B(t) = r(t) + \epsilon \tag{16}$$

with $\epsilon > 0$ (and in the region of approx 1%).

3.4 Central Bank

The central bank has two traditional tools at its disposal: i) the max leverage $\bar{\Gamma}$ and ii) the cost of borrowing CBDC r. It also sets iii) the maximum CBDC top-up β_c and cash withdrawal β. With the first, the central bank tweaks the commercial bank's leverage constraint. With the second, it determines the cost of borrowing CBDC. Together, these two affects the commercial bank's choices, i.e., i) the amount of financed projects X and the commercial bank reserves K_{bank} (because the target leverage depends on the average cost of debt) and ii) the interest paid on deposits r_B, affecting in turn consumers' choices.

In theory we could explicitly model an objective function for the central bank and solve for a (game-theoretic) equilibrium between the commercial bank and the central bank, whereby each agent chooses mutually optimal actions. At this stage, however, the model will primarily be used to shed light onto the effects of central bank choices, so the central bank variables are left free, to be provided as an exogenous input. In order to close the model and provide some first results, the following rules will be used as baselines central bank policies:

1. Constant policy: $\bar{\Gamma}$, r constant in time, and set to levels which indicatively reflect current conditions. Limit β is also constant and set to typical maximum cash withdrawal level.
2. Growth-targeting rate policy: $\bar{\Gamma}$ and β as above; $r(t)$ set as an increasing function of $X(t) - X(t-1)$.
3. Macroprudential policy: β and r as per "constant policy"; $\bar{\Gamma}$ set, at long intervals, as a decreasing function of the mean $X(t) - X(t-1)$ over the period.

4 Model Calibration & Simulation Engine

In order to run our model we need to pin down its many parameters and the system's initial conditions. We do not attempt any rigorous calibration here because our objective is to exemplify how the model can be used, rather than provide definitive answers. However, we adopt the following strategy to minimize the impact from a necessarily approximate calibration of the initial conditions.

First, we calibrate a stripped down version of the model that does not include the CBDC - we do so looking at data from the German retail market and other stylized facts (e.g. long-term returns on equity). Having set these parameters and some realistic initial values for the state variables, we then perform a "model burn-in"

i.e. we simulate the model until a steady state is reached, while keeping switched off all the CBDC-related variables[11]. Finally, we introduce the CBDC. By so doing, we ensure that any subsequent change is attributable to the introdiction of the CBDC.

4.1 Consumers

The model calibration for the consumers is illustrated in Table 1. In what follows, we describe the parameterization of variables related to i) purchases & payments, and ii) wealth allocation.

Parameter	Description	Value	Data Source
N_c	Number of consumers	1500	Model assumption
$\overline{\eta}$	Average number of daily purchases	1500	Cabinakova et al. (2019)
$\bar{P}^{\alpha}_{i,j}, \hat{P}^{\alpha}_{i,j}, \tilde{P}^{\alpha}_{i,j}$	Mean, median and max purchase value	EUR 21.11, 14.21, 1000	Cabinakova et al. (2019)
p_{online}	The proportion of online purchases	20%	E-commerce data
$g(w, y)$	Adoption function, where w (y) is the proportion of consumers (merchants) who have adopted CBDC	$0.25(w + y)$	Model assumption
$[C, B, A] \sim f(.)$	Initial wealth distribution	See text.	Deutsche Bundesbank (2019)
κ	Exogenous salary	EUR 18	Comparable to median purchase value
r_A	Return on asset	$\frac{3\%}{365}$	Moody's
τ	CBDC top-up horizon	10	Bagnall et al. (2016)
τ_c	Cash withdrawal horizon	10	Bagnall et al. (2016)
ν	Adjustment on target wealth	1	Model assumption

Table 1: Model calibration for the consumers.

4.1.1 Purchases & payments

- *Distribution of the total number of purchases.* For a total number of consumer $N_C = 1500$ in the model, we assume that the total number of purchases is

[11] In practice, it is sufficient to assume that no merchant accepts CBDC, and that no CBDC is borrowed from the central bank.

distributed according to the Poisson distribution with mean $\overline{\eta} = 1500$. The parameterization is based on Cabinakova et al. (2019) who find that there are roughly 20 billion transactions in German retail market per year (54.8 million per day). Taking into account that the German's working population is about 53 million, this imply that each consumer makes approximately 1 payment per day.

- *Distribution of purchase value.* The value of $\bar{P}^{\alpha}_{i,j}$ and $\hat{P}^{\alpha}_{i,j}$ are calibrated from Cabinakova et al. (2019).
- *Probabilities involved in the "settlement rule" (to determine which payment instrument is used).* In the case where consumers needs to perform a random draw, we assume that each asset has the same probability to be chosen.
- *Probability of online purchases.* By comparing the value of retail[12] versus e-commerce spending[13], we assume that $p_{online} = 20\%$.

4.1.2 Wealth allocation

- *Initial wealth allocation* ($[C, B, A] \sim f(.)$). Deutsche Bundesbank (2019). We generate the consumers initial deposits from a (truncated) log-normal distribution with parameters mean EUR 7,100 and median EUR 1,800; and maximum EUR 10,000. These numbers are based on the distribution of German households' portfolio structure in 2017 published in Deutsche Bundesbank (2019). We also assume each consumer initially holds EUR 200 cash and EUR 0 CBDC.[14] The consumers assets are then obtained by scaling the wealth based on a target, $\frac{A}{B} = 1 - \frac{1}{\frac{r_A}{r_B}}$, defined previously in Equation (5).
- *Interest rate of asset* We assume r_A to be $\frac{3\%}{12}$, which is comparable to the Moody's seasoned AAA corporate bond yield in 2016.[15]
- *CBDC top-up and cash withdrawal horizon.* The value of τ and τ_c are set to be 10. Based on 2017 payment behavior survey (Deutsche Bundesbank, 2019), respondent reported to 42 times in a year withdraw ATM. This would approximately mean to withdraw money every 10 days.

4.2 Merchants

We summarize the parameterization for the merchants in Table 2.

- *The number of merchants in the model,* N_m, is calibrated as follow. According to Cabinakova et al. (2019), the German retail sector in 2016 encompasses

[12] https://tradingeconomics.com/germany/retail-sales-annual
[13] https://ecommercenews.eu/ecommerce-in-germany-was-worth-e83-3-billion-in-2020/
[14] We also assume that, when the consumers sign up for CBDC, they will initially add EUR 100 balance to their wallets.
[15] https://fred.stlouisfed.org/series/AAA

Parameter	Description	Value	Data Source
N_m	Number of merchants	15	Cabinakova et al. (2019)
$\mathscr{P}_B$	The proportion of merchants who accepts card payments	20%	Bagnall et al. (2016)
$\mathscr{P}_K$	The proportion of merchants who initially accepts CBDC	10%	Model assumption
$h(z_j(t))$	Adoption function, where $z_j(t)$ is the percentage of consumers that visited merchant j and asked to pay via CBDC up to tC	$h(z) = 0.25z$	Model assumption

Table 2: Model calibration for the merchants.

around 355 thousand stores. Taking into account that the German working-age population in 2016 is about 53 million, this then implies that there is a retailer for approximately each 100 people in the country. We therefore assume a total number of merchant $N_M = 15$ for a total consumer $N_C = 1500$ in the model. An agent in our model, therefore, represents approximately 35,000 consumers or 350 retailers in the German retail market.

- *The proportion of merchants who, at time $t = 0$, accept card and CBDC payments* ($\mathscr{P}_B$ and $\mathscr{P}_K$). We set the value of $\mathscr{P}_B$ to be 20%, which corresponds to $\mathscr{P}_B \times N_C = 20\% \times 15 = 3$ merchants in our parameterization. This is stemming from the study of Bagnall et al. (2016) who find that the use of cash is strongly correlated with merchant card acceptance, and that share of cash payment (by volume) reaches more than 80% in the Germany retail market. Meanwhile, we assume that $\mathscr{P}_K$ equals to 10%, which corresponds to $\mathscr{P}_K \times N_C = 10\% \times 15 = 1$ merchant.

- *CBDC adoption rule.* We assume the function $h(.)$ in Equation (9) to be $h(z) = 0.25z$.

4.3 Commercial Bank

Table 3 summarizes the model calibration for the commercial bank. The value of initial deposits, $B(0)$, is computed from the total deposits held by consumers. The initial CBDC borrowing, $K_{bank}(0)$, is set to be EUR 1M. Moreover, by considering a leverage ratio of 20%[16], and taking into account the value of $B(0)$ and $K_{bank}(0)$, we then obtain $X(0) =$ EUR 4.2M as the bank's initial investment in risky assets. Meanwhile, we set

[16]Note that, in this paper, we define the leverage ratio as the ratio between the total debts and the equity, while the Basel leverage ratio is defined as the equity divided by the total assets. The value of leverage ratio = 20 in this paper is therefore corresponds to the value of leverage ratio = 4.76% in the Basel definition.

the value of average return on risky assets, Λ, to be $\frac{1.2\%}{365}$[17], which is comparable to the real capital gain in housing (Jorda et al., 2019). Finally, the deposit rate is initially set to be $\frac{0.5\%}{365}$.[18] The value is comparable to the deposit rate in Germany in 2016.[19]

Parameter	Description	Value	Data Source
$B(0)$	Initial deposit	$\approx$ EUR 3M	Computed from the model
$K_{bank}(0)$	Initial CBDC borrowing	EUR 1M	Model assumption
$X(0)$	Initial investment in risky assets	$\approx$ EUR 4.2M	Computed from the model
$r_B(0)$	Deposit interest rate	$\frac{0.5\%}{365}$	World bank's development indicator
Λ	Return on risky assets	$\frac{1.2\%}{365}$	Jorda et al. (2019)
ξ	Friction in liquidating	0.95	Model assumption
ϵ	Mark-up on the cost of borrowing CBD	$\frac{0.3\%}{365}$	Computed from the model
μ	Adjustment on target balance sheet	1	Model assumption

Table 3: Model calibration for the commercial bank.

4.4 Central Bank

All the parameters for the central bank are primarily used to shed light onto the effects of central bank choices, so the central bank variables are left free, to be provided as an exogenous input. We assume a constant policy approach for the central bank where banking maximum leverage $\bar{\Gamma}$, CBDC lending rate r and limit β is constant in time, and set to levels which indicatively reflect current conditions. Table 4 summarizes the parameterization for the central bank. The maximum amount of CBDC that a consumer can hold in their wallet is $\beta =$ EUR 1,000.[20] Meanwhile, the maximum daily cash withdrawal is set to be fixed at $\beta_c =$ EUR 300.[21] With regard to the maximum leverage ratio that the commercial bank can take, we set $\bar{\Gamma}$ to be 32.3, which equals to 3% in the Basel terms. Finally, we set the interest rate of CBDC to be $r = \frac{0.2\%}{365}$.

[17] Note that we divide the value by 365, as the model runs daily.

[18] Here we avoid using negative interest rate, however it should be clear that the model can be also used to investigate the negative value.

[19] https://tradingeconomics.com/germany/deposit-interest-rate-percent-wb-data.html

[20] One can also run the simulation with different values of β. For example, Bindseil and Panetta (2020) illustrate a case where the CBDC maximum allowed balance equals to EUR 3,000.

[21] Note that the definition of β for CBDC and β_c for cash is slightly different. The former sets the maximum allowed CBDC balance, while the latter corresponds to maximum daily ATM withdrawal.

Parameter	Description	Value	Data Source
β	Maximum allowed CBDC balance	EUR 1,000	Model assumption
β_c	Maximum cash withdrawal	EUR 300	Model assumption
$\bar{\Gamma}$	Maximum leverage ratio	32.3	Model assumption
r	CBDC borrowing rate	$\frac{0.2\%}{365}$	Model assumption

Table 4: Model calibration for the central bank.

4.5 Simulation Software

We have developed a cloud-based simulator with FNA (www.fna.fi) to operationalize the dynamics described above. The simulator is a web application allowing users to define the model parameters, run the model and view the results of the simulation as it progresses, as well as download the simulation results for further analysis.

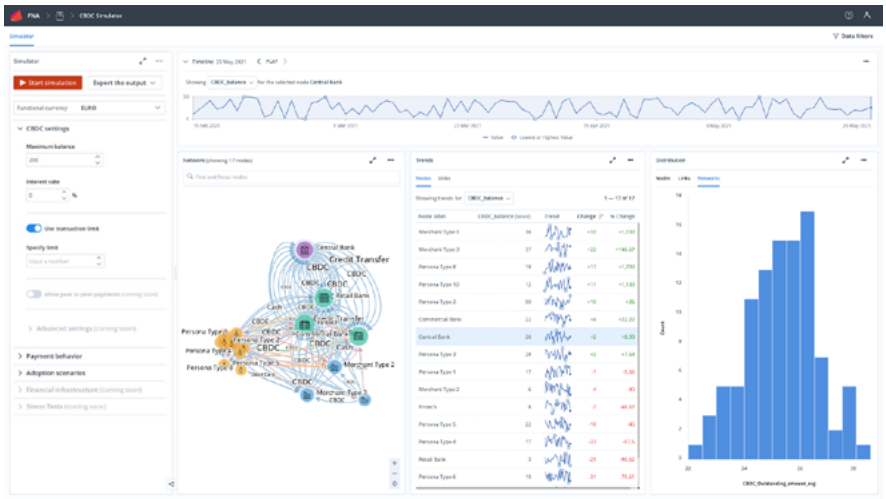

Fig. 3: A cloud-based agent-based simulator of CBDC created with FNA (www.fna.fi).

5 Results

This section illustrates model's behaviour, focusing on how a newly introduced CBDC affects the usage of alternative means of payment (and thus brings about "disintermediation"), modifies the composition of households' wealth, and alters the balance sheets of the banking sector. As mentioned, we parametrize all non-CBDC variables

as described in the above section and run the model for a "burn-in" phase till convergence. Only then, we introduce the CBDC, by letting 10% of merchants accept it.[22] In what follows, $t = 0$ therefore corresponds to the time when the CBDC is introduced. Simulations run from $t = 0$ to $t = 3500$ because by then a new equilibrium has been reached.

5.1 Adoption and Diffusion of Means of Payments

Figure 4 shows the evolution in time of the percentage of consumers (merchants) who have signed up for (accepted) CBDC. The left panel shows the whole simulation run, while the right one zooms into the transition phase, where adoption grows in a non-linear fashion.

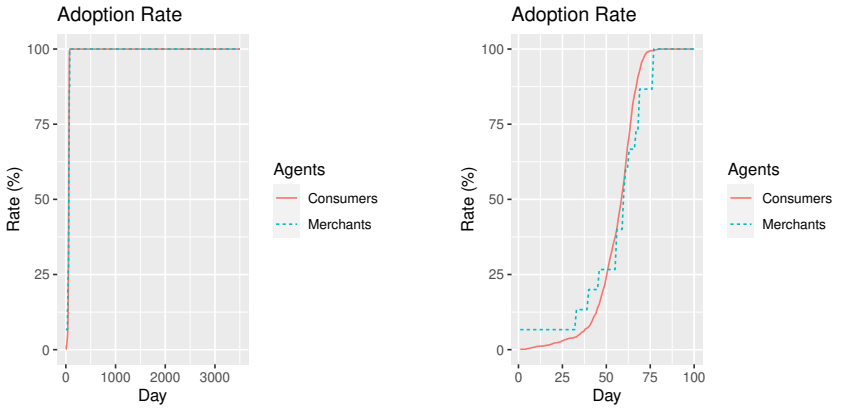

(a) Longer period (b) Shorter period

Fig. 4: The rate of adoption over time. At $t = 0$, 10% of merchants start to accept CBDC, but the number of consumers who have signed-up for CBDC is still 0.

Figure 5 illustrates the process of crowding out of traditional means of payments by the CBDC, showing the total value and the total number of transactions settled in cash, cards and CBDC. Panels (a) and (b) show that, while the relative shares of cash, card and CBDC become stable around $t = 1500$, the CBDC had already been fully adopted some time before (at $t = 75$, as shown in Figure 4). Following our particular calibration, panel (c) shows that, by the time the shares of payments stabilize, both cash and cards have declined by approximately 30%, while CBDC establishes itself as the main means of payment. The decline of card payments by 30% means that card companies would loss 30% $\times$ EUR 5 billion = EUR 1.5 billion transaction revenues in the German domestic payments market.[23]

[22] Recall that, as we consider a total number of 15 merchants, 10% therefore corresponds to 1 merchant in the model calibration.

[23] According to Germann et al. (2019), the card transactions expected revenues in the German domestic payments market for 2022 is EUR 1.5 billion.

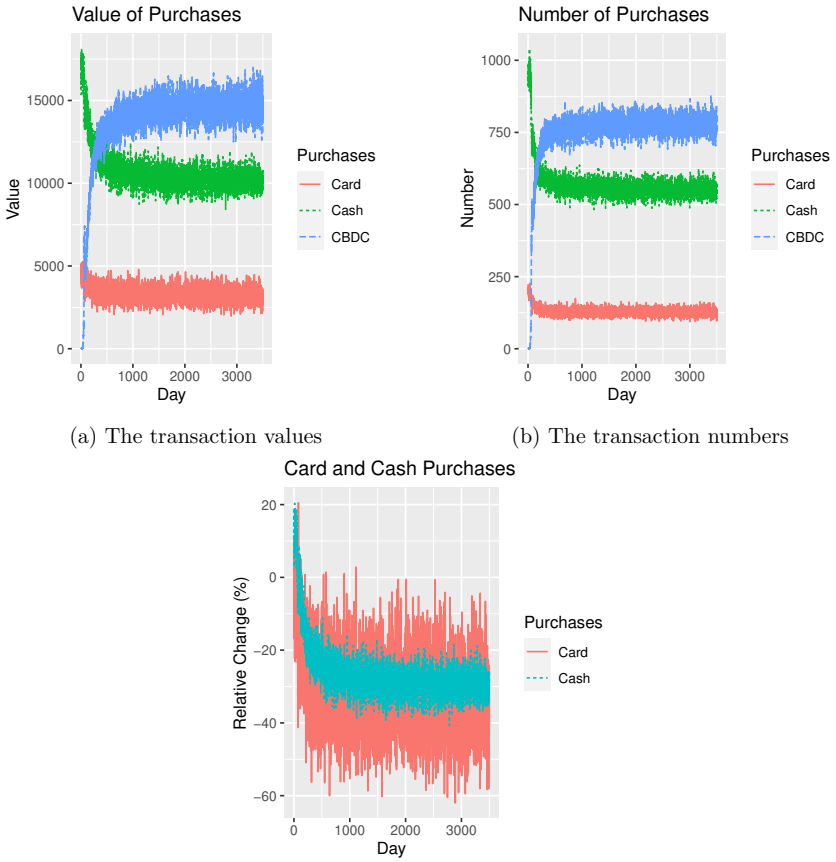

(a) The transaction values

(b) The transaction numbers

(c) The relative changes

Fig. 5: The diffusion of means of payments over time. Panel (a) shows the comparison between the total transaction values that are settled via cash, card and CBDC, while panel (b) shows the corresponding transaction numbers. Panel (c) describe the relative change in the total cash and card transaction values.

5.2 Consumer Wealth Allocation

In the previous section we saw CBDC crowding out deposits and cash *as form of payments*. Figure 6 below shows that a similar dynamics takes place for consumer's *holdings* of cash, deposits and CBDC. Panel (a) in particular illustrates how part of the rise in CBDC happens at the expenses of cash. However, as shown by panel (d), bank deposits and financial assets keep increasing. This happens because we assume that consumers receive an exogenous income that is (slightly) higher than daily purchases[24], so that the overall consumer wealth grows. However (this is not shown in the charts but it is suggested later on by Figure 6), the growth in deposits is less than what it would be in the absence of CBDC: cash holdings are impacted to the point of decreasing, deposits are impacted in the sense of increasing at a lower rate.

Panel (b) of Figure 6 shows the same variables as panel (a), but zooms into the initial transition period and reveals a somehow peculiar behavior in CBDC holdings: these first increase rapidly, then fall, then increase again. This happens because many consumers, on opening a CBDC wallet, top it up by more than they actually need, generating a sudden swell in CBDC accounts[25]. Later, though, they go over the initial CBDC 'overhang', so their aggregate holdings, although subject to small random fluctuations, follow an essentially monotonic path. All this is a not-so-subtle consequence of our exogenous calibration but the point here is that, with appropriate calibration and extensions, this model can generate patterns resembling e.g. the 'initial / irrational exuberance' that often accompanies financial innovation.

Another type of analysis that can be carried out within this framework is cross-sectional analysis, to look at how wealth (and its composition) varies in the population of agents. For example, Figure 7 shows the histograms of CBDC balance at different periods: $t = 44, 125, 1000, 3000$. Note that, as shown previously in Figure 4, the adoption rate grows substantially after $t = 44$ and reaches 100% rate at $t = 75$. We see from the figure that, on $t = 44$, most of the consumers have zero in their balance as they still have not signed up for CBDC. Meanwhile, on $t = 125$, we find that the number of consumers with zero CBDC balance has declined, and that with EUR 22-40 has substantially increased. We then observe that the shape of the distribution becomes more similar at $t = 1000$ and $t = 3000$. In particular, we see that the majority of consumers hold $\approx$ EUR 50-100 at both periods. Note that, in all figures, there are still consumers with zero CBDC balance, even though the adoption rate is already 100%. This is because some consumers may not enough deposits to top up their CBDC wallets.[26]

In a similar vein, Table 5 reports the mean and median of assets holdings, across agents and at four points in time (again $t = 44, 125, 1000, 3000$). For example, we see from the table that a consumer holds, initially at t, EUR 88 cash in average. We then see how the cash holding decreases over time and fall to EUR 70 at $t = 3000$. Moreover, we also see that the CBDC holding grows from EUR 9 at $t = 44$ to EUR

[24] As per Section section 4, the daily exogenous income is 18 EUR and the median (mean) of the daily purchases is 14.21 (21.11).

[25] The initial top-up is exogenously set at 100 euros. Subsequent top-ups are determined by individual consumers' average CBDC spending.

[26] We assume that the consumers can still sign up and leave the balance to zero.

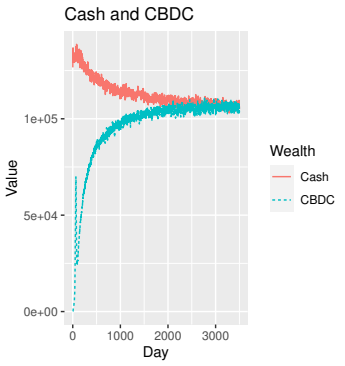

(a) Total cash and CBDC

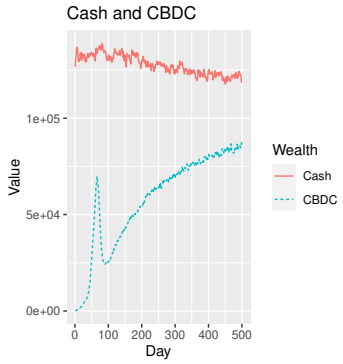

(b) (a) Total cash and CBDC but over a shorter period of time

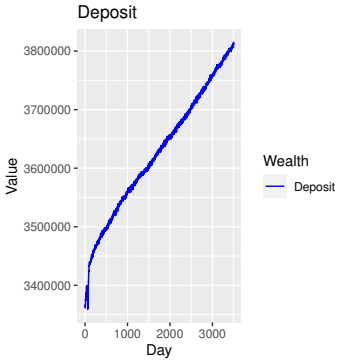

(c) Total deposits

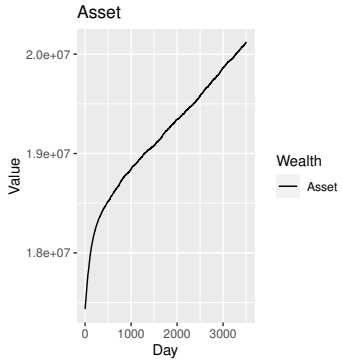

(d) Total other assets

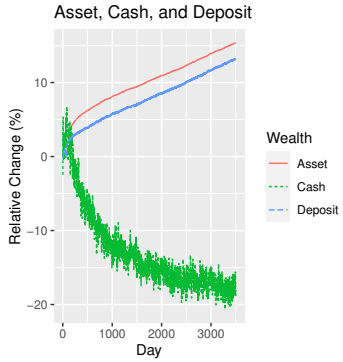

(e) Relative change of assets, cash, deposits

Fig. 6: The composition of consumer wealth over time (sum over all consumers) allocated between cash, CBDC, deposits and other assets.

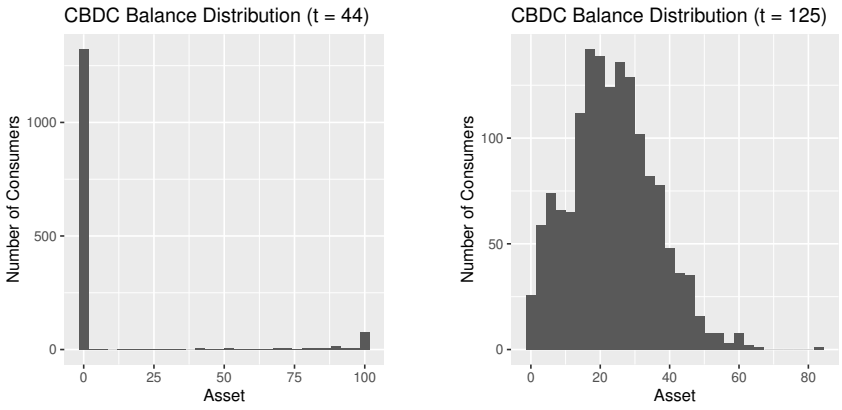

(a) $t = 44$ (b) $t = 125$

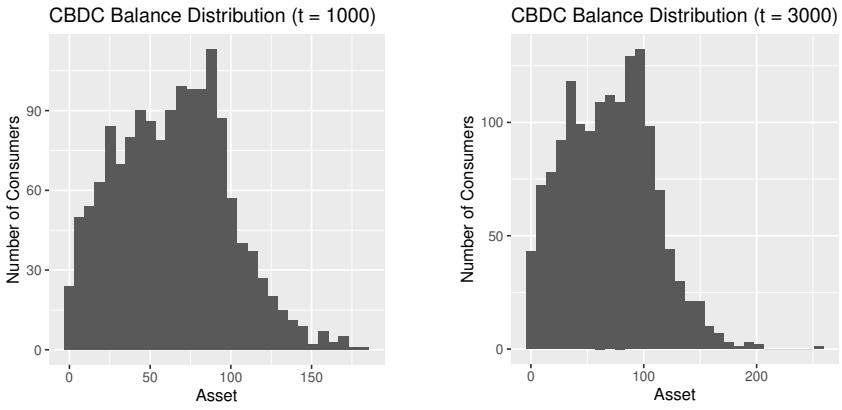

(c) $t = 1000$ (d) $t = 3000$

Fig. 7: Distribution of CBDC balances across agents (note different scale of y-axes).

71 at $t = 3000$.

Wealth	$t = 44$		$t = 125$		$t = 1000$		$t = 3000$	
	Mean	*Med*	*Mean*	*Med*	*Mean*	*Med*	*Mean*	*Med*
Cash	88	88	89	90	76	77	70	73
CBDC	9	0	23	22	65	65	71	71
Deposit	2265	1281	2292	1316	2373	1315	2543	1301
Asset	11790	6906	12014	7123	12563	7261	13413	7134

Table 5: Wealth allocation over time (mean and median across agents.

5.3 Banking Disintermediation and Balance Sheet

In the following, we study the impact of the introduction of a CBDC on the banking disintermediation. The disintermediation process, in general, involves removing the middleman (commercial bank) that sits between two parties (consumers and merchants) in a transaction. Disintermediation can be investigated, first, by looking at the total value of CBDC in circulation compared to that of bank deposits. This is illustrated by Figure 8 - panel (a) that shows the ratio of the former to the latter is 2.8%. To put this number into perspective, the impact of 2.8% to the total bank deposits in Germany would corresponds to EUR 75.04 billion.[27]

Another aspect of disintermediation will be a negative impact on bank deposits. In this regard, Figure 8 - panel (b) shows that deposits keep growing after introduction of the CBDC[28], but at a slower pace than before.[29] In more detail, the growth rate in deposits falls by approximately three quarters (from $0.75 - 1\%$ to $0 - 0.25\%$). To put this into perspective, for the German banking market the fall would correspond to approx EUR 20 billion in 20 days. In the short run, we also see a fall in deposits by approximately 0.75% in the period through $t = 75$ (by which time the CBDC is fully adopted by all consumers and merchants), related to the synchronous initial CBDC top-ups that discussed when commenting Figure 6.

[27] As described in https://fred.stlouisfed.org/series/DDOI02DEA156NWDB, the total bank deposits in Germany for 2017 is EUR 2,68 trillion. The impact of 2.8% will therefore correspond to $2.8\% \times$ EUR $2,68$ trillion. $=$ EUR 75.04 billion.

[28] This is mainly due to our assumption that consumers' (exogenous) income is above their spending

[29] The growth rate depicted in the figure is $\frac{B(t)-B(t-20)}{B(t-20)}$, i.e. is the average growth over 20 days. Daily rates are much more volatile.

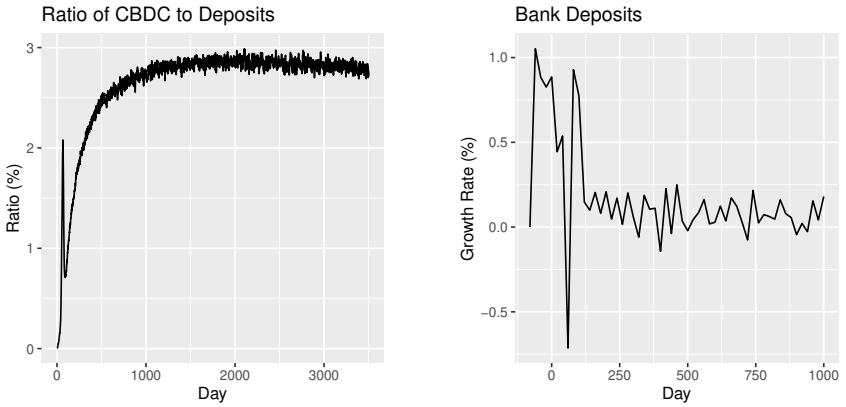

(a) Ratio of CBDC to deposits (b) Bank deposit growth rate

Fig. 8: Banking sector disintermediation. Panel (a) describes the ratio of the total value of CBDC to deposits. Panel (b) shows the growth rate of bank deposits over the previous 20 days. The CBDC is introduced to the system at $t = 0$ and is fully adopted by all consumers and merchants at $t = 75$.

6 Discussion

We proposed an agent-based simulation model to study the economic impact of introducing a retail central bank digital currency (CBDC): a general purpose and non-interest-bearing central bank liability, which competes with cash and card as a means of payments. Our model features 4 classes of agents, namely consumers, merchants, a commercial bank and a central bank, who interact in time according to defined rules. We used data on the German economy to calibrate the model.

Firstly, we looked at the adoption rate of CBDC and the change in the diffusion of means of payments. We showed a non-linearity in the adoption function such that the rate would grow substantially following the adoption by a small proportion of consumers and merchants in the system. As CBDC is being adopted, we found that both cash and card payments would decline by approximately 30%. The latter indicates that card companies would loss EUR 1.5 billion transaction revenues in the German domestic payments market. Secondly, we studied the composition of consumer wealth that is allocated between cash, CBDC, deposits and other assets. Over time, we found that consumers would allocate their wealth for a higher value of CBDC and a lower value of cash. We also saw a slower growth rate of deposits and other asset following the CBDC adoption. Thirdly, we discussed the banking disintermediation and balance sheet. We found that the total value of CBDC in circulation is 2.8% of the total deposits, which corresponds to EUR 75.04 billion of the total deposits in the German banking system. We also found that the growth rate of the bank deposits would decline by 0.75% in 20 days, which corresponds to the fall of approximately EUR 20 billion for the German banking market.

The model presented in this paper is still a prototype, ready for refinements and

extensions. First the parameters shown here, meant to approximately represent the German retail market, could be pinned down more precisely, and/or be tailored to gauge the effects of a CBDC on other economies. Second, a variety of policy experiments could be carried out, to assess the impact of different CBDC 'configurations' or design. Third, in the current version of the model, the CBDC does not bear interests, so consumers' demand is determined only by retail payment needs. Small changes to the model could allow for an interest-bearing CBDC competing with deposits and other assets as a store of value; and similarly, for other payment instruments such as electronic money and stable coins. Finally, by adding a richer variety of consumers and more articulated bank behaviour, the model could be extended to explore issues related to financial inclusion and financial stability.

References

Bagnall, J., D. Bounie, K. P. Huynh, A. Kosse, T. Schmidt, and S. Schuh (2016). Consumer cash usage: A cross-country comparison with payment diary survey data. *International Journal of Central Banking 12*(4).

Barontini, C. and H. Holden (2019). Proceeding with caution - a survey on central bank digital currency. *BIS Papers No 101*.

Beyeler, W. E., R. J. Glass, M. L. Bech, and K. Soramäki (2007). Congestion and cascades in payment systems. *Physica A: Statistical Mechanics and its Applications 384*(2), 693–718.

Bindseil, U. (2021). Tiered cbdc and the financial system. *European Central Bank Working Paper No 2351* (January).

Bindseil, U. and F. Panetta (2020). Central bank digital currency remuneration in a world with low or negative nominal interest rates. *VoxEU*.

Boar, C. and A. Wehrli (2021). Ready, steady, go? – results of the third bis survey on central bank digital currency. *BIS Papers No 114* (January).

Bolt, W. and K. Soramäki (2008). Competition, bargaining power and pricing in two-sided markets. *DNB Working Paper No. 181* (September).

Byck, S. and R. Heijmans (2021). How much liquidity would a liquidity-saving mechanism save if a liquidity-saving mechanism could save liquidity? A simulation approach for canada's large-value payment system. *Journal of Financial Market Infrastructures 9*(3).

Cabinakova, J., F. Knümann, and F. Horst (2019). The costs of cash payments in the retail sector. *Deutsche Bundesbank Report*.

Central Bank of the Bahamas (2019). Project sand dollar: A bahamas payments system modernisation initiative.

Deutsche Bundesbank (2019). Household wealth and finances in germany: results of the 2017 survey. *Deutsche Bundesbank Monthly Report* (April), 13–42.

Farmer, J. D., A. M. Kleinnijenhuis, P. Nahai-Williamson, and T. Wetzer (2020). Foundations of system-wide financial stress testing with heterogeneous institutions. *Bank of England Staff Working Paper No. 861* (May).

Federal Reserve Bank of Boston (2020). The federal reserve bank of boston announces collaboration with mit to research digital currency.

Fernández-Villaverde, J., S. D., L. Schilling, and H. Uhlig (2020). Central bank digital currency: central banking for all? *National Bureau of Economic Research Papers* (February).

Galbiati, M. and K. Soramäki (2011). An agent-based model of payment systems. *Journal of Economic Dynamics and Control 35*(6), 859–875.

Geanakoplos, J., R. Axtell, D. J. Farmer, P. Howitt, B. Conlee, J. Goldstein, M. Hendrey, N. M. Palmer, C.-y. Yang, N. Haven, and B. J. Geanakoplos (2012). Getting at systemic risk via an agent-based model of the housing market. *102*(1358), 53–58.

Germann, F., R. Höll, and M. Niederkorn (2019). A perspective on german payments. *McKinsey & Company Global Banking Practice* (September).

Gross, J. and J. Schiller (2021). A model for central bank digital currencies: implications for bank funding and monetary policy. *SSRN Electronic Journal*.

Hałaj, G. (2018). System-wide implications of funding risk. *Physica A: Statistical Mechanics and its Applications 503*, 1151–1181.

Hellqvist, M. and J. Koskinen (2005). Stress testing securities clearing and settlement systems using simulations. In *Liquidity, Risks and Speed in Payment and Settlement Systems – a Simulation Approach*, pp. 323–350. Bank of Finland.

Jorda, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics 134*(3), 1225–1298.

Keister, T. and C. Monnet (2020). Central bank digital currency: stability and information. *Mimeo*.

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 20*(104).

Koponen, R. and K. Soramäki (1998). Intraday liquidity needs in a modern interbank payment system: A simulation approach. *Scientific Monographs from Bank of Finland*.

Kumhof, M. and C. Noone (2018). Central bank digital currencies - design principles and balance sheet implications. *Bank of England Staff Working Paper No. 725* (May).

Leinonen, H. and K. Soramäki (2004). Simulation: A powerful research tool in payment and settlement systems. *Journal of Financial Transformation*, 79–84.

Markose, S., E. Tsang, and S. M. Jaramillo (2005). The red queen principle and the emergence of efficient financial markets: an agent based approach. In *Lecture Notes in Economics and Mathematical Systems*, Volume 550, pp. 287–303. Springer, Berlin, Heidelberg.

McLafferty, J. and D. Edward (2013). Liquidity saving in chaps: A simulation study. In *Simulation in Computational Finance and Economics: Tools and Emerging Applications*, pp. 23.

Panetta, F. (2018). 21st century cash: central banking, technological innovation and digital currency. In E. Gnan and D. Masciandaro (Eds.), *SUERF Conference Proceedings 2018/2*, pp. 23–32.

Reuters (2021). RPT-update 1-ecb's digital euro could run on instant payment platform, visco says.

Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. E. Beyeler (2007, jun). The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications 379*(1), 317–333.

Turrell, A. (2016). Agent-based models: understanding the economy from the bottom up. *Bank of England Quarterly Bulletin Q4*, 173–188.

Wilensky, U. and W. Rand (2015). An introduction to agent-based modeling. *MIT Press*.

4. Resources

CBDC Think Tank	https://cbdctt.com
CBDC Insider	https://cbdcinsider.com
OpenCBDC Sandbox	https://opencbdcs.com