MISSING KEY:

The Challenge of Cybersecurity and Central Bank Digital Currency

Atlantic Council
GEOECONOMICS CENTER

ACKNOWLEDGMENTS

This report is written and published in accordance with the Atlantic Council Policy on Intellectual Independence. The authors are solely responsible for its analysis and recommendations. The Atlantic Council and its donors do not determine, nor do they necessarily endorse or advocate for, any of this report's conclusions.

ISBN-13: 978-1-61977-236-6

Cover: iStock

June 2022

© 2022 The Atlantic Council of the United States. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Atlantic Council, except in the case of brief quotations in news articles, critical articles, or reviews. Please direct inquiries to:

Atlantic Council, 1030 15th Street NW, 12th Floor, Washington, DC 20005

MISSING KEY:

The Challenge of Cybersecurity and Central Bank Digital Currency

Lead Authors:
Giulia Fanti and Kari Kostiainen

Contributing Authors:
William Howlett, Josh Lipsky, Ole Moehr, John Paul Schnapper-Casteras, and Josephine Wolff

Atlantic Council
GEOECONOMICS CENTER

The **GeoEconomics Center** works at the nexus of economics, finance, and foreign policy with the goal of helping shape a better global economic future.
The Center is organized around three pillars - Future of Capitalism, Future of Money, and the Economic Statecraft Initiative.

Table of Contents

Foreword

The challenge of securing the dollar dates back to the earliest days of the United States. Benjamin Franklin famously printed currency with the phrase "to counterfeit is death"—and colonial England used fake currency to try to devalue the Continental Dollar during the American Revolution.

In the modern era, security issues have multiplied with the rise of the Internet and the threat of cyberattacks. The United States Federal Reserve (Fed) considers cybersecurity a top priority and sees securing both the dollar and the international financial system as a core national security challenge. We are entering a new era of security and currency, one that requires responsible innovations in digital currency. This report examines the novel cybersecurity implications that could emerge if the United States issues a government-backed digital currency—known as a central bank digital currency (CBDC) or "digital dollar."

This topic is fast-moving, consequential, and still somewhat nascent.

CBDCs have quickly landed on the international policy landscape. As of June 2022, according to Atlantic Council research, 105 countries representing 95 percent of the global GDP are researching and exploring the possible issuance of CBDCs. In the United States, spurred on by various domestic and international factors, the Fed has begun studying the issue and published a white paper in January 2022 that examines the potential benefits and risks of issuing a CBDC. In February 2022, the Federal Reserve Bank of Boston, in collaboration with the Massachusetts Institute of Technology, released test code and key findings on what a possible US CBDC might look like. But the government has so far demurred on whether it will actually issue a digital dollar, calling upon Congress to authorize such a major decision. Further complicating matters is the rapid ascendance of privately issued crypto dollars, sometimes referred to as stablecoins, which now surpass $130 billion in total market capitalization. As Fed Vice Chair Lael Brainard testified to the US House of Representatives' Committee on Financial Services in May 2022, the recent collapse of the stablecoin TerraUSD raises new questions about the ways in which a CBDC could stabilize the digital asset ecosystem.

The security of CBDCs has real-world import and is one of the major challenges to overcome if a CBDC is to be issued in the United States. Not just because of the classical counterfeiting scenarios or the possibility of a hacker looting the digital equivalent of Fort Knox, but also because a government-administered digital currency system could—depending on how it is designed—collect, centralize, and store massive amounts of sensitive data about individual Americans and granular details of millions of everyday transactions. For example, a CBDC could contain large volumes of personally identifiable information ranging from what prescription drugs you buy or where you travel each day. This could become a rich trove of data that could be stolen by advanced hackers or nation-states (similar to reams of personal data collected from federal employees that was stolen in 2016). Separately, other security issues could arise, for example, misuse or exfiltration of data by inside employees, smaller-scale identity theft, or "gray" charges via opaque fees. However, as our analysis shows, many of these risks already exist in the current system and could be mitigated through an effectively designed CBDC.

The debate around CBDCs in the United States is also, relatively speaking, in its infancy, with the Fed and Treasury Department often taking the lead thus far, and several CBDC-related bills percolating through Congress. Part and parcel of the conversation about how and whether to develop a CBDC in the United States is what it will look like and how secure it could be. These intertwined questions of policy, design, and security should be an increasing focus of the conversation, both among federal agencies and between the executive branch and Congress. The United States can, and should, play a leading role in international standard setting. US President Joseph R. Biden, Jr.'s recent executive order highlighted the importance of digital assets protecting democratic values.

This report introduces key concepts, potential design trade-offs, and some policy principles that we hope can help federal stakeholders make foundational decisions around the future of CBDCs in the years ahead. While it is too early for a CBDC to be designed with ideal cybersecurity, efforts to dismiss a CBDC as uniquely and categorically vulnerable to cyberattacks have overstated the risk. This report puts forward a road map for policy makers to build secure CBDCs.

Executive Summary

This report examines the novel cybersecurity implications that could emerge if the United States or another country issues a Central Bank Digital Currency (CBDC). Central banks consider cybersecurity a major challenge to address before issuing a CBDC. The United States Federal Reserve (Fed) sees securing both the dollar and the international financial system as a core national security imperative. According to Atlantic Council research, currently 105 countries have been researching and exploring the possible issuance of CBDCs, with fifteen in pilot stage and ten fully launched.[1] Of the Group of Twenty (G20) economies, nineteen are exploring a CBDC with the majority already in pilot or development. This raises immediate questions about cybersecurity and privacy. A government-issued digital currency system could, but does not necessarily need to, collect, centralize, and store massive amounts of individuals' sensitive data, creating significant privacy concerns. It could also become a prime target for those seeking to destabilize a country's financial system.

This report analyzes the intertwined questions of policy, design, and security to focus policy makers on how to build secure CBDCs that protect users' data and maintain financial stability. Our analysis shows that privacy-preserving CBDC designs are not only possible, but also come with inherent security advantages, compared to current payment systems, that may reduce the risk of cyberattacks. Divided into three chapters, the report:

1. provides a brief background on the Fed's process as a baseline for central banks' current cybersecurity measures;

2. explores the novel cybersecurity implications of different potential CBDC designs in depth; and

3. outlines legislative and regulatory principles for policy makers in the United States and beyond to set the conditions for secure CBDCs.

Payment systems' status quo: how the Federal Reserve currently secures payments

Current wholesale and retail payment systems face a complex cybersecurity landscape and represent a major point of attack for both organized crime and state-sponsored actors. Cybersecurity risks posed by CBDCs must be assessed relative to this landscape.

A targeted attack on wholesale payment infrastructures, such as the Fed's domestic funds transfer system, Fedwire, could cause major global financial shocks, including severe liquidity shortfalls, commercial bank defaults, and system-wide outages that would affect most daily transactions and financial stability. There would also be secondary effects, including severe market volatility. To minimize the risk of cyberattacks and reduce the impact of successful hacks, the Fed's current measures include regular contingency testing for high-volume and high-value Fedwire participants; redundancy requirements, such as backup data centers and out-of-region staff; and transaction value limits. Other risks for the wholesale payments infrastructure include attacks on the Society for Worldwide Interbank Telecommunications (SWIFT) messaging system. After recent attacks revealed significant vulnerabilities, SWIFT and its member banks have taken several steps to shore up their defenses, focusing on stronger security standards and quicker response.

Key cybersecurity risks for retail payment systems include credit and debit fraud, which collectively caused nearly $25 billion in damages in 2018 worldwide; fraud by system insiders affecting platforms like the Automated Clearing House (ACH) in the United States; and user error, such as falling prey to phishing scams. Risk management strategies for retail payments often rely on voluntary industry standards, such as the Payment Card Industry Data Security Standard (PCI DSS). To counter phishing and other types of user error, ACH and other platforms require unique user credentials and offer merchants

1 "Central Bank Digital Currency Tracker," Atlantic Council, last updated June 2022, https://www.atlanticcouncil.org/cbdctracker/.

additional steps like micro validation, tokenization and encryption, and secure vault payments.

In sum, the various technical systems administered by the Fed, industry associations, and private banks already face considerable cybersecurity challenges.

Cybersecurity of CBDCs—threats and design options

While a CBDC would be subject to many of the same cybersecurity risks as the existing financial systems, deployment of a CBDC would also create new risks. Depending on the choice of CBDC design, potential new cybersecurity risks include (but are not limited to):

- Increased centralization of payment processing and sensitive user data. It is possible a central bank would store user activity and transactions.
- Reduced regulatory oversight of financial systems
- Increased difficulty reversing fraudulent or erroneous transactions
- Challenges in payment credential management and key custody
- Susceptibility to erroneous or malicious transactions enabled by complex, automated financial applications
- Increased reliance on third parties (e.g., non-banks)

The exact set of new cybersecurity risks depends largely on the digital currency variant that a country chooses for its CBDC system. Each digital currency variant also provides different properties in terms of system scalability, system robustness, user privacy, and networking requirements. Since each currency design variant presents different trade-offs in terms of performance, security, and privacy, the choice of which digital currency design variant to deploy as a CBDC is a policy choice for finance ministries, central banks, and legislatures. It should be driven by a thorough analysis of the relative technical trade-offs.

This report reviews various possible digital currency design variants and compares the benefits and risks of each. **Our analysis also challenges the prevailing thinking in several ways and outlines the following findings.**

Finding 1: The design space for CBDCs is larger than the often-presented trifecta of centralized databases, distributed ledgers, and token models.

- For example, this means that both ledger and token-based payments can embody robust privacy protection through certain cryptographic measures.

Finding 2: CBDCs can enable both strong user privacy and (some level of) regulatory oversight at the same time.

- It is possible to design systems where users enjoy reasonable levels of payment privacy and regulatory authorities can at the same time advance other important policy goals.

Finding 3: A privacy-preserving currency design can inherently provide security advantages.

- In a privacy-preserving CBDC deployment that initially declines to collect or subsequently restricts sensitive user data even from trusted system insiders, breaches will have significantly less severe security consequences.

Finding 4: It is critical to use best practices from system design, such as proven consensus protocols and cryptographic primitives.

- Distributed security protocols, such as those used to secure distributed ledgers, can introduce subtle new design challenges and security trade-offs. This report encourages the use of well-tested protocols with provable security guarantees as key components of CBDC deployments.

Principles for future legislation and regulation

With most governments, including in the United States, still weighing whether to develop a CBDC, this report identifies key principles to help guide policy makers and regulators on how to deploy a CBDC with robust cybersecurity protections in mind.

Principle 1: Where possible, use existing risk management frameworks and regulations.

- Depending on the CBDC design, policy makers and regulators should assess which areas of a new CBDC ecosystem will be covered by current laws and regulations and where novel statutes—or new technical frameworks—might be necessary to provide adequate protection.
- When crafting new regulations for a CBDC, policy makers and regulators should set the conditions for a

safe digital currency ecosystem that enables financial intermediaries to innovate and compete.

Principle 2: **Privacy can strengthen security.**

- Privacy-preserving CBDC designs can have security benefits because they reduce the risk and potential harmful consequences of cyberattacks associated with data exfiltration and the centralization of detailed personally identifiable information.
- CBDCs can offer cash-like privacy, while potentially providing reasonable oversight options to regulatory authorities. A CBDC's level of privacy is a legislative and political choice that will filter through to the digital currency's design and determine its cybersecurity profile.

Principle 3: **Test, test, and test some more.**

- Governments should ensure that they have full access to, and can directly oversee, security testing and audits for all CBDC implementation instances. To enable extensive testing and security audits, the US Congress should consider the appropriations accordingly as part of next year's budget process and allocate a pilot project.
- Open-source CBDC code bases may be valuable for various reasons, including because they allow for more participation in the security testing process, especially when combined with longer-term bug bounty programs. Nonetheless, they still require due attention, funding, and staffing to maintain and monitor the code base over the long run.

Principle 4: **Ensure accountability.**

- The overall framework governing CBDCs needs to establish clear rules and policies surrounding accountability for errors, breaches, and resulting consequences (both technical and financial).
- For CBDCs that rely on distributed ledger technology (DLT), it is paramount to clearly establish accountability requirements among validators on the blockchain.

Federal Reserve chair Jerome Powell testifies before a U.S. House Financial Services Committee hearing on Capitol Hill in Washington, U.S., March 2, 2022. Source: REUTERS/Tom Brenner

Principle 5: **Promote interoperability.**

- To increase the resiliency of countries' existing financial systems, policy makers should develop rules to ensure that a CBDC is interoperable with the country's relevant financial infrastructure.
- To strengthen the security of CBDC systems, US leadership is critical to promote global interoperability between CBDCs through international coordination on regulation and standard setting through fora like the Group of Seven (G7), the G20, the Financial Stability Board (FSB), and the Financial Action Task Force (FATF).

Principle 6: **When new legislation is appropriate, make it technology neutral.**

- The US Congress can help study and oversee the application of federal cybersecurity laws to a potential CBDC with the goal of developing laws that apply evenhandedly to different technologies over time.
- Congress may consider using incentives and accountability for CBDC development or set security requirements by empowering a federal agency to develop a cybersecurity framework for a CBDC as part of a pilot project.

Background:
How the United States Currently Secures Its Payment Systems

Cybersecurity is an area of concern not only for CBDCs but also the current financial and payment systems. Any study of CBDCs' cybersecurity must assess them relative to this current infrastructure and recognize how they will interact to alter and potentially remedy existing vulnerabilities. Additionally, it must draw lessons from how central banks currently handle payments' cybersecurity.

The Fed has recognized the immense risks posed by cyberattacks to the current financial system. Asked in April 2021 about the chances for a systemic breakdown like the 2008 financial crisis, Federal Reserve Chairman Jerome Powell said that "the risk that we keep our eyes on the most now is cyber risk." He specifically singled out a scenario in which a "large payment utility…breaks down and the payment system can't work" or "a large financial institution would lose the ability to track the payments that it's making."[2] At a conference in October, Loretta J. Mester, president of the Federal Reserve Bank of Cleveland, argued "there is no financial stability without cybersecurity."[3] As the issuer of the world's reserve currency, the Fed's cybersecurity models hold outsized importance for the global economy. The Fed's standards have also become models for cybersecurity across central banks.

PAYMENTS OVERVIEW

The current payment system comprises three categories: retail, wholesale, and cross-border.[4] Retail payments are what the vast majority of Americans interact with: purchasing groceries with a credit card, buying Cracker Jack at a baseball game with a five-dollar bill, or shopping online with payment service providers. The wholesale system operates in the background, serving as the plumbing of the financial system by enabling the transfer and settlement of funds between financial institutions. Cross-border payments are between different countries and require international coordination to bridge national systems.

All three of these systems could be impacted or overhauled by CBDC. CBDC is the digital form of a country's fiat currency that is also a claim on the central bank. Instead of printing money, the central bank issues electronic coins or accounts backed by the full faith and credit of the government. This differs from current "e-money" because it is a direct liability of the central bank, like paper cash. A CBDC could take multiple forms: a retail CBDC would be issued to the public to enable fast and secure payment, while a wholesale CBDC would only be accessible by banks and would facilitate large-scale transfers. According to the Atlantic Council's CBDC tracker, forty-five of the 105 countries pursuing a CBDC are focused on its retail use, while eight are exclusively developing it for a wholesale purpose, and twenty-three are doing both (with the remaining twenty-nine undecided).[5] On the cross-border payments front, multiple partnerships between countries, such as Project Dunbar among South Africa, Singapore, Malaysia, and Australia, are piloting cross-border payments using CBDCs.

The key components of the United States' current payment systems are described below.[6]

- Fedwire is the Fed's domestic and international funds transfer system that handles both messaging and settlement.
- Clearing House Inter-Payments System (CHIPS), privately operated and run by its member banks, handles dollar-denominated domestic and international funds transfers.[7]

2 Scott Pelley and Jerome Powell, "Jerome Powell: Full 2021 60 Minutes Interview Transcript," CBS News, April 11, 2021, https://www.cbsnews.com/news/jerome-powell-full-2021-60-minutes-interview-transcript/.

3 Federal Reserve Bank of Cleveland President Loretta J. Mester, "Cybersecurity and the Federal Reserve," speech to the Fourth Annual Managing Cyber Risk from the C-Suite Conference, October 5, 2021, https://www.clevelandfed.org/newsroom-and-events/speeches/sp-20211005-cybersecurity-and-the-federal-reserve.aspx.

4 Eswar S. Prasad, *The Future of Money: How the Digital Revolution Is Transforming Currencies and Finance* (Cambridge, Massachusetts: The Belknap Press of Harvard University Press, 2021), 45–48.

5 "Central Bank Digital Currency Tracker."

6 See the appendix of this report for a detailed analysis of US payment system providers' current cybersecurity measures.

7 "CHIPS," Clearing House, accessed January 14, 2022, https://www.theclearinghouse.org/payment-systems/chips.

The Marriner S. Eccles Federal Reserve Board Building in Washington, DC.

- The Society for Worldwide Interbank Telecommunications (SWIFT), operated as a consortium by member financial institutions, is a global messaging system that interfaces with Fedwire and CHIPS for the actual settlement of payments.[8]

- FedNow will complement the Fed's Fedwire with instant, around-the-clock settlement and service. A full rollout is planned over the next two years.

- The Automated Clearing House (ACH) is a network operated by the National Automated Clearing House Association (Nacha) that aggregates US transactions for processing and enables bank-to-bank money transfers.

While CBDCs will likely play a role in all three levels of the payment system, this background chapter as well as the report's appendix predominantly examine risks to payment systems in which central banks are involved. That currently means the wholesale system. The Fed's approach to securing wholesale payments sheds light on its current cybersecurity practices and how it might handle a CBDC. We also briefly examine the retail payment system to understand cyber risks that a retail CBDC could impact.

A targeted attack on wholesale payment infrastructures, such as Fedwire, could cause major global financial shocks, including severe liquidity shortfalls, commercial bank defaults, and system-wide outages that would affect most daily transactions and financial stability. There would also be secondary effects, including severe market volatility. To prevent cyberattacks and reduce the impact of successful hacks, the Fed's current measures include regular contingency testing for high-volume and high-value Fedwire participants; redundancy requirements, such as backup data centers and out-of-region staff; and transaction value limits. Other risks for the wholesale payments

8 Financial Crimes Enforcement Network, *Feasibility of a Cross-Border Electronic Funds Transfer Reporting System under the Bank Secrecy Act*, US Department of the Treasury, October 2006, https://www.fincen.gov/sites/default/files/shared/CBFTFS_Complete.pdf.

infrastructure include attacks on the SWIFT messaging system. After recent attacks revealed significant vulnerabilities, SWIFT and its member banks have taken several steps to shore up their defenses, focusing on stronger security standards and quicker response times.

Key cybersecurity risks for retail payment systems include credit and debit fraud, which collectively caused nearly $25 billion in damages in 2018 worldwide; fraud by system insiders affecting platforms like ACH in the United States; and user error, such as falling prey to phishing scams.[9] Risk management strategies for retail payments often rely on voluntary industry standards, such as the PCI DSS. To counter phishing and other types of user error, ACH and other platforms require unique user credentials and offer merchants additional steps like micro validation, tokenization and encryption, and secure vault payments.

Information security is generally assessed along three core principles known as the CIA triad: confidentiality, integrity, and availability.[10] Confidentiality requires that data are only accessible to those who are authorized.[11] For payments, this means that data about participants and their transactions are kept private. Countermeasures to ensure confidentiality focus on areas like authentication, encryption, and educating users.[12] Integrity means that data are "correct, authentic, and reliable" and can thus be trusted to not have been tampered with.[13] This is accomplished via hashing and controlling access.[14] In payments, integrity is linked to the need for non-repudiation: the payor cannot deny sending the payment, and the payee cannot pretend to have not received it.[15] Finally, availability means that the system is up and running, allowing users to have timely and reliable access.[16] In payments, this could be hampered by an attack on a specific institution or by the failure of supporting infrastructure like data centers. Securing availability can be done by hardening systems against attacks and building in redundancy.[17]

LOOKING AHEAD TO CBDCS

Chapter 1 assesses the cybersecurity risks facing CBDCs and how design choices will shape vulnerabilities using a framework derived from the CIA triad but customized to the challenges of CBDCs. Understanding how CBDCs will fit into the existing landscape is crucial for turning this insight into actionable steps for policy makers, which we explore in Chapter 2.

9 "Credit Card Fraud Statistics," SHIFT Credit Card Processing, last updated September 2021, https://shiftprocessing.com/credit-card-fraud-statistics/.

10 "The Three Essentials Pillars of Cybersecurity: Preventing Losses from Cyber Attack," Lexology, https://www.lexology.com/library/detail.aspx?g=03734e1f-98d0-47ef-908f-f29ad6f69a7b.

11 Debbie Walkowski, "What Is the CIA Triad?" F5 Labs, July 9, 2019, https://www.f5.com/labs/articles/education/what-is-the-cia-triad.

12 Ibid.

13 Ibid.

14 Ibid.

15 Ibid.

16 Ibid.

17 Ibid.

Chapter 1: Cybersecurity of CBDCs—Threats and Design Options

This chapter discusses the cybersecurity of CBDCs. A central theme, which pervades all aspects of this chapter, is how CBDCs may centralize data and control over the financial system. Although the current financial system is already relatively centralized (e.g., in the United States, more than 50 percent of banking assets in 2022 are controlled by just four banks),[18] CBDCs have the potential to significantly increase centralization by storing a single ledger or similar data repository that aggregates transaction data from all participants. The ledger could even include data from payment modalities that are currently difficult to monitor, such as cash. Such dramatic centralization of CBDCs could have downstream effects that are difficult to predict or manage. For example, a database containing an entire nation's financial transactions would represent an unprecedented target for cybercriminals. It can also provide unscrupulous regimes with a mechanism for mass surveillance. Such threats can be mitigated in part through technical design choices, but every design comes with implications (and trade-offs) regarding security, privacy, performance, and usability, to name a few. This chapter discusses a landscape of possible design variants, while highlighting the relevant trade-offs.[19]

We start our discussion by introducing the different roles that would be involved in a typical CBDC deployment, their primary tasks, and trust assumptions. After that, we introduce a threat model for CBDCs by discussing the main security requirements and involved threat actors. Then, we review common digital currency variants and analyze them with respect to the established threat model. We complete our analysis with a comparison that shows the main advantages and drawbacks of different currency designs. Finally, through case studies, we show how a few noteworthy CBDC pilot projects fit into our classification. The key contributions of this chapter are as follows.

Key contributions of this chapter

Systemize knowledge: We define a framework for systematically analyzing and comparing digital currency designs. We show the main pros and cons of common digital currency variants and explain how noteworthy existing CBDC pilot projects fit into our classification. We also identify potential cybersecurity risks involved in each currency variant.

Highlight recent research advances: As part of our review, we also highlight recent developments from the research community and possible digital currency design alternatives that are not yet typically considered in most CBDC reports. Such designs can enable improved user privacy or transaction validation scalability, for example.

Clarify common misconceptions: Throughout our discussion, we also point out common misconceptions, recurring harmful practices, or otherwise bad patterns related to the design and deployment of digital currencies.

ROLES AND TRUST ASSUMPTIONS

Currency issuer. Every CBDC system needs an entity that creates money. We call this role currency issuer. In most envisioned CBDC deployments, this role would be played by a central bank. In a private digital currency, this role could also be played by a private company. The currency issuer should be trusted by all system participants for the correctness of money creation. That is, the money created by the issuer is considered valid by everyone involved in the system. This entity does not necessarily need to be trusted for all other aspects of the system, such as user privacy or payment validation.

Payment validator. CBDC systems require entities that keep the system running and provide the needed infrastructure for other participants. One such infrastructure role is the payment validator that approves payments and records them into data storage, such as a database or ledger. The role of the payment validator could be distributed among several nodes for increased security and performance, as will be discussed later in this chapter. The role of the payment validator could be taken by the central bank, or alternatively, it could be delegated to another public authority or to commercial banks. The payment validator needs to be trusted to verify the correctness of payments, but not necessarily for other properties, such as money creation or user privacy.

18 "Large Commercial Banks," Federal Reserve Statistical Release, accessed March 13, 2021, https://www.federalreserve.gov/releases/lbr/current/default.htm.

19 Federal Reserve Bank of Boston, The Federal Reserve Bank of Boston and Massachusetts Institute of Technology release technological research on a central bank digital currency, press release, February 3, 2022, https://www.bostonfed.org/news-and-events/press-releases/2022/frbb-and-mit-open-cbdc-phase-one.aspx#resources-tab.

Figure 1a. Main Roles Involved in a Retail CBDC System

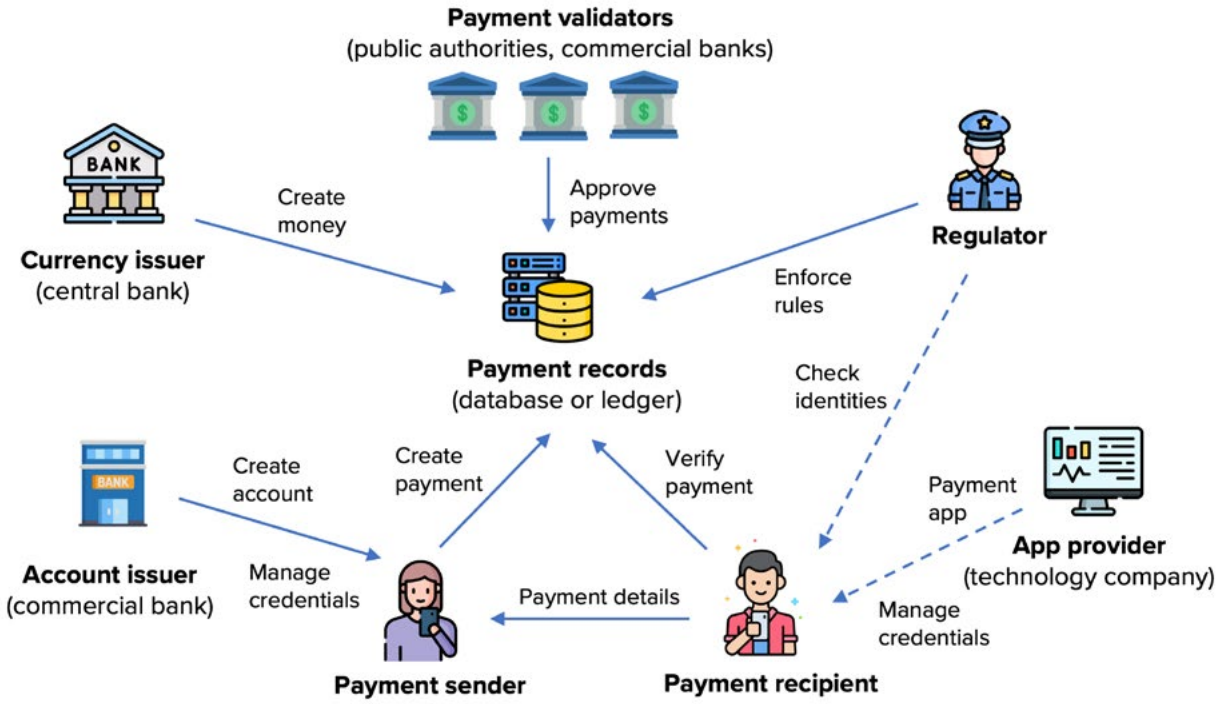

Figure 1b. Main Roles Involved in a Wholesale CBDC System

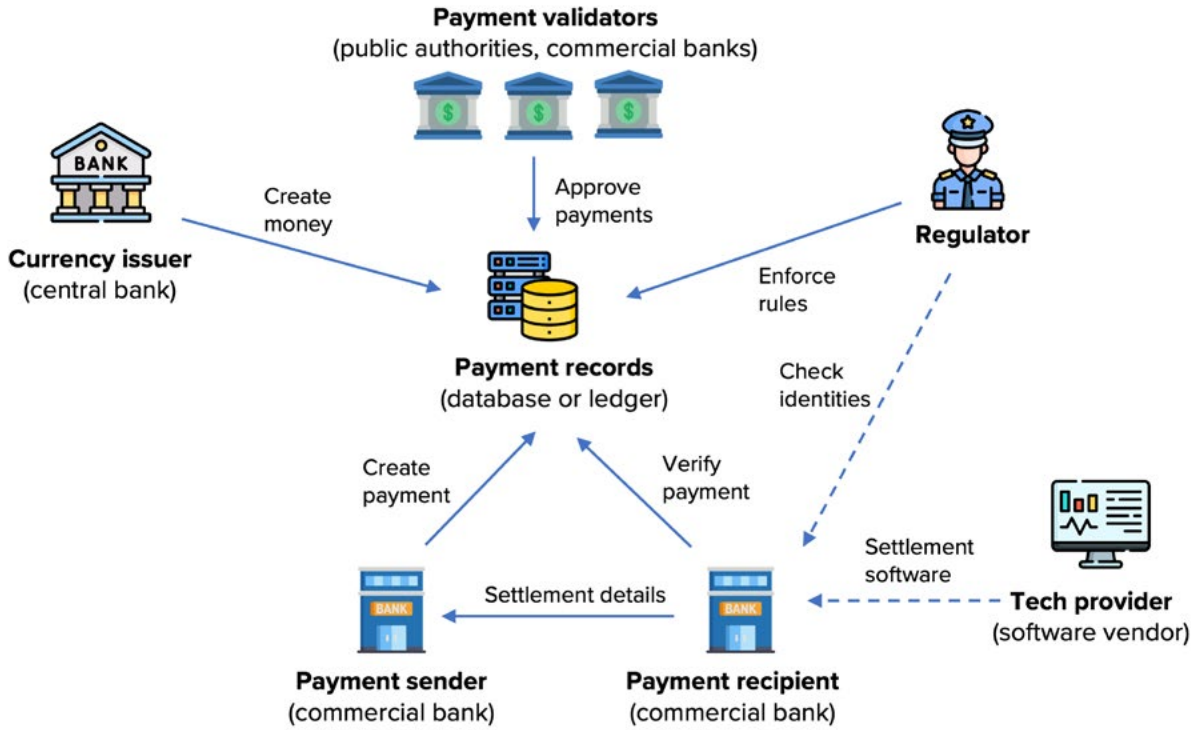

Source: Figure created by Kari Kostiainen with icons licensed from Freepik Company.

Account provider. Another infrastructure role in a typical retail CBDC system is an account provider that allows users to register, obtain payment credentials (e.g., in the form of a digital wallet), and start making CBDC payments. In most retail CBDC deployments, the account issuer would need to verify the identity of the user before account creation. Most likely, central banks would not want to interface with users directly and, therefore, this role would be better served by commercial banks that already have existing customer relationships. The account provider, such as a commercial bank, would need to be trusted for the verification of users' identities. In a custodial solution, the account provider could be also trusted with the management of users' payment credentials and it could control users' monetary assets. In a non-custodial solution, the account provider would not control any monetary assets on behalf of the users. The role of account provider may not be needed in a wholesale CBDC deployment where the end users are financial institutions like commercial banks.

Payment sender and recipient. We consider two types of end users: payment senders and payment recipients. In a retail CBDC system, such users could be private individuals, commercial companies, or other legal entities. Such users would typically perform payments through a client device such as a smartphone that holds the payment credentials obtained from the account provider. For specific use cases like visiting tourists other solutions are likely to be needed for obtaining payment credentials. In a wholesale CBDC, the payment sender and recipient could be commercial banks performing an inter-bank settlement. Payment senders and recipients are generally not trusted by other system participants. Instead, it is assumed that users may behave arbitrarily or even fully maliciously.

Regulator. Another role that we consider is the regulator. The task of the regulator is to ensure that all payments in the system conform to requirements such as anti-money laundering rules. For example, in the United States, the recipients of a cash payment worth more than $10,000 are required to report the payment details to the Internal Revenue Service (IRS). In a CBDC deployment, all payments that exceed a similar threshold amount could be automatically forwarded to the regulator for audit. While the regulator is trusted to examine specific payments and report non-conforming payments, in a well-designed CBDC system, all details of all payments do not necessarily need to be visible to the regulator. For example, receiving $50 fully anonymously (i.e., such that even the regulator cannot see the payment details of the transaction) should be possible. We discuss the challenges involved in realizing such privacy-preserving regulation later in this chapter.

Technology provider. In a retail CBDC, the needed payment application could be provided by a technology company. For example, in the digital yuan pilot in China, the CBDC payment functionality is integrated into popular smartphone payment applications, such as Alipay from Ant Group and WeChat Pay from Tencent. In addition to providing the payment application, in a custodial deployment, the technology provider may assist the user in payment credential management. The end users (i.e., payment senders and recipients) need to trust the technology provider for the correctness of the payment application and potentially also for the management of payment credentials. In a wholesale CBDC, the payment senders and receivers (commercial banks) could obtain the needed (settlement) technology from an external software vendor.

Figures 1a and 1b below illustrate the typical relationships between these roles in retail and wholesale CBDC deployments, respectively.

THREAT MODEL

To understand the cybersecurity implications of CBDCs, it is important to first specify the threat model. In this section, we will highlight the security requirements and the threat actors that are relevant to CBDCs.

Requirements

CBDCs should satisfy a number of properties, both security and performance related. These requirements are intertwined: different design variants can have different implications for each of these requirements.

Integrity. The integrity of a financial system refers to its ability to ensure that money transfers and creation is correct. In other words, it should not be possible to create or delete money out of thin air. It should also not be possible to transfer funds that do not belong to the sender.

Authentication and authorization. Only the legitimate owner of money should be able to transfer said money. In current payment systems, this is typically achieved through a two-step process. Authentication refers to the process of verifying a user's identity.[20] Authorization refers to the process of verifying the transaction details, such as the recipient's identity and the amount to be paid. In some CBDC design variants, these two processes can be intertwined, so we address them jointly in this report.

20 The initial process of linking a digital identifier to a user can be achieved through offline channels, for example. A detailed discussion of this topic is beyond the scope of this report.

Confidentiality. Transactions should not be visible to unauthorized parties (e.g., telecommunications providers). Confidentiality is typically achieved via encryption of data in transport over untrusted channels. Such techniques are widely used in the banking industry today, and we do not expect them to vary significantly across different CBDC variants (though they may need to be updated due to emerging technologies, such as quantum computing). Because of this, we will not analyze confidentiality separately in the remainder of this document.

Privacy. Whereas confidentiality aims to protect data from unauthorized parties, privacy aims to protect user information (e.g., payment transaction details) from authorized parties, such as payment validators. While these two concepts are closely related, we treat them as separate. Deciding what level of privacy to provide is a political decision as well as a technical one, and has repercussions for the architecture and design of the CBDC.

Incorporating privacy protections into a CBDC design is important for two main reasons. The first reason is that the privacy of end users is valuable in itself. CBDCs will inevitably aggregate tremendous amounts of financial data, and consequently some national banks have indicated that their goal is not to build a tool of mass surveillance.[21] Additionally, the successful adoption of CBDC technology may require that the deployed system meets the privacy expectations of end users. In a recent survey on the digital euro, participants rated privacy as the most important feature of a possible CBDC deployment.[22] The second reason is that a system with strong privacy protections is also inherently more secure. If a system that collects huge amounts of sensitive user data does not include privacy protections and is breached, then all the sensitive information will be disclosed to the attacker and, potentially, to other unauthorized parties, which violates confidentiality. In a privacy-preserving design that hides sensitive user data even from trusted system insiders, a similar breach or insider attack will have significantly less severe consequences for security and confidentiality.

Resilience. The system should be robust to faults, or failures, of different components of the system. Typical faults include infrastructure failures (e.g., a server crashes), software-level failures (e.g., a program stops executing), and protocol-level failures (e.g., a validator node misbehaves). Faults can be either accidental (e.g., random infrastructure failures) or intentional (e.g., caused by misbehaving nodes).

TAKEAWAY

PRIVACY-CONSCIOUS DESIGN CAN ALSO PROVIDE SECURITY BENEFITS

If a CBDC deployment without privacy protections is breached, either by an external attacker or a malicious insider, then all the sensitive user information is disclosed to unauthorized parties. In a privacy-preserving CBDC deployment that hides sensitive user data even from trusted system insiders, breaches will have less severe security consequences.

An important aspect of resilience is availability. System availability is often specified in terms of uptime; a common goal is "five nines," i.e., the system is operational 99.999 percent of the time. As a result, the system must be able to process payments even if some parties are offline, including back-end infrastructure, the payment sender, or the payment recipient.

Another relevant dimension of resilience is transaction revertability. Fraudulent transactions are very common in financial systems. Ideally, if a transaction can be shown to be fraudulent, authorized parties, such as payment validators, should be able to revert the transaction, i.e., add the paid amount back to the payment sender's account balance and deduct the paid amount from the recipient's balance.

Network performance and costs. The system must be highly performant to process nation-scale financial transactions. Common performance metrics include throughput (number of transactions that can be processed per second) and latency (time to transaction confirmation). For comparison, the Visa credit card network currently processes 1,700 transactions per second on average and is capable of processing up to 24,000 transactions per second.[23] Meanwhile, typical transaction latencies for digital payments are in the order of seconds.

In exchange, CBDCs will inherently incur communication (or bandwidth) and computation costs. These costs are divided between the back-end infrastructure and end users. In general, a CBDC is expected to impose high costs on back-end infrastructure, both in terms of computation and communication. As such, we do not focus further on back-end resource costs in this report. However, certain potential designs (e.g., privacy-preserving ledgers) require access to the entire ledger, in encrypted form, to verify the validity of transactions. This imposes significant bandwidth requirements on end users, as

21 David Chaum, Christian Grothoff, and Thomas Moser, *How to Issue a Central Bank Digital Currency*, Swiss National Bank Working Papers, March 2021. https://www.snb.ch/n/mmr/reference/working_paper_2021_03/source/working_paper_2021_03.n.pdf.

22 *Eurosystem Report on the Public Consultation on a Digital Euro*, European Central Bank, April 2021, https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf.

23 "Visa Acceptance for Retailers," Visa, accessed May 16, 2022, https://usa.visa.com/run-your-business/small-business-tools/retail.html.

well as substantial computational requirements. These costs must be weighed against the associated privacy benefits.

Governance. The maintenance of a CBDC may involve the participation of multiple parties, including application developers, hardware manufacturers, cloud service providers, and transaction validators. It is important to ensure that these parties have well-designed guidelines for managing operations and conflicts. In addition, all parties should be incentivized to behave correctly and reliably. For example, in the case of distributed transaction validation pipelines, validators should be incentivized to validate transactions promptly and correctly (e.g., in the order they were received), and there should be clear policies in place for managing unfulfilled commitments.

Layers of the technical stack

Attackers can exploit different components of a CBDC to achieve their goals. In this section, we outline the CBDC technical stack, illustrated to the right. In other words, these are the conceptual components that an attacker could target using different vulnerabilities and offensive capabilities. These layers are not exhaustive and attackers can launch cross-layer attacks.

Figure 2. CBDC Technical Stack

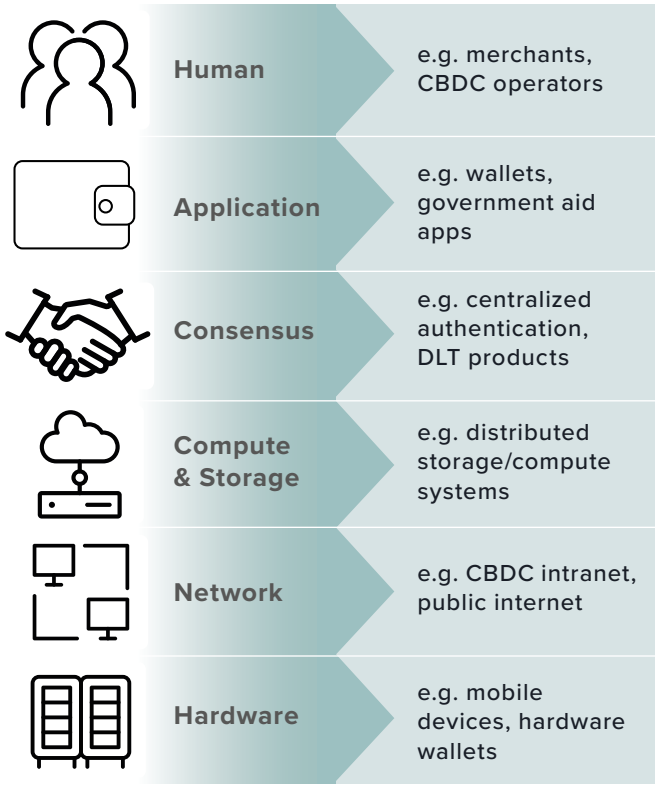

Source: Authors

Human. Although end users are not part of a technical CBDC implementation, they can be exploited to affect system security at large. Users can be both a vector for launching attacks as well as victims. Examples of relevant attacks include fraud and money laundering. Operators of the CBDC can also pose vulnerabilities, e.g., through phishing attacks to gain access to the CBDC's control mechanisms.

Application. CBDCs are expected to usher in an ecosystem of new applications that can interface seamlessly with the digital payment system. Potential use cases include mobile applications for seamless disaster relief, more efficient tax processing, and everyday transaction processing. Many of these applications will likely be developed independent of underlying CBDC infrastructure, just as mobile application developers are typically independent of device issuers. This has several security implications. In particular, it may be difficult to control the security specifications and properties of applications. Developers can introduce vulnerabilities (consciously or not) that can be exploited to steal money or exfiltrate data. While application-level threats or failures may not be directly the fault of the CBDC, they can affect the viability of the CBDC as a whole, as seen in the early release of the eNaira in Nigeria, for example.[24] The application ecosystem is, therefore, an important layer in the CBDC stack from a security perspective.

Consensus. In order to provide redundancy against unforeseen factors like faulty devices, compromised infrastructure, and resource outages, many proposed CBDC designs involve the use of consensus protocols: decentralized processes for determining the validity of financial transactions among multiple payment validators, for example. Consensus protocols can be designed with varying degrees of robustness to adversaries of varying strengths. At a high level, they provide robustness through redundancy: transactions are approved only pending the approval of multiple parties, according to specific, carefully designed protocols. The participants in consensus protocols could be different stakeholders in the systems (e.g., different banks running validator nodes) or they could be different servers controlled by the central bank but running on different infrastructure (e.g., in different data centers). For example, the Swedish e-krona uses distributed ledger technology (DLT) for consensus in which different stakeholders like banks run their own payment validator nodes.[25]

24 Alexander Onukwue, "Nigeria's eNaira Digital Currency Had an Embarrassing First Week," Quartz, October 28, 2021, https://qz.com/africa/2080949/nigerias-enaira-android-wallet-deleted-days-after-launch/.

25 "E-kronapiloten – test av teknisk lösning för e-krona" ["The e-Krona Pilot – Test of Technical Solution for the e-Krona"], Sveriges Riksbank, last updated April 6, 2021, https://www.riksbank.se/sv/betalningar--kontanter/e-krona/teknisk-losning-for-e-kronapiloten/.

Attacks on the consensus protocol typically involve the corruption of one or more parties. Good protocols are designed to be robust up to some threshold number of corruptions. However, consensus protocols are notoriously subtle; to provide true robustness to malicious faults, they should be accompanied by mathematical security guarantees. Further, even when those security guarantees exist, they rely on assumptions about the adversary that may not hold in practice (e.g., many protocols assume the adversary can only corrupt up to one-third of all validator nodes).

Computation and storage. CBDCs require back-end infrastructure to maintain a secure and functional payment system. For example, they may require distributed computation nodes to parallelize transaction processing in the face of stringent performance requirements. To the extent possible, ledger storage may also be distributed to reduce the load on any single node. However, some security mechanisms are easier to parallelize than others. For example, ledger-based systems typically require the full ledger to ascertain transaction validity; hence splitting the ledger into shards can affect the system's ability to correctly validate transactions.

Network. The validation of transactions, issuance, deletion of money, and all other events in a CBDC will be communicated to the relevant parties via an underlying network. This network will very likely rely at least in part on private infrastructure to communicate updates among payment validators and CBDC internal parties. Interactions between account providers and end users will likely occur on the public Internet. These networks can be used to launch attacks such as denial of service, censorship attacks, or even partitioning attacks that cause different parts of the network to have different views of the global state. This causes the network layer to interact with the consensus layer.

Hardware. CBDCs will ultimately run on hardware, including mobile devices, hardware wallets, and servers that maintain the state and functionality of the system. Hardware can become an attack vector through insecure firmware and/or vulnerabilities that are hard coded into the products (e.g., backdoors in a hardware wallet). Such vulnerabilities tend to be difficult to exploit by all but the most sophisticated adversaries.

Threat actors

Security is defined with respect to a particular adversary. In a CBDC, there are several potentially adversarial actors of interest. We consider the following, in increasing order of strength.[26]

Table 1: Which Layers of the CBDC Stack Can Different Adversaries Access or Corrupt?

	Users	Third Parties	System Insiders	Nation States
Human	●	●	●	●
Application		●	○	●
Consensus			●	○
Compute & Storage		○	●	○
Network		○	○	○
Hardware		●	○	●

Source: Table created by Giulia Fanti.

Note: Solid circles indicate (the potential for) full access, whereas outlined circles indicate the potential for partial access.

Users. Users are typically limited in their ability to affect the internal mechanics of the CBDC. They are generally able to access and exploit applications only to the extent that they can manipulate other users. End users may be motivated to steal money from other users.

Third parties. Various types of third parties can threaten a CBDC, including scammers, application developers, or hardware manufacturers. Such adversaries are generally more powerful than typical end users, with more resources to attack the CBDC at layers ranging from hardware to application. For example, they may release malicious applications into the ecosystem, or manufacture backdoored hardware wallets. Their motives may range from stealing money to destabilizing the currency (e.g., particularly at the behest of a nation-state).

System insiders. Insiders refer to individuals (or groups of individuals) who have access to the internal operations of a CBDC, including infrastructure operators or CBDC developers; their capabilities range from modifying system-critical code to exfiltrating data to bringing down key infrastructure (e.g., unplugging servers). Such attackers are notoriously difficult to defend against. Their motivations can be political, financial, or even personal. Common goals of malicious insiders include stealing resources or simply bringing the system to a halt.

26 Note that in security analysis, threat modeling typically considers an adversary's means (what are their capabilities?) as well as their motives (what are they trying to achieve?). We, therefore, discuss both.

A key lock is placed in cyberspace.

Foreign nation-states. Foreign nation-states are among the most powerful adversaries that a CBDC must defend against. Such adversaries may have effectively limitless resources to spend on offensive tactics, including the development of zero-day attacks as well as deployment of sophisticated attacks on applications, operating systems, and hardware. Additionally, they may coerce third-party producers of hardware or software to hard code backdoors into products, thus giving easier downstream access. Such attacks can affect payment validator nodes, end user wallets, and custodial wallets hosted by account providers, to name a few. Their motivations are typically assumed to be political in nature.

Attack matrix

Different threat actors have different capabilities for infiltrating a CBDC. Table 1 indicates which attackers have access to which portions of the CBDC stack. Here, solid circles indicate that there exists the potential for full corruption of at least some portion of a given layer, whereas half-filled circles indicate the potential for partial corruption. Notice that all of the adversaries have only partial access to the network layer because CBDCs will rely in part on the public Internet. As such, full corruption is believed to be infeasible even for foreign nation-states. On the other hand, hardware is most easily corrupted through supply chain attacks, which can be executed by third parties as well as nation-states.

CBDC DESIGN VARIANTS

In this section, we discuss major design choices related to cybersecurity for CBDC systems. The space of CBDC designs is vast, with each design presenting its own trade-offs.[27] We present six digital currency variants that could form the basis of a CBDC system. This review does not attempt to cover all possible designs, but rather to give representative examples of different styles of digital currency schemes. For each design variant, we summarize the security, privacy, and performance trade-offs according to our requirements from the previous section. The design variants we discuss are reflected in Figure 3; orange boxes represent design variants, and blue boxes represent differentiating factors. Additionally, each design variant is annotated with one or more new cybersecurity challenge that arises in this CBDC design compared to the current financial system. These challenges are summarized below.

Database with account balances (status quo)

We start our review with a simple payment system that we call database with account balances. This design variant captures the payment approach used by the existing credit card payments, mobile payments, and bank account transfers. We assume that both the payment sender and the payment recipient have already established an account and obtained the needed payment credentials. We also assume a database (payment

27 James Lovejoy et al., "Project Hamilton Phase 1: A High Performance Payment Processing System Designed for Central Bank Digital Currencies," Federal Reserve Bank of Boston, February 3, 2022. https://www.bostonfed.org/-/media/Documents/Project-Hamilton/Project-Hamilton-Phase-1-Whitepaper.pdf.

NEW CYBERSECURITY CHALLENGES FOR CBDCS

The design variants discussed here pose various cybersecurity challenges that differ from challenges seen in the current digital financial system.

1 **Financial data can be more centralized.** Some design variants rely on a single, centralized database of financial transactions that is visible to system operators. This presents a central point of failure and a unified target for potential attackers. Although such databases exist with digital payments today (e.g., credit cards), CBDCs present an even greater potential for data centralization, and hence increased cybersecurity risk.

2 **Regulatory agencies have less visibility into data.** Some design variants prevent regulatory or law enforcement agencies from accessing transaction data, typically because said data is encrypted or stored only on local devices. This reduces regulators' visibility into financial transaction flows compared to the current digital financial system and has implications for tracking illicit transactions, for example.

3 **Security hinges on the integrity of third-party validators.** Some design variants use third-party validators (e.g., banks, telecommunications providers) to validate transactions.[1] Transaction integrity is dependent on a (super-)majority of these validators not being compromised. This poses new challenges in terms of auditing and monitoring validators, as well as coordinating incident responses across validators, who may have different policies and procedures for dealing with breaches.

4 **Client key custody becomes more complicated.** Some design variants require transactions to remain encrypted to provide client privacy. Custodial key management solutions, which are commonly used in the current financial system, would, therefore, compromise the promised privacy guarantees because the custodian could access client financial data. This requires client-side key management tools, which can present significant usability challenges. This problem has materialized in many cryptocurrencies and remains prevalent.[2]

5 **Security relies on trusted hardware manufacturers.** Some design variants use trusted hardware to enforce transaction integrity. This places an increased supply chain risk specifically with trusted hardware manufacturers compared to the current financial system.

6 **Transaction revocation is more difficult.** Some design variants prevent an authority from unilaterally revoking fraudulent or contested transactions. This could be because client keys are stored locally, because there are multiple validators, or because data is encrypted so the central database is unable to ascertain the amount and endpoints of a contested transaction.

7 **Programmable transactions can amplify the scope and scale of errors.** Applications built on CBDCs are expected to rely on programmable transactions, or smart contracts (these are explained in more detail at the end of this chapter in the section on Other Design Choices). Incorrectly specified smart contracts could result in misdirected funds at a massive scale, especially if these smart contracts are deployed naively. When coupled with Risk 6 (difficulty revoking transactions), this could lead to substantial financial losses.

1 Technically, it would be possible for all validators to be run by the central bank. However, most deployments have chosen to run validation with a coalition of third parties.

2 Tim Copeland, "96 Private Keys Stolen from Vulcan Forged in $140 Million Theft," Block, December 13, 2021, https://www.theblockcrypto.com/post/127270/96-private-keys-stolen-from-vulcan-forged-in-140-million-theft.

records in Figures 1a and 1b that maintains an account balance for each user.

To initiate a payment, the sender first requests the payment details, such as account number, from the payment recipient. In the case of card payment, this would happen during interaction with the recipient's payment terminal. In the case of a bank account transfer, the payment details could be obtained manually or by scanning a QR code. Then, the payment sender creates a payment request that defines the identity of the recipient and the payment amount, signs the payment request using their payment credentials, and sends it to the payment validators. The payment validators check from the database that the sender has sufficient funds associated with their account, and if that is the case, update the account balances of the sender and the recipient accordingly. This process may be distributed among multiple nodes for resilience and performance reasons, as in the recent Project Hamilton proposal.[28] Finally, the payment validators send a payment completion acknowledgment to the payment recipient.

28 Lovejoy et al. "Project Hamilton Phase 1: A High Performance."

Figure 3. CBDC Design Variants Discussed in This Chapter

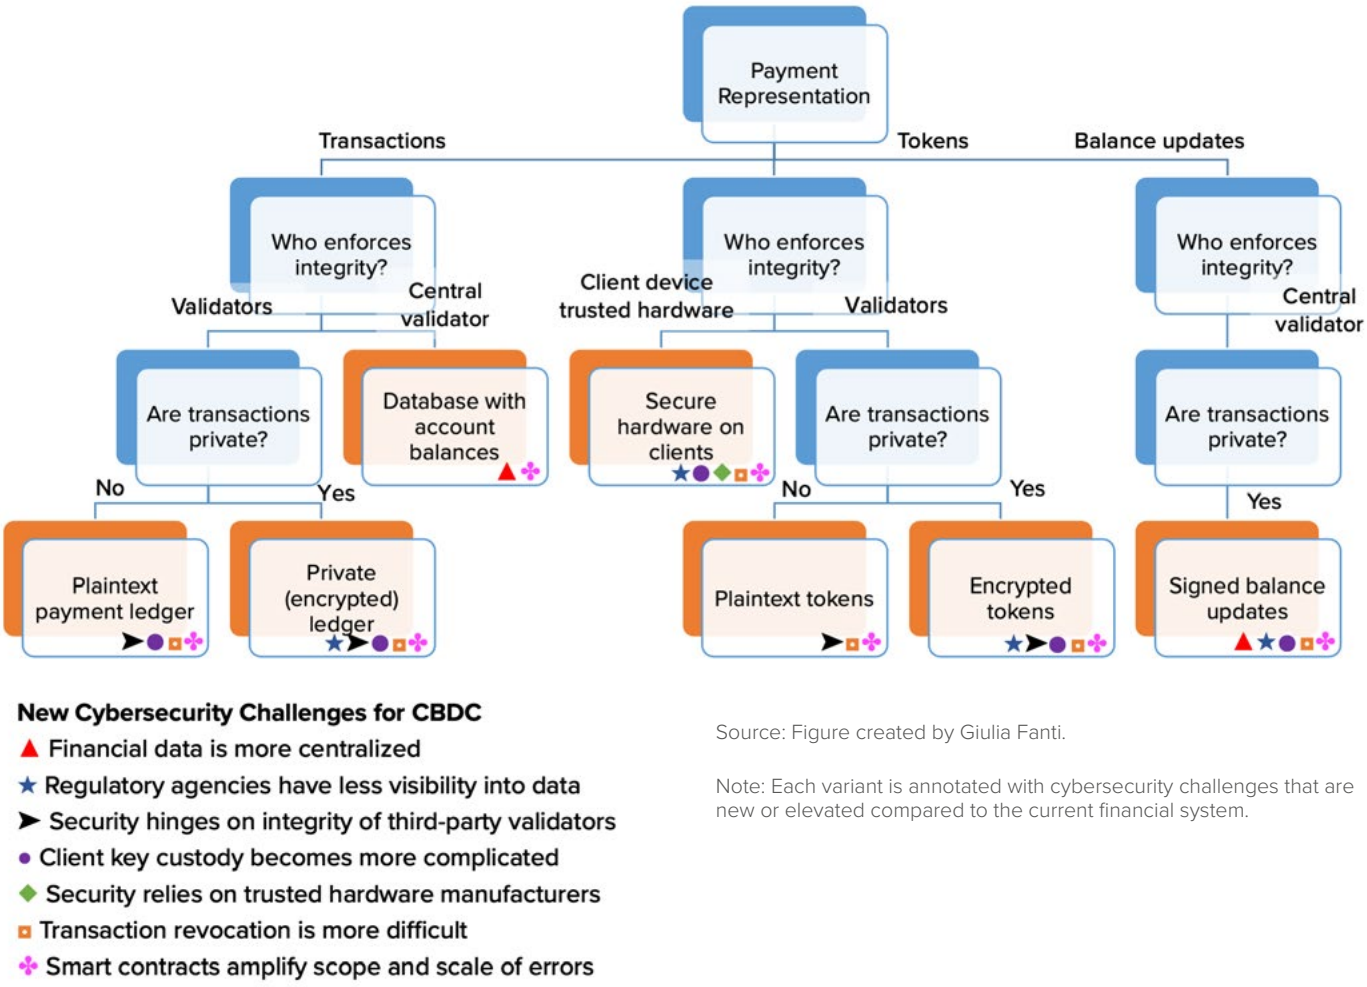

Source: Figure created by Giulia Fanti.

Note: Each variant is annotated with cybersecurity challenges that are new or elevated compared to the current financial system.

TAKEAWAY

THE DESIGN SPACE FOR DIGITAL CURRENCIES IS LARGE

The discussion in many CBDC reports focuses on currency designs that are based on a centralized database, distributed ledger, or token model. We argue that the design space for digital currencies is larger than that. As will be discussed below, a digital currency can also be realized as signed balance updates or as a set of trusted hardware modules, and both the distributed ledger variant and the token model can support privacy-preserving transactions in addition to plaintext ones.

In this approach, all payment details necessary for validation are visible to the payment validators. The payment database stores the latest account balance for each user and such account balances are not disclosed to the public (only validators know the balance of each user account).

Analysis

Integrity. The integrity of the database is entirely governed by the payment validator(s), who must (collectively) check that users do not overdraw their accounts. Assuming the currency issuer and payment validators perform these operations correctly, no money can be created out of thin air and no money will disappear from the system. To violate payment integrity, either an adversarial insider would need to manipulate the operations of the currency issuer or a sufficient number of payment validators' nodes or an external adversary would need to compromise these entities through remote attacks.

Authentication and authorization. Users are authenticated upon logging into the system. Payments are authorized when an (authenticated) sender approves a transaction within a secure payment application. The easiest way for an adversary to break payment authorization is to compromise the initial authentication process, for example, through phishing attacks or malware. These threats can be mitigated through multi-factor authentication (MFA), including the use of hardware tokens.

Privacy. The database model inherently provides no privacy to users. In terms of privacy, this design variant is comparable to current credit card and smartphone payments where the payment processors learn all transaction details. Any party with access to the database (i.e., payment validators) can see all transaction details: sender, receiver, amount, and time. If privacy is desired, it must be accomplished through non-technical means, such as implementing strict access control policies that prevent internal operators from accessing this data without approval. Hence, the primary attacks on privacy will be at the human layer, by corrupting operators and processes.

Resilience. To process a payment, the sender only needs to submit the payment to the validator nodes. The receiver does not need to be online and can retrieve the funds the next time they access their wallet. However, to preserve availability, the validator infrastructure must be active at all times to confirm incoming transactions. Attacks on availability in this model are likely to target underlying infrastructure layers (e.g., network, storage, and/or compute). Transaction revocation is straightforward and can be executed unilaterally by the database operator (similar to credit card payments today).

Network performance. In terms of throughput, this design is very scalable and flexible. In particular, it can be implemented in a fully centralized fashion. This removes a major bottleneck to scaling throughput: communication bandwidth constraints. In this setting, we can feasibly achieve throughput comparable to existing financial services like banks or credit cards.

TAKEAWAY

CBDC DEPLOYMENT MIGHT CENTRALIZE USER DATA COLLECTION

The main difference between a database with account balance CBDC and the current financial system is that the CBDC may result in a greater centralization of user data and financial infrastructure. This can have advantages, such as greater efficiency in implementing monetary policy. It can also have disadvantages, including the privacy threat of storing a single database containing users' (or banks') every transaction.

CASE STUDY: JAM-DEX (Jamaica)

In 2021, the Bank of Jamaica ran a pilot of retail CBDC with vendor eCurrency Mint Inc. The Bank of Jamaica specifically chose to avoid blockchain technology for this pilot not because of technical misgivings, but in order to seamlessly interface with existing payment structures within the nation.[1] Over the course of the pilot, the bank issued CBDC to banks and financial institutions as well as small retailers and individuals. After continuing these trials in early 2022 to test interoperability and transactions between clients and wallet providers, the Bank of Jamaica announced a phased launch of the Jamaican Digital Exchange (JAM-DEX) in May 2022.[2]

Advantages

Centralized databases are a mature technology and can in many cases be more easily integrated with existing infrastructure.

Risks

A primary risk is related to privacy; this architecture exposes all users' transactions in plaintext to the Bank of Jamaica. Even if the bank itself does not abuse this information, the transaction database poses an attractive target for hackers. The consequences of a data breach in a centralized setting may be very serious.

1 Natalie Haynes, "A Primer on BOJ's Central Bank Digital Currency," Bank of Jamaica, accessed March 31, 2022, https://boj.org.jm/a-primer-on-bojs-central-bank-digital-currency/.

2 Bank of Jamaica, "Bank of Jamaica's CBDC Pilot Project a Success," Jamaica Information Service, December 31, 2021, https://jis.gov.jm/bank-of-jamaicas-cbdc-pilot-project-a-success/.

This model has potentially the lowest communication costs overall. If implemented as a centralized service, transactions do not need to be validated by multiple parties. This reduces back-end communication costs. End users do not need to store any data except their own; this minimizes user-facing communication costs. Note that a "centralized design" can still boost throughput through parallelization.[29]

Governance. The governance requirements for this design are equivalent to those of the current financial system. In particular, as the system is centralized, there is no need to manage the threat of misbehaving validators. However, there is still a need for well-documented policies governing incidents at various layers of the stack, including insider attacks.

29 Federal Reserve Bank of Boston, The Federal Reserve Bank of Boston and Massachusetts Institute of Technology.

Distributed ledger with plaintext transactions

Another popular design variant that we consider captures the way payments work in the currently popular public blockchain systems like Bitcoin and Ethereum. We call this approach distributed ledger with plaintext transactions. As above, the payment process starts such that the sender obtains the payment address of the recipient. The payment sender prepares a transaction that includes the payment details (sender and recipient identities, payment amount) in plaintext, authorizes the payment by signing the transaction with their payment credentials, and sends it to the validators. The validators check that the sender has sufficient funds (such a check is trivial because all payment details are in plaintext in the transaction) and then append the payment transaction into a ledger that records all the transactions of the system.

In a public payment scheme like Bitcoin and Ethereum, the sender (or any other third party) can verify that the payment was approved by checking that it appears in the public ledger. For a CBDC deployment, most likely the ledger would be private and only accessible by authorized parties like the payment validators, currency issuer, and regulator. In such private ledger deployment, the payment recipient could verify the completion of the payment by querying the payment validators, instead of verifying the payment directly from the ledger.

Analysis

Integrity. The integrity of a ledger with plaintext transactions relies on two properties. First, regular transactions must not draw upon funds that have already been spent. This is verified by checking the transaction source against the set of all unspent transactions from the ledger. To bypass such a check, a sufficient number of validator nodes need to be manipulated or compromised. Second, the payer must be authorized to spend the transaction; the next paragraph explains how to verify this. In the case of minting new money, the first condition is not relevant, as the money is being created; authorization is still essential, though.

Authentication and authorization. This design variant can involve separate authentication and authorization processes, but they can also be merged. Payment authorization requires a cryptographic signature on the transaction. Hence, for each payment, validators must verify that the signature is valid (which is itself a form of authentication, as cryptographic keys are meant to be linked to a specific user) and authorized to spend the money in question. The easiest attack on authorization is for an adversary to steal a user's private keys, for example via phishing attacks. More recent "ice phishing" attacks trick users into signing a transaction that delegates the right to spend a user's tokens.[30]

Privacy. Ledgers with plaintext transactions do not inherently provide privacy to the transaction sender or receiver. Payments will be visible to any party with access to the ledger, including (at least) account issuers. At best, the system can provide pseudonymity with respect to parties that have access to the ledger; in other words, users are represented by pseudonymous public keys, and privacy is maintained only as long as these keys cannot be linked to a real-world identity. However, pseudonymity guarantees are known to be easily broken;[31] moreover, providing pseudonymity with respect to account issuers inherently complicates Anti-Money Laundering (AML) and Know Your Customer (KYC) efforts, and may, therefore, be less favored.[32]

Resilience. To process a payment, the sender only needs to submit the payment to the validator nodes. Notably, the receiver does not need to be online and can retrieve the funds the next time they access their wallet. However, to preserve availability, the validator infrastructure must be active at all times to confirm incoming transactions. In this design variant, transaction revocation can be more complex. For example, suppose a transaction sender requests that the transaction be revoked by appealing to their bank (which happens to be operating a validator node). However, the transaction receiver may argue to their bank (also a validator) that the transaction should stand. In this case, no bank has the authority to unilaterally revoke the transaction, absent legal or policy frameworks for handling such situations. Such challenges can be mitigated if a central authority (in this case, the central bank) is given the authority to revoke transactions and freeze assets. However, this requires the central bank to be directly involved in dispute resolution. Moreover, it changes the core threat model by involving a central trusted party in the validation process, thereby introducing a central point of failure.

Network performance. Ledger-based designs inherently require sequential processing that can limit throughput. In particular, validators must verify that each transaction is not drawing on previously spent funds. The only fully safe way to ensure this is by serially processing every transaction. Although there has been work in the research community showing how to

30 Microsoft 365 Defender Research Team, "'Ice Phishing' on the Blockchain," Microsoft, February 16, 2022, https://www.microsoft.com/security/blog/2022/02/16/ice-phishing-on-the-blockchain/.

31 Josyula R. Rao and Pankaj Rohatgi, "Can Pseudonymity Really Guarantee Privacy?" 9th USENIX Security Symposium Paper, 2000, https://www.usenix.org/events/sec2000/full_papers/rao/rao_html.

32 Bank of England, "Central Bank Digital Currency: Opportunities, Challenges and Design," Discussion Paper, March 12, 2020, https://www.bankofengland.co.uk/paper/2020/central-bank-digital-currency-opportunities-challenges-and-design-discussion-paper.

CASE STUDY: Digital Won (South Korea)

In 2021, the Bank of Korea announced plans to pilot a digital won. This pilot study, which started in late 2021, is an example of a distributed ledger with plaintext transactions. It is running on a Klaytn ledger,[1] which uses a custom DLT consensus protocol that was initially proposed for the Ethereum blockchain.[2] The validators in this blockchain are currently being run by various companies, including banks and payment providers. The technology is being provided by GroundX, which is the blockchain unit of Korean communications giant Kakao.

Advantages

The use of DLT technology can provide better integrity against certain adversaries. Specifically, decentralized validation protects against the threat of corrupt insiders arbitrarily modifying, rejecting, or creating transactions.

Risks

The DLT consensus protocol used by Klaytn, while derived from well-established consensus protocols, is relatively untested and has not been publicly peer reviewed (to the best of our knowledge). Indeed, early versions of this consensus protocol had design errors that affected the integrity and robustness of the system.[3] DLT consensus protocols are notoriously subtle to design, and care should be taken with new, untested protocols.

User privacy may be limited, as Klaytn user accounts are associated with (internally visible) user-selected addresses. However, exploring privacy implications is one of the objectives of Phase 2 of the pilot study, scheduled to terminate in June 2022.

1 "Consensus Mechanism," Klaytn, accessed May 16, 2022, https://docs.klaytn.foundation/klaytn/design/consensus-mechanism.

2 ConsenSys, "Scaling Consensus for Enterprise: Explaining the IBFT Algorithm," June 22, 2018, https://consensys.net/blog/enterprise-blockchain/scaling-consensus-for-enterprise-explaining-the-ibft-algorithm/.

3 Roberto Saltini, "IBFT Liveness Analysis," 2019 IEEE International Conference on Blockchain, 245–252, 10.1109/Blockchain.2019.00039.

achieve high throughput in such a setting,[33] these systems tend to add implementation complexity. Alternatively, account issuers (e.g., banks) may be willing to parallelize the processing of smaller transactions to achieve higher throughput, at the risk of allowing double-spending. This risk can be managed through non-technical means, such as insurance.

Communication costs will depend in part on architectural decisions. In the lowest-trust setting, validators should download the entire ledger to verify correctness. Some designs involve a tiered system, where certain nodes store the full ledger history, whereas others store only the system state that is relevant to them (e.g., a single user's set of unspent transactions); in these tiered systems, so-called light clients may store only information relevant to their own needs, and outsource transaction verification to third parties to avoid the storage and bandwidth costs of maintaining the full ledger. In a CBDC, users are likely to want this light client functionality to transact from lightweight devices like a mobile phone; in this case, account provider(s) may play the role of the trusted third party, much as in the current financial system.

Governance. This design introduces the need for independent validators. As such, it is important to establish policies that govern situations in which one or more validators misbehave (e.g., approving invalid transactions, changing the order of transactions, or not meeting promised availability or latency guarantees). These policies can be retroactive, punishing entities that misbehave. They can also be proactive, by establishing mechanisms that incentivize validators to correctly and promptly validate transactions. Common examples of such mechanisms include transaction fees, which reward validators for each transaction processed, and block fees, which reward validators for processing a batch of transactions. A third possibility is to allow validators to accrue interest from a reserve pool, which is invested independent of the currency; this was the approach suggested for Libra, now Diem, Meta's proposed digital currency.[34] Another important governance issue is related to the interface between central banks and independent validators, such as banks or other financial institutions. For example, in the event of policy changes internally to the CBDC, do validators have a say, or will changes be imposed unilaterally by the central bank? How much information should validators share with each other and with the central bank, and at what timescales? We touch on these questions in Chapter 2.

33 Vivek Bagaria et al., *Prism: Deconstructing the Blockchain to Approach Physical Limits*, CCS '19, November 11-15, 2019, London, United Kingdom, https://dl.acm.org/doi/pdf/10.1145/3319535.3363213; Haifeng Yu et al., "OHIE: Blockchain Scaling Made Simple," 2020 IEEE Symposium on Security and Privacy, https://ieeexplore.ieee.org/iel7/9144328/9152199/09152798.pdf; and Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman, "Asymptotically Optimal Validated Asynchronous Byzantine Agreement," proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, July, 2019, 337–346, https://doi.org/10.1145/3293611.3331612.

34 Zachary Amsden et al., *The Libra Blockchain*, MIT Sloan School of Management, revised July 23, 2019, https://mitsloan.mit.edu/shared/ods/documents?PublicationDocumentID=5859.

Distributed ledger with private transactions

The next design alternative that we consider captures how private blockchain systems like Monero and Zcash work. We call this design alternative distributed ledger with private transactions. In this approach, the payment sender prepares the payment transaction such that payment details like identities and amounts are hidden. In practice, payment details can be hidden using encryption or cryptographic commitments. Additionally, the payment sender computes a zero-knowledge proof that allows the payment validators to verify that the transaction updates the user's funds correctly without learning the payment details. More precisely, the zero-knowledge proof shows to the verifier that the sender has sufficient funds, the balances of the sender and the recipient are updated correctly by the transaction, and the proof is created by the legitimate owner of the funds (i.e., payment integrity and authorization hold).

The payment sender uploads such private transactions to the payment validators who will verify the zero-knowledge proof without learning any payment details. If the proof is correct, the validators include the transaction in the ledger. As above, the payment recipient can verify the completion of the payment either by reading it directly from a public ledger (as is done in systems like Zcash and Monero) or by querying the payment validators (as would be more likely in a CBDC deployment with a private ledger).

TAKEAWAY

STRONG USER PRIVACY IS POSSIBLE

Recent reports on CBDCs imply that a CBDC would inherently provide weaker privacy than cash.[1] To some extent, we agree that such a view is justified. In any CBDC realization, a payment transaction would leave some digital trace (e.g., a communication channel opened between the payer and the payment infrastructure). However, we argue that such a view is also an oversimplification. The use of modern cryptographic protections, such as encryption, commitments, and zero-knowledge proofs, enables digital currency designs where even the payment validators who process and approve transactions do not learn the identities involved in the payment or the payment amount or cannot link payments from the same individual together. For many practical purposes, such strong privacy protection is comparable to the privacy of cash.

1 "Central Bank Digital Currencies: A Solution in Search of a Problem?" Economic Affairs Committee, UK Parliament, January 13, 2022, https://committees.parliament.uk/committee/175/economic-affairs-committee/news/160221/central-bank-digital-currencies-a-solution-in-search-of-a-problem-report-published/.

Analysis

Integrity. The integrity of a private ledger relies on the same two properties as public ledgers: the transaction should draw on valid funds, and the sender should be authorized to send (or create) the funds in question. In this model, payments are accompanied by cryptographic proof (e.g., a zero-knowledge proof) that proves the funds can be spent. Hence, verifying integrity involves checking that a zero-knowledge proof is valid. Creating and checking these proofs incurs additional computational overhead compared to plaintext ledgers, but these overhead costs have been falling in recent years thanks to innovations in applied cryptography.[35]

Authentication and authorization. Unlike public ledgers, this design variant aims to break the linkage between users and their transactions. As with public ledgers, payment authorization requires a signature or a similar cryptographic operation using a key or credential that is only known to the owner of the assets (payment sender). The cryptographic operation is such that system insiders, such as payment validators, cannot link this payment authorization to the identity of the payment sender. Therefore, in this design variant there is no explicit authentication process.

Privacy. Ledgers with private transactions are designed to protect both the transaction sender and receiver. Such approaches can prevent an observer from linking the sender or receiver to a given transaction, while also hiding the amount of a given transaction. In this model, the ledger is still fully available to all validation nodes, but in encrypted form. Notably, these techniques do not protect against privacy attacks at the network layer—only at the consensus and application layers. Generally speaking, ledger-based private transactions cannot be easily reverted, because the payment validators do not learn the identities of the transacting parties in the fraudulent transaction.

Resilience. As with public ledgers, only validators and the sender need to be online to process a transaction. The receiver does not need to be online and can retrieve the funds the next time they access their wallet. In particular, validators should be active at all times to ensure the system remains operational. As with plaintext distributed ledgers, transaction revocation can be complicated by the presence of multiple validators. Additional challenges arise in the case of revoking transactions on distributed private ledgers because validators do not have visibility into the amounts and/or parties involved in a transaction.

35 Paige Peterson, "Reducing Shielded Proving Time in Sapling," Electric Coin Co., December 17, 2018, https://electriccoin.co/blog/reducing-shielded-proving-time-in-sapling/.

Network performance. As before, the sequential processing associated with ledgers can limit throughput. In particular, validators must verify that each transaction is not drawing on previously spent funds. As mentioned above, checking zero-knowledge proofs does incur some extra computational overhead compared to checking the validity of plaintext transactions. Today, the Zcash cryptocurrency uses schemes that require a few seconds to generate a proof (needed to create a new transaction), whereas transaction validation takes only milliseconds.[36] These schemes are close to the state-of-the-art today. This additional processing time primarily affects transaction latency rather than throughput.

The communication costs of this model are high. As with plaintext ledgers, transaction validation requires access to the (encrypted) system state to validate transactions. This requires validator nodes to download large quantities of data in continuation. However, unlike in plaintext ledgers, light clients are difficult to implement as existing designs effectively break the promised privacy guarantees.

Governance. This design has all the same governance requirements as the ledger with plaintext transactions, particularly regarding the interactions between private validators and the central bank. Although the ledger is encrypted in this setting, many types of validator misbehavior can be detected just as easily as in the plaintext setting. For example, validators who validate conflicting transactions (thereby violating integrity) can still be detected as their digital signatures are visible to other validators and can be linked to the originator.

Plaintext payment tokens

The next design variant that we consider is a token-based payment system. In such a system, the payment sender withdraws digital coins (that function as payment tokens) from the currency issuer. This withdrawal operation is authenticated using the sender's credentials and the currency issuer updates the account balance of the user based on the withdrawn amount. Each coin (token) has a specific denomination and a unique serial number.

To create a payment, the payment sender passes an appropriate number of coins (tokens) to the payment recipient who verifies that each coin is correctly signed and that their total denomination corresponds to the expected payment amount. To prevent double-spending of coins, the payment recipient deposits the coins to the payment validators immediately. The payment validators maintain records of the already used serial numbers and check that the serial numbers in the deposited coins have not been already used. After that, the payment validators add the amount of the deposited coins to the balance of the payment recipient and inform the recipient that the payment has been accepted.

Analysis

Integrity. The integrity of a digital cash scheme relies on the correctness of the following two operations. First, when the payment sender withdraws coins from the currency issuer, the issuer must update the account balance of the user with an amount that matches the denomination of the withdrawn coins. Second, when the payment recipient deposits the received coins, the payment validators must check that the serial numbers of the coins have not already been used, and then update the account balance of the recipient with the denomination of the deposited coins. Assuming that the currency issuer and payment validators perform these operations correctly, no money can be created out of thin air and no money will disappear from the system. To violate payment integrity, either an insider adversary would need to manipulate the operation of the currency issuer or a sufficient number of payment validators, or an external adversary should be able to compromise these entities through remote attacks.

Authentication and authorization. In this design variant, anyone who holds coins (tokens) is able to authorize a payment by simply passing coins to a payment recipient. Sender authentication occurs when a user withdraws coins. Recipient authentication occurs when the user deposits received tokens. It is noteworthy that, unlike in most digital currency solutions, payment authorization does not require an explicit cryptographic operation like signing (only passing tokens from one entity to another). To break payment authorization, the adversary would need to steal coins from the user. Assuming that the coins are stored in the user's wallet hosted on their smartphone, this might be possible by either stealing the device or tricking the user to install malicious software on the device.

Privacy. Plaintext payment token systems do not provide privacy for the end users. The payment validators learn the payment amount and identity of the payment receiver during the coin deposit operation. Due to unique serial numbers, payment validators can link deposit operations to previous withdrawal operations, and thus also learn the identity of the payment sender.

Resilience. To perform a payment, the payment sender needs to contact the payment recipient, and the payment recipient needs to be online in order to deposit the received coins. In principle, the payment recipient can accept coins fully offline (and deposit them later), but in such a case, there is no double-spending protection, and thus payment acceptance is not safe. Because payments are processed by distributed validators in this design variant, transaction revocation may be more complicated.

36 Ibid.

CASE STUDY: E-Krona (Sweden)

In 2019, Sveriges Riksbank began planning the possible design of a CBDC, called e-krona, and investigating the regulatory implications of such a deployment. In 2020, together with Accenture as the technology provider, Riksbank started a CBDC pilot where one possible design alternative was tested.[1]

The piloted design follows the plaintext payment token approach where users withdraw coins (tokens), then make payments by passing them to the payment recipient who deposits them back to the payment infrastructure to verify the coins have not already been used (double-spending protection).

Advantages
One advantage of the piloted design is that it is easy to scale. Separate payment validators can verify separate ranges of coin serial numbers without having to run a complicated and expensive consensus protocol. This makes payment verification fast and easy to scale for a large number of parallel validators.

Risks
Compared to ledger and database variants, a token or coin-based design places a higher burden on the user for wallet management. If the wallet that stores the coins is lost, the user will lose all funds. In most other currency variants it is sufficient to securely manage and back up one key that is used to authorize payments.

Additionally, the piloted design provides no privacy protection for the users, and, therefore, the payment infrastructure operator who runs the validator nodes learns the identities of the payment recipient and sender, and the amount of each payment.

1 Sveriges Riksbank, E-Krona Pilot Phase 1, April 2021, https://www.riksbank.se/globalassets/media/rapporter/e-krona/2021/e-krona-pilot-phase-1.pdf.

Network performance. Digital cash solutions are easy to scale for high throughput. The payment validators need to check a signature and serial number for each deposited coin. Because each payment and coin deposit is essentially independent of each other, such operations can be easily run by independent payment validators in parallel. For example, each payment validator can be responsible for one range of possible serial numbers. This is in contrast to ledger-based solutions where typically all payment validators need to communicate and share a common view of all payments in the system, which makes scaling more complicated.

Communication costs in this model are low. The payment senders need to download a number of coins that depend on the total amount of payments that the sender wants to make.

Governance. This design is effectively centralized and, therefore, poses similar governance requirements to databases with account balances.

Privacy-preserving payment tokens

A privacy-preserving variant of the above token-based payment system was proposed by David Chaum.[37] As above, the payment sender withdraws coins from the currency issuer. The main difference is that the coin withdrawal process leverages a cryptographic technique called a blind signature. The user who withdraws the coins picks random serial numbers for each coin, and the use of blind signatures allows the currency issuer to sign the coins without learning their serial numbers.

To create a payment, the payment sender passes an appropriate number of coins to the payment recipient who forwards them to payment validators for double-spending checks and for updating the payment recipient's account balance. The main difference from plaintext tokens is that such a payment validation scheme preserves the privacy of the payment sender. The payment validators learn the payment amount and the identity of the payment recipient, but due to the use of blind signatures, the validators cannot link the deposit operation to a previous withdrawal operation and thus they cannot learn the identity of the payment sender.

37 David Chaum, "Blind Signatures for Untraceable Payments: Advances in Cryptology," proceedings of the Springer-Verlag Crypto '82 conference, 3, 1983, 199–203.

Analysis

The main differences between this currency variant and the previous one is the level of end-user privacy that is achieved, as well as the authentication process.

Authentication and authorization. Unlike in plaintext payment tokens, this design variant does not reveal the sender's identity to the validator(s). As such, there is not an explicit sender authentication process at the time of payment (only at the time of coin withdrawal). The identity of the payment recipient is authenticated at the time of payment so that the recipient's account balance can be updated accordingly. Payment authorization is similar to the plaintext setting (passing coins from the sender to the recipient).

Privacy. This currency variant provides privacy for the payment sender. The payment recipient can accept coins fully anonymously (i.e., without knowing the identity of the sender) and when the coins are deposited, the payment validators who may communicate with the currency issuer cannot link them to the identity of the sender either, due to the use of blind signatures during coin withdrawal. Private payment token systems do not provide privacy for the payment recipient. When the received coins are deposited, the recipient must authenticate their identity to the payment validators so that the validators can update the account balance of the recipient correctly. Such systems also do not ensure payment amount privacy, since the payment validators learn the denominations of the deposited coins, and thus the amount of the payment.

Signed balance updates

Next, we consider a hybrid payment approach proposed in recent research.[38] This approach combines centralized signing used in digital cash schemes with the account model and zero-knowledge proofs commonly used in private ledger transactions. We call this approach signed balance updates.

To join the system, each user creates a cryptographic commitment to a randomly chosen serial number and their current account balance value and requests the payment validators to sign this commitment. To create a new payment, both the payment sender and the payment recipient create new commitments to fresh serial numbers and the updated account balances that add the payment value to the recipient's balance and deduct the payment value from the sender's balance. The payment sender will also create a zero-knowledge proof that shows that both commitments are updated with the correct amount and the payment sender has sufficient funds in their current commitment. The payment recipient then sends the new commitments and the proof to the payment validators.

Similar to digital cash, the payment validators maintain a database of already used serial numbers. The validators will verify the proof and check that the serial numbers associated with the commitments have not already been used. If that is the case, the payment validators sign the new commitments (that represent balance updates) and return them to the payment sender and recipient who can consider the payment completed.

CASE STUDY:
Swiss National Bank (Switzerland)

In 2021, the Swiss National Bank (SNB) released a working paper that outlines one possible design for a CBDC system.[1] This working paper follows the private payment token approach with the use of cryptographic blind signatures during coin (token) withdrawal. To the best of our knowledge, there is no pilot project yet, but the working paper indicates that this currency variant is also being considered.

Advantages
Compared to the plaintext payment token scheme (used in the e-krona pilot), the main advantage is added privacy. More precisely, it is possible to perform payments where the identity of the payment sender remains private to the payment validators. For example, in a practical retail setting this would mean that the payment validators learn the payment amount and the identity of the merchant who accepts the payment, but not the identity of the customer who made the payment. Good scalability is another noteworthy advantage.

Risks
As discussed above, a token-based design places a higher burden on the user for wallet management, compared to ledger and database variants.

1 David Chaum, Christian Grothoff, and Thomas Moser, "How to Issue a Central Bank Digital Currency," SNB (Swiss National Bank) Working Papers, March 2021, https://www.snb.ch/en/mmr/papers/id/working_paper_2021_03.

38 Karl Wüst, Kari Kostiainen, and Srdjan Capkun, "Platypus: A Central Bank Digital Currency with Unlinkable Transactions and Privacy Preserving Regulation," Cryptology ePrint Archive, October 27, 2021, https://eprint.iacr.org/2021/1443.

Analysis

Integrity. The integrity of payments relies on two mechanisms. First, the payment sender creates a zero-knowledge proof that allows the payment validator to verify that the cryptographic commitments that represent account balance values are updated correctly. Assuming that the payment sender cannot forge such a proof, the integrity of each individual payment holds. Second, the payment validators check that each commitment serial number is used only once. This prevents double-spending the same funds multiple times in the system. So, there is no double-spending as long as the payment validator who approves the payment is not compromised. Here we assume that the used zero-knowledge scheme cannot be forged, and thus the only way to violate integrity is to compromise (a sufficient number of) payment validators (either remotely or locally through insider attacks).

Authentication and authorization. As with private distributed ledgers, there is no explicit authentication process at transaction time, as this design variant aims to break the link between users and their transactions. Payment authorization is based on zero-knowledge proofs. Each proof shows that the payment sender holds a private key (payment credential) that is associated with the used commitments. Payments cannot be created by unauthorized parties, as long as they cannot steal the payment credentials of legitimate users. As before, typical attack vectors for stealing user credentials would include stealing the user's device and tricking the user to install malicious software on their device.

Privacy. This approach provides sender privacy, recipient privacy, amount privacy, and payment unlinkability at the protocol level. The identities of the payment sender and recipient and the payment amount are hidden from the payment validators (and all other parties) because the used commitments hide all such details. Also, the used zero-knowledge proofs leak no information to the payment validators. Since fresh serial numbers are randomly chosen for each commitment, such payments also provide unlinkability. This means that payment validators, or another party, cannot connect one payment with another. (Linking of payments and construction of transaction graphs is a common technique used to de-anonymize ledger-based payments.) Network-level de-anonymization of users remains a potential privacy threat.

Resilience. To create a payment, both the payment sender and the payment recipient need to be online. The payment sender needs to communicate with the recipient and the validators. Due to the strong privacy protections provided by this design, fraudulent transactions cannot be easily reverted; in this regard, this design variant functions similar to cash.

Network performance. This approach provides good scalability. There can be several payment validators who are each in charge of separate ranges of commitment serial numbers and validate payments independently. A simple consistency check is needed between two validators (one who checks the sender commitment serial number and another who checks the recipient commitment serial number). The communication requirements of this scheme are moderate. Users upload commitments and proofs that they create, and download commitments signed by the payment validators. Users do not need to download the entire ledger that contains all transactions.

Governance. As before, the encryption in this design is primarily protecting the privacy of user transactions, not validator actions. As such, this design is effectively centralized and poses similar governance requirements to databases with account balances.

Secure hardware on clients

Finally, we consider a design alternative that assumes that every client has a trusted hardware module, such as a smart card or secure chip on a smartphone. This trusted hardware module maintains an account balance for the owner of the module. Payment is simple: the trusted hardware modules of the sender and the recipient execute a protocol where the payment amount is deducted from the balance in the sender's module and the same amount is added to the balance in the recipient's module.

Analysis

Integrity. The integrity of such a solution relies on the assumption that every hardware module used in the system remains uncompromised. If even only one of the users is able to break their own module (to which they naturally have physical access), such malicious users can double-spend the same funds an unlimited number of times. Also, if an external adversary is able to compromise even one of the deployed hardware modules, unlimited double-spending is possible. Another risk is a malicious hardware vendor or supply chain attack. If some of the deployed hardware modules are already malicious during the deployment phase, these design variants cannot guarantee the integrity of the currency. Due to such reasons, this variant is commonly seen as too risky for many deployments.

Authentication and authorization. User authentication can be conducted when transferring funds to the secure hardware; at transaction time, the hardware itself acts as an identifier. Similarly, simple payment authorization could be based on physical access to the trusted hardware module. That is, anyone who has the module can perform a payment. Such authorization would be vulnerable to module theft. Another approach is to require local user authentication for each payment. For example, the owner of the trusted hardware token provides a PIN code or fingerprint to the hardware module to authorize a payment.

Privacy. While this approach provides weak integrity guarantees, it offers strong privacy protections. Because payments happen directly between the sender and the recipient, there is no information leakage to validators or any other parties. Thus, such payments are fully anonymous and unlinkable (and leave no electronic trace to any payment infrastructure). Therefore, this design variant provides similar privacy guarantees as cash payments.

Resilience. This design variant supports offline payments. That is, payments are possible between the sender and the recipient even if both parties are offline, as long as they can communicate with each other (e.g., using a local communication channel such as near-field communication; NFC). Performing safe offline payments without trusted hardware is currently an open problem, and thus no other design variant discussed in this chapter provides similar offline-payment capability. In this design variant, fraudulent transactions cannot be easily reverted (similar to cash).

Network performance. Such design is extremely scalable, as there is no centralized authority like the payment validators who would need to approve each payment. Such payments would need only minimal communication between the payment sender and the recipient.

Governance. The use of secure hardware introduces new challenges related to the responsibilities of hardware manufacturers. For example, policies must be put in place for managing the implications of possible security vulnerabilities (intentional or otherwise) in trusted hardware modules.

Summary

While the design space is large, many central banks have narrowed their scope to three of the discussed design variants: databases with balances, distributed ledgers with plaintext transactions, and variants of digital cash. Although there are no central banks that have committed to the other three design choices (to the best of our knowledge), there could be hybrid architectures that allow for combinations of technologies.

Figure 4. Breakdown of Current Adoption/Exploration of Different CBDC Design Variants Globally

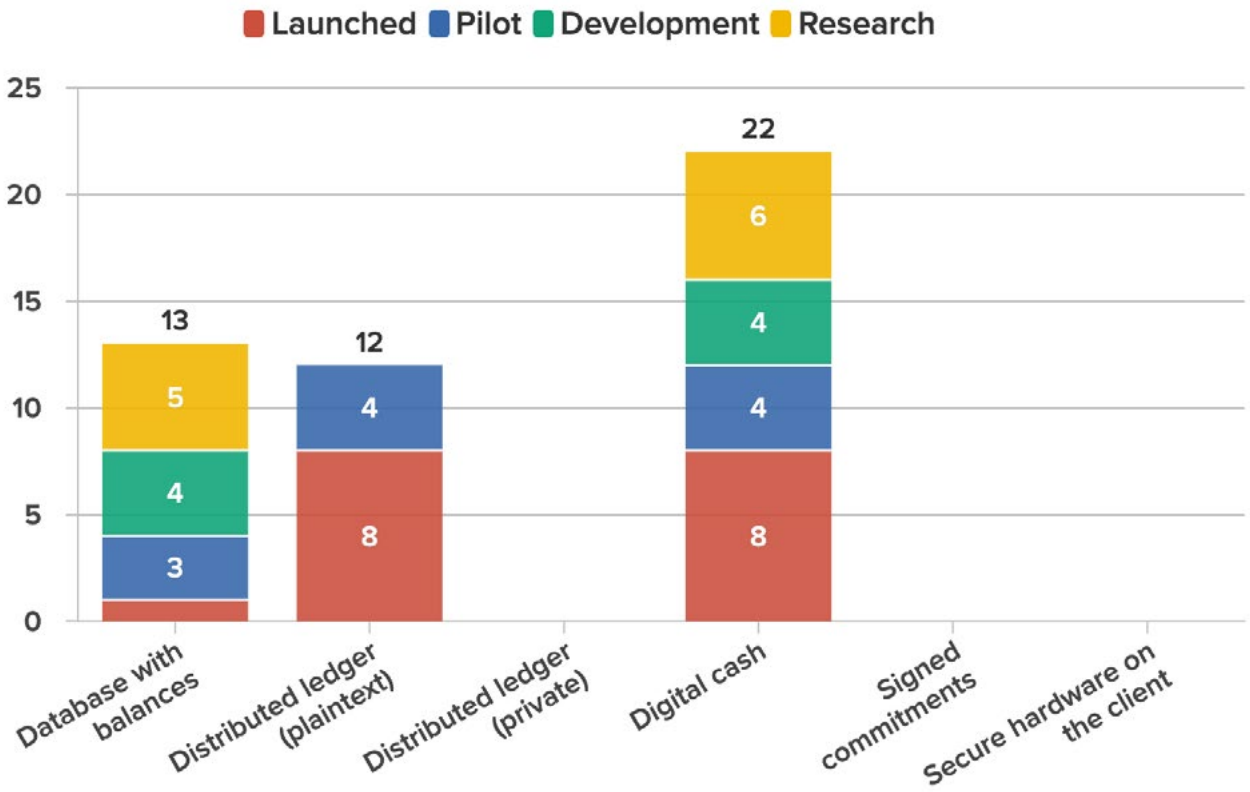

Source: Data were taken from the Atlantic Council's CBDC tracker at https://www.atlanticcouncil.org/cbdctracker/, accessed on June 10, 2022.

The table below summarizes our analysis in this chapter. Due to space limitations, governance considerations are not included in the table; a discussion of governance can be found in the main text.

For a discussion on differences in governance models, see the discussion of individual design variants in this chapter.

Table 2. Summary of Currency Variant Analysis

	Integrity	Authorization	Privacy	Resilience	Performance
Database with account balances	Relies on a centralized validator	Possession of payment key	Toward validators: none. Toward third parties: communication leakage.	Payment sender needs to be online. Infrastructure needs to be operational. Transaction can be reverted.	Low communication requirements. for all parties. High throughput Low latency.
Distributed ledger with plaintext transaction	Relies on a set of validators collectively	Possession of payment key	Toward validator: none. Toward third parties: communication leakage.	Payment sender needs to be online. Infrastructure needs to be operational. Transaction can be reverted.	Validators need to run consensus protocol. Low communication cost for users. Moderate throughput and latency.
Distributed ledger with private transactions	Relies on a set of validators collectively	Possession of payment key	Toward validators: full. Toward third parties: communication leakage.	Payment sender needs to be online. Infrastructure needs to be operational. Transaction cannot be reverted.	Validators need to run consensus protocol. User may need to download entire ledger. Moderate throughput and latency.
Plaintext payment tokens	Relies on a set of validators	Possession of coins	Toward validators: none. Toward third parties: full.	Payment recipient needs to be online. Infrastructure needs to be operational. Transaction can be reverted.	Validators do not need consensus. Users need to download and upload coins. High throughput Low latency.
Privacy-preserving payment tokens	Relies on a set of validators	Possession of coins	Toward validators: sender privacy. Toward third parties: full.	Payment recipient needs to be online. Infrastructure needs to be operational. Transaction can be reverted.	Validators do not need consensus. Users need to download and upload coins. High throughput Low latency.
Signed balance updates	Relies on a set of validators	Possession of payment key	Toward validators: full. Toward third parties: communication leakage.	Payment recipient needs to be online. Infrastructure needs to be operational. Transaction cannot be reverted.	Validators do not need consensus. Users download and upload account commitments. High throughput Low latency.
Secure hardware on clients	Relies on correctness of hardware modules	Possession of hardware module	Toward validators: n/a. Toward third parties: full.	Offline payments are possible. Transactions cannot be reverted.	Only local communication between users. Very high throughput. Low latency.

Note: Text highlighted in green represents a well-supported requirement or an advantage of the analyzed currency variant. Text highlighted in orange represents a requirement that is not well supported or an aspect of the currency variant that is a disadvantage compared to other variants.

ADDITIONAL KEY DESIGN CHOICES

The previous section described possible design alternatives for a CBDC and their analysis. In this section, we discuss other common design choices that (possibly) span different designs, including consensus, wallets, and privacy together with compliance.

Consensus mechanism

System designers must choose a consensus mechanism, which determines how transactions are confirmed by the validator node(s). The choice of consensus mechanism requires understanding trade-offs between robustness and efficiency. At one extreme, we have a single validator to confirm the validity of each transaction. This is efficient because it requires no coordination between multiple validators, but it is not robust; if the validator goes offline or misbehaves, integrity and/or availability are lost. At the other extreme, we can design consensus schemes with hundreds or thousands of validators, as in public cryptocurrencies like Bitcoin or Ethereum. Such approaches tend to be much less efficient, as they are fundamentally limited by the bandwidth and latency of the underlying network. However, these mechanisms tend to be much more robust to misbehaving or unavailable validator nodes. In practice, we expect CBDCs are likely to operate in an intermediate regime, with, for example, tens of validators.

Fault models. In this regime, two types of robust consensus mechanisms are typically considered: crash fault-tolerance and Byzantine fault-tolerance. Crash fault-tolerance means that the protocol is robust to some fraction of validators going offline, for example, due to a disruption in power or network infrastructure. Byzantine fault-tolerance is a stronger concept; in addition to tolerating crash faults, it is additionally robust to a fraction of validators actively misbehaving, for example, by deviating arbitrarily from protocol. Byzantine fault-tolerance requires additional communication costs compared to simple crash fault-tolerance; this accordingly increases latency and can reduce throughput. However, in a CBDC, the financial incentives for misbehavior are high; there is a compelling case to be made for building in robustness to Byzantine faults.

When evaluating consensus mechanisms, it is essential to consider the precise security assumptions and guarantees of each mechanism and ensure that back-end infrastructure is designed to match those assumptions. For example, many Byzantine fault-tolerant consensus protocols are robust up to some fraction of malicious parties (e.g., one-third or half). This means that (for example) up to one-third of the validators can be compromised without affecting the system's integrity. These attractive security guarantees have led some countries to consider adopting such consensus mechanisms (e.g., the digital euro).

TAKEAWAY

USE OF PROVEN PROTOCOLS IS IMPORTANT

Byzantine fault-tolerant consensus mechanisms (e.g., DLT) are notoriously difficult to design and implement securely. Consensus protocols should be carefully evaluated (e.g., through peer review), and run on fully independent infrastructure to give meaningful security guarantees

Deployment considerations. Despite the apparent security benefits, Byzantine fault-tolerant distributed validation protocols should be implemented with care. For the security guarantees to be meaningful, it is essential that validators be independent. That is, the corruption of one validator should minimally (or not at all) affect the likelihood of another validator being corrupted. At a minimum, this means that validator nodes should be run on servers running from different locations and using different power sources and network infrastructure. Ideally, they should be hosted and managed by independent entities. This is meant to avoid situations where, for example, an adversary manages to compromise the integrity of a single validator, and then uses the same exploit to compromise the remaining validators. In such settings, the security guarantees provided by Byzantine fault-tolerant consensus would be vacuous.

Another important consideration is the ability to identify misbehaving nodes in a consensus protocol. That is, suppose some fraction greater than half of validators misbehave. In such scenarios, it is important to be able to identify which nodes misbehaved to punish them appropriately. However, some consensus protocols make such reidentification difficult (e.g., PBFT-MAC), whereas others naturally support it (e.g., LibraBFT).[39] Such questions of consensus protocol forensics are an important consideration when selecting a consensus mechanism.

Consensus and fairness. The choice of consensus mechanism can also have implications for the fairness of the CBDC. For example, some consensus mechanisms choose a single validator node to be the "leader" at each instant; the leader's job is to order incoming transactions and commit them to the ledger. However, such leader-based protocols can undermine fair transaction ordering; the leader can be bribed to place some transactions before others, leading to the risk of financial manipulation. It is an active research area today to identify

39 Peiyao Sheng et al., "BFT Protocol Forensics," CCS '21: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, November 2021, 1722–1743, https://doi.org/10.1145/3460120.3484566.

(efficient) consensus protocols that preserve the natural ordering of transactions in the presence of malicious validators.[40]

Wallets

In most currency variants reviewed above, the payment sender authorizes the payment by signing a transaction with their payment credential, such as a digital signature key, obtained from an account provider (typically a commercial bank). The secure storage and use of the payment credential are important — if unauthorized parties learn the payment credential, they can spend the user's funds; or if the legitimate user loses the payment credential, they are no longer able to spend their funds. The data storage and computing environment where the credential is stored and used is commonly called a digital wallet.

Custodial wallets. One possible deployment alternative is one where the account providers (commercial banks) support the users in the management of their payment credentials. Such deployments are commonly called custodial wallets or custodial solutions. For example, a bank can create the payment credential and provision it to the user's digital wallet on their smartphone. If the user loses the device, and thus the payment credential, the bank can issue a new payment credential and send it to the user (upon successful authentication).

Another custodial alternative is that the payment credential is stored only at the bank. In such a solution, payment creation requires contacting the bank with user authentication. One benefit of this approach is that if the user's device is lost or stolen, the payment credential is not leaked. Another benefit is that if the user has multiple devices, the payment credential can be conveniently accessed from any of the other devices without the user having to replicate or synchronize credentials across multiple devices. The main drawback of custodial solutions is that a malicious insider at the bank is able to use the user's payment credential without permission. Also, if the relevant IT system of the bank is compromised, or if the bank is subject to a data leak, a large number of payment credentials may be leaked.

Non-custodial wallets. The alternative approach is a non-custodial wallet, where the user maintains the payment credential themself. The payment credential can be created and stored on the user's smartphone that hosts the wallet software. A simple wallet software could store the payment credential on the normal data storage like flash memory. The main benefit of non-custodial wallets is that the payment credential is not directly accessible to any other party besides the owner of the funds. The downside of non-custodial wallets is that the user needs to manage backups themself. Safe backups can be difficult to organize in practice (paper backups may get lost, online backups are not safe, and many users might forget to create a backup altogether).

Secure hardware. Because smartphones can get lost, stolen, or infected with malware, a more secure approach is to store the payment credential inside a protected environment like hardware-assisted Trusted Execution Environment (TEE). Most modern smartphones support a TEE technology called ARM TrustZone, while PC platforms like laptops typically support TEE technology called Intel SGX. Both TEE technologies allow storage and use of the payment credential such that the credential is not accessible by any other software except the wallet software on the same device. While TEE wallets increase the robustness of credential storage significantly, recent research has shown that TEEs can be vulnerable to sophisticated attacks like side-channel analysis.

Another option is to store the payment credential in a separate hardware token, such as a USB dongle. Hardware tokens offer strong security guarantees because the payment credential is physically isolated from potentially unsafe devices like the user's smartphone or laptop. A common challenge with hardware tokens is how to safely back up a payment credential. Another typical challenge is the limited user interface on small hardware dongles, and thus safe payment detail input or verification can be difficult with hardware tokens.

In all wallet solutions, safe storage of the payment credential relies on the trustworthiness of the hardware that hosts the wallet software. Nation-state adversaries could coerce hardware manufacturers to implement backdoors that would allow the adversary to learn any secrets stored on that hardware. While such attacks are only possible from the most powerful adversaries, such threats should be considered as part of an extensive threat profile for CBDCs.

Social engineering attacks. Social engineering attacks are currently one of the most widely used and successful attack vectors in IT systems. In traditional email phishing, the victim receives a benign-looking but malicious email from the adversary. The goal of the email is to convince the victim to enter their login credentials into a fake website controlled by the adversary.

Phishing attacks are becoming increasingly common also in the context of decentralized cryptocurrencies. A possible attack vector tricks the victim into revealing the "recovery seed" of their wallet.[41] If the adversary obtains such information, they can recreate the victim's payment credentials and steal all of the victim's funds (most hardware wallets support a recovery

40 Mahimna Kelkar et al., *Order-Fairness for Byzantine Consensus*, August 9, 2020, https://eprint.iacr.org/2020/269.pdf.

41 M. Moon, "Crypto Scammers Stole $500K from Wallets Using Targeted Google Ads," Engadget, November 4, 2021, https://www.engadget.com/crypto-scammers-google-ads-phishing-campaign-100044007.html.

seed so that the legitimate owner of the wallet can recover funds in case the hardware token is lost or damaged). Another possible attack is to trick the user into performing a fraudulent transaction that transfers some of their cryptocurrency assets to the adversary.[42] Such spoofing attacks work even if the adversary does not obtain the user's payment credentials—the adversary merely tricks the victim into using their credentials to the benefit of the adversary. As payments cannot be easily reverted in current decentralized cryptocurrencies, such attacks are difficult to recover from.

Similar adversarial strategies could apply to future CBDC deployments. The fact that CBDC will be more centralized may alleviate some concerns (e.g., the possibility to revert fraudulent transactions can be built into certain currency designs), but in general, the same attack concepts apply to both decentralized and centralized systems. While certain countermeasures and defensive techniques are well-known (e.g., multi-factor authorization, safe wallet UI design practices), such attacks most likely cannot be fully eliminated by technical means alone. As in any large and complex IT system, humans and social engineering remain a viable threat vector, and security awareness and user training are probably required to limit the effectiveness of these threats. How exactly such attacks may manifest in future CBDCs will depend on how such systems will be implemented, what kinds of wallet user interfaces will become common, and various other similar factors. Therefore, performing a detailed analysis on this topic is not yet possible, but designers of CBDC systems are, nonetheless, advised to consider such adversarial strategies.

Smart contract design and deployment

Many CBDC deployments (especially of the retail variety) are expected to support applications that allow users to interact with the underlying CBDC infrastructure. Examples include payment or banking applications. These applications will likely be backed by smart contracts, which are computer programs that govern the transfer of digital assets. First proposed in the context of decentralized cryptocurrencies, the concept is much more general and can be applied to centralized financial services as well. To give an example, a smart contract could be used to programmatically implement disaster relief programs by specifying that every registered citizen should receive $500 at a particular date and time. Smart contracts can also specify conditions under which transfers should occur; for example, a smart contract could specify that every time a user (Alice) receives a payment from another user (Bob), 30 percent of that payment will be transferred to Alice's family member (Carol). Finally, smart contracts can be composed with each other to create complex dependencies between events in a financial system.

Smart contracts have been particularly powerful in the cryptocurrency space because of standardization: although different parties have differing goals and requirements and can use different programming languages, all parties utilize the same set of rules for specifying and processing smart contracts. The most widely known set of such rules is the Ethereum Virtual Machine (EVM). The EVM has enabled the deployment of complicated applications between parties that would otherwise require either manual effort or custom-built services governing the logic and automated transfer of funds.

Smart contracts raise a number of issues that are likely to pose new cybersecurity challenges for CBDCs.

Managing vulnerabilities or errors. Because smart contracts are computer programs, it is inevitable that many smart contracts will have bugs: errors in the logic or implementation of the contract. Software bugs can lead to (sometimes catastrophic) security vulnerabilities. The main concern in a CBDC is that these errors could erroneously transfer large amounts of money to the wrong recipient, or enable malicious agents to steal money by exploiting vulnerabilities in a smart contract. While software vulnerabilities have always been a concern for financial institutions, the main risk that arises with smart contracts is greater scale: smart contracts enable large-scale, nearly instantaneous transactions, which can also set off a chain of downstream-dependent transactions. In decentralized systems (e.g., cryptocurrencies), smart contract bugs have been particularly problematic because contracts are immutable, meaning they cannot be changed once they are deployed. In a more centralized CBDC setting, contracts do not necessarily have to be immutable. However, as mentioned earlier, some CBDC design variants make it more difficult to revert transactions (Figure 3); these designs may complicate full recovery from bugs.

To some extent, these vulnerabilities can be managed through a combination of technical and procedural means. On the technical side, software engineering best practices call for testing all code before deployment. In other words, smart contracts should be evaluated under a wide range of inputs to evaluate whether they contain vulnerabilities prior to deployment. While there are tools for software (and smart contract) testing, these are not error proof. The problem is particularly complicated for smart contracts because of the complex dependencies between them. An input to one contract may not obviously cause errors but could have cascading effects that cause errors in another contract that is invoked through a chain of downstream dependencies.

42 Charlie Osborne, "Microsoft Warns of Emerging 'Ice Phishing' Threat on Blockchain, DeFi Networks," ZDNet, February 17, 2022, https://www.zdnet.com/article/microsoft-warns-of-ice-phishing-on-blockchain-networks/.

On the procedural side, CBDCs could implement staging or test environments where smart contracts are deployed and evaluated at a small scale before being fully deployed on the main CBDC network. This is again analogous to best practices in software engineering for deploying new updates to complex software systems.

Privacy

Any deployed CBDC system is likely to collect significant amounts of data since such systems would process a large number of payments every day. Some central banks have indicated that their goal is not to build a tool of mass surveillance[43] and, therefore, CBDC deployments should carefully consider and incorporate at least some privacy protections.

System designers have a number of mechanisms for preserving privacy in a CBDC. Process and policy can be an important tool for enforcing privacy with respect to system insiders. Here, it is important to follow the principle of "least privilege": operators should be given access only to the data they require to do their jobs. For example, an account operator should not have access to portions of the ledger that are not relevant to its own customers. Even when access control policies are in place, insiders within the currency issuer (or account issuer) can still have access to large quantities of sensitive financial data. Digital currency variants with built-in privacy protections, as described and analyzed in the previous section, provide a significantly stronger foundation for user privacy, as in such design the privacy of end users does not rely on the trustworthiness of system insiders.

Privacy and compliance. At the same time, most countries have rules and laws like AML regulations that would need to be appropriately enforced. For example, in the United States, it is mandatory to report the receipt of more than $10,000 in cash payments to the IRS. Obviously, policy makers do not want a CBDC system to become widely used for illicit activities or to create materially new problems for the enforcement of criminal law (more than what exists with cash). Another example of concern is that if holding large amounts of CBDC money is made safe and easy, users might be tempted to migrate their savings from commercial banks to a CBDC format.[44] There are some concerns that this could threaten the safe operation of commercial banks (e.g., increase the possibility of bank runs during financial crises) and, under certain conditions, the stability of the monetary system.

For various reasons, it may be desirable to create a system where users enjoy some measure of privacy, but at the same time authorities are still able to enforce laws such as how much CBDC money can be spent, received, or held. These two requirements are to some extent in tension since most currency variants are able to provide only one but not the other. For example, a database that holds account balances, or a ledger that records plaintext transactions, is easy to regulate but provides no end-user privacy. A ledger that records private transactions, similar to systems like Zcash or Monero, provides privacy but is hard to regulate because payment details like identities and amounts are not disclosed even to infrastructure nodes like validators who process the payments.

One of the most promising approaches to provide both privacy and compliance is to use cryptographic zero-knowledge proofs to construct payments that preserve user privacy (as much as possible) but can be verified to conform to specific regulatory rules. One example is a solution where the amount of each payment is hidden (e.g., encrypted) and each transaction must be accompanied with a zero-knowledge proof that shows to the regulator that all the payments received by the same user within the current time period (e.g., one month) are combined below a certain allowed limit (like $10,000).[45]

Another example is a solution where the zero-knowledge proof shows that the updated account balance of the recipient is below a certain limit (say, $50,000) without revealing the exact account balance to the payment validators or the regulator. The first technique could mimic the current rules regarding the

TAKEAWAY

PRIVACY AND COMPLIANCE CAN COEXIST

Providing users with strong privacy protections and regulators with the extensive oversight they may desire are two inherently conflicting requirements. However, recent research developments have shown that it is possible to design digital currencies where these two requirements may coexist, at least to some extent. For example, it is possible to realize a digital currency where payment details remain fully private as long as the total value of all payments by the same individual does not exceed a certain predefined threshold value (say, $10,000 per month). In such a system, fully private payments are allowed up to a certain monthly limit and if the individual exceeds that limit, the regulatory authority is able to see the details of payment transactions. As discussed further below, privacy issues must be squarely addressed at the legislative level.

43 Bank of England, "Central Bank Digital Currency."

44 Stan Higgins, "Central Bank Digital Currencies Could Fuel Bank Runs, BIS Says," CoinDesk, updated September, 13, 2021, https://www.coindesk.com/markets/2018/03/12/central-bank-digital-currencies-could-fuel-bank-runs-bis-says/.

45 Karl Wüst et al., "PRCash: Fast, Private and Regulated Transactions for Digital Currencies," https://fc19.ifca.ai/preproceedings/5-preproceedings.pdf.

reporting obligation for large cash payments, while the second technique could be used to address excessive migration of bank deposits to protect the stability of banks.

In addition to zero-knowledge proofs, other privacy-preserving techniques have also been studied and proposed in the research literature. Fully homomorphic encryption and private set intersection are two examples. Such techniques are not used or required in the design variants that are the focus of this chapter but may enable new privacy-preserving currency designs in the future.

Privacy and network traces. Third-party adversaries with access to the network layer (e.g., ISPs) or compute layer (e.g., cloud service providers) can potentially de-anonymize transactions.[46] This threat can (and should) be mitigated in part by encrypting all traffic between validators and end users. This is not possible in permissionless cryptocurrencies, where all transactions are meant to be publicly broadcast. In a CBDC, though, there is no reason for third parties to have access to transaction packet contents.

Privacy and performance. Cryptographic privacy protections can have important implications for other security and efficiency properties. For example, zero-knowledge proofs increase the computational overhead of creating and validating transactions. This overhead can impact latency and, if implemented poorly, throughput (e.g., if the system requires interactive zero-knowledge proofs). More generally, the use of cryptography limits the kinds of operations that can be performed by validators on encrypted transactions. Thus, the system needs to be designed much more carefully to anticipate the kinds of computation that may be necessary down the line, for example, related to regulatory compliance.

Summary. Cryptography-based privacy solutions like zero-knowledge proofs for AML/KYC compliance are still an active area of research. The performance implications of initial research proposals need further validation, and more sophisticated solutions are likely to be proposed in the near future. However, fortunately, the initial results indicate that reconciling regulation and privacy is not a fully impossible task and central banks can consider such solutions that ensure both as part of their technology road map.

Cybersecurity frameworks

Over the last decades, several best practices and expert recommendations regarding how to build and deploy secure IT systems have been collected in various cybersecurity frameworks and standards. The ISO 27000 series and the National Institute of Standards and Technology's Cybersecurity Framework (NIST CSF) are two popular examples.

Regarding the design and build phases of IT systems, such cybersecurity frameworks may, for example, mandate how the system design process should be documented or what kind of testing methods should be used. The frameworks can also provide security-related checklists that system designers and programmers may follow. Regarding the operation of a deployed IT system, such frameworks may provide organizational advice, such as who should be allowed to access confidential user data or how the organization should respond to possible security incidents. Many experts agree that leveraging such frameworks can be useful (if the added cost is acceptable for the project at hand).

For the creation of a CBDC, cybersecurity frameworks can provide similar benefits (and costs) as for other IT systems. Careful application of a chosen cybersecurity standard can, for example, help to ensure that the design process is documented appropriately, the software testing phase is performed based on industry best practices, and appropriate measures are in place to respond to possible security incidents.

However, the existing cybersecurity frameworks and standards do not provide advice regarding some of the most challenging and fundamental design choices related to the creation of a CBDC. As we have discussed earlier in this chapter, each digital currency variant provides a different security, privacy, and performance trade-off and comes with its unique set of risks and challenges. The currently available cybersecurity frameworks do not explicitly help system designers make critical choices such as which digital currency variant to choose. Therefore, the designers of future CBDC systems may need to consult a broader set of resources (such as the analysis presented previously in this chapter) during the design process.

CHAPTER 1 SUMMARY

In this chapter, we have analyzed the cybersecurity aspects of CBDCs. Our discussion first identified the main roles and entities involved in a CBDC deployment. After that, we discussed possible threat models and the key security requirements. Using such a framework, we then analyzed various possible digital currency design alternatives and compared their main advantages, drawbacks, and cybersecurity challenges. The main takeaways of this chapter are as follows.

46 Alex Biryukov, Dmitry Khovratovich, and Ivan Pustogarov, "Deanonymisation of Clients in Bitcoin P2P Network," CCS '14: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, November 2014, 15–29, https://doi.org/10.1145/2660267.2660379.

TAKEAWAY

THE MAIN TAKEAWAYS OF THIS CHAPTER

- **CBDC deployment may introduce new cybersecurity risks.** While a CBDC would be subject to many of the same cybersecurity risks as the existing financial systems, deployment of a CBDC would also create new cybersecurity risks, such as increased centralization, reduced regulatory oversight, increased difficulty of reversing fraudulent transactions, challenges in payment credential management, malicious transactions enabled by automated financial applications, and increased reliance on non-bank third parties. The exact set of risks depends largely on the design and deployment of a given CBDC.

- **The design space for CBDCs is large.** While most CBDC reports identify centralized databases, distributed ledgers, and token models as possible digital currency designs, our discussion shows that the design space for digital currency systems is actually larger than that. Currencies can also be realized as signed balance updates or as a set of trusted hardware modules. Both ledger and token-based payments can be made private through cryptographic protections.

- **CBDC deployment might centralize user data collection.** The main difference between a (centralized) CBDC deployment and the current financial system is that the CBDC may result in a greater centralization of user data and financial infrastructure. This can have advantages, such as new options for implementing monetary policy, but it can also have serious privacy and security disadvantages.

- **Privacy-preserving designs can also be more secure.** If a CBDC deployment without privacy protections gets breached either by an external attacker or malicious insider, then large amounts of sensitive user information are disclosed to unauthorized parties. In a privacy-preserving CBDC deployment that initially declines to collect or subsequently restricts sensitive user data even from trusted system insiders, breaches will have significantly less severe security consequences.

- **Strong user privacy protection is possible.** While some recent reports imply that CBDCs would inherently reduce the privacy of users, our review of recent research developments has shown that it is possible to design a digital currency system where transaction details are hidden even from the payment validators and infrastructure. We argue that such systems would provide a level of privacy that is comparable to cash.

- **Privacy and compliance can coexist.** User privacy protection and enforcement of compliance rules are at odds with each other, and simple system designs can typically achieve only one or the other. Our review of recent research advancements indicates that it is possible to design systems where users enjoy reasonable levels of payment privacy and regulatory authorities can at the same time enforce common compliance rules.

- **The use of proven protocols is important.** Distributed security protocols, such as Byzantine fault-tolerant consensus protocols, are notoriously difficult to design securely. Our discussion shows that several current CBDC pilot projects rely on consensus protocols that lack strong, peer-reviewed security proofs. We discourage the use of such potentially unsafe protocols as key components of CBDC deployments.

Chapter 2: Policy Recommendations—Principles for Future Legislation and Regulation

At this early stage of CBDC research and development, the precise nature of the cybersecurity risks presented by CBDCs will depend significantly on the design and implementation decisions made by governments, legislatures, and central banks around the world. In the US context, we do not have a concrete decision on a CBDC design, let alone a definitive prototype, set of corresponding public policies, or authorizing legislation. That makes it somewhat more challenging to make detailed recommendations for strengthening cybersecurity at this early juncture.[47] This chapter identifies key principles to help guide policy makers and regulators as they continue to explore and potentially deploy a CBDC with robust cybersecurity protections in mind.

PRINCIPLE 1: WHERE POSSIBLE, USE EXISTING RISK MANAGEMENT FRAMEWORKS AND REGULATIONS

Cybersecurity policy around CBDCs need not entirely reinvent the wheel. There are already a variety of laws, safeguards, and requirements in place to protect the traditional banking sector and consumers from cyberattacks, some of which might directly apply (in the case of a CBDC administered by a nationally chartered bank) or which might serve as a useful model for future adaptation.

For example, in the United States, a combination of bank and non-bank regulators, federal statutes, state laws, and private sector standards shape cybersecurity in the traditional financial services sector.[48] These include the Gramm-Leach-Bliley Act of 1999 (on data privacy and security practices), the Sarbanes-Oxley Act of 2002 (reporting requirements), the Fair and Accurate Credit Transactions Act of 2003 (regarding identity theft guidelines), and the Bank Service Company Act of 1962 (regarding onsite examinations and proactive reporting of cybersecurity incidents).[49] The Federal Deposit Insurance Corporation (FDIC) alone offers detailed guidance and resources on cyber risks and examinations for banks.[50]

Depending on how a CBDC was designed and deployed, some of these laws might apply directly or indirectly. For example, particularly to the extent a two-tier CBDC would be administered or held by banks, regulators would likely have to carefully review compliance with existing security frameworks and standards. Likewise, to the extent that a CBDC is administered or held by a fintech company—such as a mobile payments app, neobank, or hot wallet—then a number of existing laws would probably apply.[51] In some instances, it will be prudent to streamline or deconflict preexisting regulations that overlap and apply to CBDCs in needlessly complex ways.

As a first step, policy makers and regulators should assess which areas of a new CBDC ecosystem will be covered by current regulations and where novel statutes—or new technical frameworks—might be necessary to provide adequate protection. Examples of existing cybersecurity frameworks include the NIST CSF, which "provides a comprehensive framework for critical infrastructure owners and operators to manage cybersecurity risks,"[52] and the Committee on Payments and Market Infrastructures and the Board of the International Organization of Securities Commissions' Guidance on Cyber Resilience for Financial Market Infrastructures.[53] The G7's "Fundamental Elements of Cybersecurity for the Financial Sector" offers

47 We will closely watch the Federal Reserve Bank of Boston and MIT's CBDC project for code samples.

48 M. Maureen Murphy and Andrew P. Scott, Financial Services and Cybersecurity: The Federal Role, US Library of Congress, Congressional Research Service, R44429, updated March 23, 2016, https://crsreports.congress.gov/product/pdf/R/R44429. See also Jeff Kosseff, "New York's Financial Cybersecurity Regulation: Tough, Fair, and a National Model," *Georgetown Law Technology Review*, April 2017, https://georgetownlawtechreview.org/new-yorks-financial-cybersecurity-regulation-tough-fair-and-a-national-model/GLTR-04-2017/.

49 Andrew P. Scott and Paul Tierno, "Introduction to Financial Services: Financial Cybersecurity," US Library of Congress, Congressional Research Service, IF11717, updated January 13, 2022, https://sgp.fas.org/crs/misc/IF11717.pdf.

50 "Banker Resource Center, Information Technology (IT) and Cybersecurity," Federal Deposit Insurance Corporation, accessed February 15, 2022, https://www.fdic.gov/resources/bankers/information-technology/.

51 See generally Chris Brummer, *Fintech Law in a Nutshell* (St. Paul, Minnesota: West Academic), 461–538. Brummer summarizes the cybersecurity regulations that apply to fintech services, for example, the Cybersecurity Act of 2015, the Gramm-Leach-Bliley Act, and other rules.

52 Tarik Hansen and Katya Delak, "Security Considerations for a Central Bank Digital Currency," FEDS Notes, Board of Governors of the Federal Reserve System, February 3, 2022, https://doi.org/10.17016/2380-7172.2970.

53 Committee on Payments and Market Infrastructures and the Board of the International Organization of Securities Commissions, *Guidance on Cyber Resilience for Financial Market Infrastructures*, June 2016, https://www.bis.org/cpmi/publ/d146.pdf.

policy makers another measuring stick to compare a CBDC's necessary regulations against.[54] Chapter 1 of this report details the benefits of using these frameworks, but stresses that none of the current schemes fully address the most challenging and fundamental choices related to designing a secure and resilient CBDC. Therefore, we encourage policy makers to begin collaborating with industry associations and leveraging international fora to update current frameworks using resources such as this report. The European Union (EU) provides a useful example as its current banking and stability provisions will cover certain aspects of new FinTech innovations.[55]

Regulators may have to balance old and new regulation, as well as weigh potentially competing policy values, such as security, innovation, competition, and speed of deployment. When crafting new regulations for a CBDC, policy makers and regulators should set the conditions for a safe digital currency ecosystem that enables financial intermediaries to innovate and compete.[56] For a two-tier retail CBDC system, which according to the Atlantic Council's CBDC Tracker is the most popular architecture choice,[57] regulators will have to devise rules for private payment service providers (PSPs) that extend beyond commercial banks to cover activities by nontraditional financial firms involved in operating the CBDC. Alipay and WeChat's important role as technology providers in the rollout of China's e-CNY underscores this point.[58] A CBDC's design choices will determine where policy makers and regulators need to step in to provide new frameworks that protect participants from cyber risks. For example, as explained in Chapter 1, a CBDC with token-based wallets would both place a higher burden on consumers to keep their money safe and require policy makers to develop "a regulatory framework for custodial wallets with the necessary consumer and insolvency protections."[59] Related to wallets' vulnerabilities, policy makers should consider putting in place consumer protections for data custody, including rules on the storage redundancy (and data retention limits) of transaction records and wallet balances. Doing so could insulate consumers and banks from the long-term impacts of breaches, technical failures, and fraud, enabling more rapid recovery and response from such incidents.

Given that certain CBDC designs might put a potentially higher burden on consumers to protect themselves against cyber fraud and theft, governments should engage PSPs and consumer protection groups to roll out cyber risk education campaigns well before launching a CBDC. As discussed in the background chapter, the credit card industry offers a cautionary tale of phishing and other cyber scams' severe costs for consumers and the industry. A successful educational campaign would raise awareness among CBDC users about how to identify and protect themselves against a wide variety of cyberattacks. In addition to learning appropriate cyber hygiene when using wallets and other CBDC applications, consumers must be informed of their legal rights and responsibilities that come with holding and transacting in digital currency.[60] At the same time, a CBDC must not offload all (or most) of the responsibility for cybersecurity onto its users.

China's official app for digital yuan is seen on a mobile phone next to 100-yuan banknotes in this illustration picture taken October 16, 2020. Source: REUTERS/Florence Lo/Illustration

PRINCIPLE 2: PRIVACY CAN STRENGTHEN SECURITY

One of Chapter 1's key findings is that privacy-preserving CBDC designs may also be more secure because they reduce the risk and potential harmful consequences of cyberattacks associated with data exfiltration, for example. CBDCs with stronger privacy rules may generate and store less sensitive data in the first place. In turn, potential attackers have a smaller incentive to infiltrate the system. If an attack is successful, the impact would be less severe. Our research

54 G7 (Group of Seven), "Fundamental Elements of Cybersecurity for the Financial Sector," October 2016, https://www.ecb.europa.eu/paym/pol/shared/pdf/G7_Fundamental_Elements_Oct_2016.pdf.

55 See, for example, Juan Carlos Crisanto and Jermy Prenio, *Regulatory Approaches to Enhance Banks' Cyber-Security Frameworks*, FSI Insights on policy implementation No. 2, Financial Stability Institute, August 2017, https://www.bis.org/fsi/publ/insights2.pdf.

56 "Digital Currency Consumer Protection Risk Mapping," Digital Currency Governance Consortium White Paper Series, World Economic Forum, November 2021, 17, https://www3.weforum.org/docs/WEF_Digital_Currency_Consumer_Protection_2021.pdf.

57 "Central Bank Digital Currency Tracker."

58 Arjun Kharpal, "China's Digital Currency Comes to Its Biggest Messaging App WeChat, Which Has over a Billion Users," CNBC, January 6, 2022, https://www.cnbc.com/2022/01/06/chinas-digital-currency-comes-to-tencents-wechat-in-expansion-push.html.

59 "Digital Currency Consumer Protection Risk Mapping," 18.

60 Ibid., 17.

also shows that CBDCs can offer cash-like privacy while potentially providing more efficient oversight options to regulatory authorities. To build a CBDC, policy makers in the US Congress and their colleagues around the world should carefully examine the relationship between privacy and security. They should weigh the findings of this report before making foundational decisions about a CBDC's level of privacy that will filter through to the digital currency's design and determine its cybersecurity profile.

As part of the privacy question, policy makers must decide when, whether, and how users will prove their digital identity to access a potential CBDC. This report outlines how different CBDC designs can rely, among other access solutions, on conventional digital versions of current identification credentials, knowledge-based cryptographic keys, or a mix of different approaches. Policy makers' decisions regarding digital identities are broader than CBDCs, but the design choices will once again determine what type of CBDC architectures are possible. Thus, policy makers should include considerations about the cybersecurity profile of a potential CBDC when deliberating the future of digital identification.[61] Should the US Congress, for example, decide to create an entirely new digital identity infrastructure, such a system would need to be integrated at the outset with the cybersecurity frameworks of a potential digital dollar. Moreover, as explained in the below principle on interoperability, US policy makers would need to ensure that any domestic digital identity schemes are compatible with future global standards. To mitigate risks of accepting and sending foreign transactions, US policy makers and regulators would need to work with their global counterparts to make sure any transactions involving third countries comply with the appropriate US digital identity standards and safeguards. As a result, global standard-setting efforts to create secure, interoperable CBDC ecosystems could also help lead a push on harmonizing international digital identity regulations. The G7's "Roadmap for Cooperation on Data Free Flow with Trust," which focuses on "data localization, regulatory cooperation, and data sharing," could provide a high-level blueprint for harmonizing countries' digital identity approaches.[62]

To address privacy risks from a CBDC's increased centralization of payment processing and sensitive user data, governments must establish clear rules around who has access to which data, for what specific reason, and for how long. This includes explicitly delineating responsibilities of Anti-Money Laundering/Combating the Financing of Terrorism (AML/CFT) compliance between the private and public sector stakeholders of a CBDC. There are a range of actions that Congress could take in authorizing legislation for a possible CBDC. The Biden administration's 2022 Executive Order on Ensuring Responsible Development of Digital Assets directs the "Attorney General, in consultation with the Secretary of the Treasury and the Chairman of the Federal Reserve" to "provide to the President . . . an assessment of whether legislative changes would be necessary to issue a United States CBDC, should it be deemed appropriate and in the national interest."[63] This assessment, and any related or competing legislation that members of the US House of Representatives or the US Senate draft in the coming months, could advance important requirements regarding the overlap between security and privacy.

Specifically, Congress could consider the following measures in legislation related to a CBDC:

- **Original Collection:** Delineate or limit what personal information/consumer data is originally collected from consumers as part of a CBDC system and in daily transactions—and what should not be collected. For example, limit data related to the underlying item purchased, the location of the transaction (GPS coordinates), or other metadata available to the Fed or other actors in a disintermediated system.

- **Subsequent Deletion:** Set out a data retention or deletion policy, for example, requiring the periodic deletion (and/or meaningful anonymization) of CBDC data after a set period of time.

- **Universal Searches:** Establish internal security standards (including logs and audit procedures) about which personnel can search repositories of CBDC data—as well as how often and how extensively they may do so, and under what forms of supervision. (By way of comparison, other government databases have experienced problems when a rogue government employee has complete discretion to perform universal search queries across millions of sensitive records, for example, about a former spouse, an ex-girlfriend, or fellow employee.

- **Fourth Amendment:** Apply Fourth Amendment protections (and federal case law about unreasonable searches and seizures), including to personally identifiable information contained in CBDC repositories. Practically, this would mean that prosecutors would need a warrant to access certain personal records.

61 "Privacy and Confidentiality Options for Central Bank Digital Currency," Digital Currency Governance Consortium White Paper Series, World Economic Forum, November 2021, 17, https://www3.weforum.org/docs/WEF_Privacy_and_Confidentiality_Options_for_CBDCs_2021.pdf.

62 Fumiko Kudo, Ryosuke Sakabi, and Jonathan Soble, "Every Country Has Its Own Digital Laws. How Can We Get Data Flowing Freely between Them?" World Economic Forum, May 20, 2022, https://www.weforum.org/agenda/2022/05/cross-border-data-regulation-dfft/.

63 "Executive Order 14067 of March 9, 2022: Ensuring Responsible Development of Digital Assets," Code of Federal Regulations, 87 FR 14143, https://www.whitehouse.gov/briefing-room/presidential-actions/2022/03/09/executive-order-on-ensuring-responsible-development-of-digital-assets/.

- **Subpoenas and Review:** When civil subpoenas are applicable, consider transparency mechanisms and procedures that would allow citizens to seek review before a CBDC system or administrators discloses personally identifiable information.

- **Remedies:** Consider penalties or remedies that should be available if and when a privacy violation should occur (particularly when it is severe or pervasive).

- **Reports:** Require annual reports on privacy-related issues (including a review of breaches or relevant inspector general reports), for example, to the Privacy and Civil Liberties Oversight Board, with a courtesy copy to relevant House or Senate oversight committee(s).

PRINCIPLE 3: TEST, TEST, AND TEST SOME MORE

Governments should ensure that they have full access to, and can directly oversee, security testing and audits for all CBDC implementation instances. There are also security and procurement benefits to making the relevant code bases open-source, which the Federal Reserve Bank of Boston has chosen to do with its current collaboration with MIT's Digital Currency Initiative.[64]

When it comes to selection of a technical platform for pilot CBDC programs, policy makers should carefully consider the key contractual terms they negotiate with those vendors for who will own and have access to the code base and who will be responsible for testing and auditing that code. Regulators may find advantages to using multiple implementations and code bases to avoid relying on a single vendor (or a single, closed source code base) in a way that may lead to a single point of failure, but for each instance or implementation, governments will have to carefully negotiate these code ownership, maintenance, and testing responsibilities.

The importance of testing was highlighted in the recent executive order on cybersecurity as a tool for efficiently and automatically identifying vulnerabilities.[65] In the context of a CBDC, testing will be important at multiple layers of the stack. For example, at the hardware and application layers, wallet software and hardware should be tested for vulnerabilities that could enable attackers to steal funds from users, exfiltrate data, or prevent the execution of transactions. At the same time, central banks may be using smart contracts to govern the dissemination of funds. Smart contracts, which digitally facilitate the execution and storage of an agreement, will be critical to many future CBDC applications. Take government stimulus payments as a use case. For aid distribution to be governed by smart contracts in the future, in-person reviews conducted by engineers, who read the smart contract code and grant approval, may not be sufficient to ensure accountability. Bugs in smart contracts, which could incorrectly execute the dissemination of funds, have caused massive losses in cryptocurrencies already.[66] To complement in-person reviews, there is a strong case for instituting automated reviews for verifying smart contracts. One option is a technique called formal methods. In addition, regulators should consider lessons from other smart contract designs, for instance, in the Ethereum ecosystem, to craft policies for a gradual rollout that are designed to catch implementation or design errors early in smart contracts' development. Smart contract testing is itself an active area of research and may hinge upon the specific code in use.

At the consensus layer, third-party vendors may provide software that implements the database management software and/or consensus management software for validators. This software should be thoroughly tested for call sequences that can induce faults in the liveness and/or correctness of the system.

Especially in the early days of pilot programs, CBDCs will require extensive testing and security audits. Governments will either require in-house expertise to conduct these audits or contract with additional vendors to perform the necessary testing and security assessment. Open-source CBDC code bases may allow for more participation in the security testing process, especially when combined with longer-term bug bounty programs, but still require due attention to the security testing process. To enable this extensive testing and security audits, the US Congress must consider the appropriations accordingly as part of the budget process.[67]

PRINCIPLE 4: ENSURE ACCOUNTABILITY

Establishing accountability across all parts of a CBDC's technical design is a necessary precondition for a secure and resilient CBDC ecosystem in the face of cyberattacks. The previous principle illustrated the importance of testing software (including smart contracts) prior to deployment. However, testing alone is not enough. Every major piece of software deployed in practice has bugs, and the same will be true of CBDCs. Given this, CBDCs need to establish clear rules and policies

64 "Central Bank Digital Currencies," Federal Reserve Bank of Boston, accessed February 15, 2022, https://www.bostonfed.org/payments-innovation/central-bank-digital-currencies.aspx.

65 "Executive Order 14028 of May 12, 2021: Improving the Nation's Cybersecurity," Code of Federal Regulations, 86 FR 26633, https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

66 Simon Joseph Aquilina et al., "EtherClue: Digital Investigation of Attacks on Ethereum Smart Contracts," *Blockchain: Research and Applications*, 2 (4) (2021), 100028, https://doi.org/10.1016/j.bcra.2021.100028.

67 See Principle 6 below for additional details on pending congressional legislation.

surrounding accountability for errors, and resulting consequences. For example, if a CBDC deploys a smart contract that allows citizens to withdraw twice as much money as was initially intended, who is responsible? The developer of the smart contract? The company that hired the developer? The central bank? Such accountability policies should be determined ahead of time, along with a plan for dealing with eventual challenges and disclosing the relevant vulnerabilities if and when they arise. Similar problems might occur with certain CBDC designs that make it impossible to revoke fraudulent or contested transactions. Policy makers should establish clear lines of responsibility for public authorities, PSPs, and users to cover potential losses and refund payments. To minimize the risk of attackers using hardware vulnerabilities to infiltrate CBDCs, policy makers might also consider processes to certify hardware suppliers and collaborate with the private sector to secure all parts of the supply chain.

Another important need for accountability arises at the consensus layer. Particularly with CBDCs that rely on DLT technology,[68] it is paramount to clearly establish accountability requirements among validators on the blockchain. In DLT-based CBDCs, security hinges on most of the participating validators behaving correctly. If one validator or node is compromised, that compromise may have exploited a vulnerability that remains unpatched among other validators as well. For this reason, robust reporting requirements must ensure that all other stakeholders learn about security breaches as quickly as possible to reduce the risk of attackers exploiting the same vulnerability across multiple validators. This, in turn, mitigates the risk of validators approving faulty transactions. Concretely, there may be a need for baseline requirements to determine how quickly validators should notify other stakeholders upon discovery of a breach or malfunction. While analogous requirements exist for trade finance, the timescales for notifying other parties of a breach are much slower. In a DLT-based CBDC, validators' accountability, particularly with regard to reporting and vulnerability disclosure, becomes much more urgent because of the potential for cascading effects across the blockchain and CBDC ecosystem.

Liability considerations

Another set of important questions for policy makers to answer revolves around the issue of liability for CBDCs and who will be legally responsible for covering the costs of cybersecurity incidents (i.e., theft of consumer data or funds). The liability question illustrates different options policy makers have at their disposal to approach CBDC cybersecurity regulation. This is an area where existing financial regulation for traditional banking lays out clear and largely pro-consumer rules for financial fraud and theft. On the one hand, policy makers could aim to implement similarly specific consumer protection-oriented rules for CBDC implementation at the outset of their development, especially because these rules will not inhibit specific innovations in the technical design of the CBDCs. By placing some responsibility or liability for fraud on the operators of a CBDC implementation, policy makers can incentivize the groups designing these systems to invest in greater security and oversight without dictating exactly how those goals should be achieved. This approach potentially allows for greater flexibility than security regulations that dictate specific standards or controls, but it also might provide less concrete security guidance to the vendors responsible for designing these systems.

Standard setting

A different approach for policy makers would be to set concrete technical standards for CBDCs that include security and privacy protections. These standards do not yet exist, and they are unlikely to evolve until there are specific CBDC implementations that have been piloted for a longer period to move policy makers toward a concrete decision on CBDCs.

In many circumstances, it may be more effective for the federal government to consult with—or expressly rely upon—private or nonprofit consortiums that develop and maintain technical standards. Policy makers and industry stakeholders may find some useful road maps in the existing standards, like the EMV standard for chip credit cards or the Data Security Standard published by the Payment Card Industry Security Standards Council. Voluntary technical security standards and protocols, like SSL and TLS, provide another model for standards development, though because they are not mandated or accompanied by liability regimes that incentivize their implementation, these models may be of more limited utility for securing a large-scale CBDC implementation. For early stage, small-scale pilot projects, however, voluntary technical standards may suffice to help provide some security guidance to initial vendors and provide some early data on which standards are most effective at preventing security breaches.

PRINCIPLE 5: PROMOTE INTEROPERABILITY

In a domestic context, policy makers should develop rules to ensure that a CBDC is interoperable with the country's relevant financial infrastructure and can serve as an "effective substitute."[69] This will increase the resiliency of countries' financial systems against failures due to cyberattacks and is a key benefit of adopting a CBDC.

68 Eighteen central banks are currently exploring CBDCs using DLT.

69 "CBDC Technology Considerations," Digital Currency Governance Consortium White Paper Series, World Economic Forum, November 2021, 9, https://www3.weforum.org/docs/WEF_CBDC_Technology_Considerations_2021.pdf.

US Treasury Secretary Janet Yellen confers with colleagues at the 2021 G7 finance ministers summit at Lancaster House in London, Britain June 4, 2021. Source: Stefan Rousseau/PA Wire/Pool via REUTERS

To strengthen the security of CBDC systems, it is also critical to promote global interoperability between CBDCs through international coordination on regulation and standard setting. Through its body of research, the Atlantic Council has long stressed the need for US leadership "to shape the trajectory of CBDC"[70] and specifically develop strong international cybersecurity standards through fora, including the G20, Financial Action Task Force (FATF), and Financial Stability Board (FSB), to "ensure countries create digital currencies that are both safe from attack and can safeguard citizens' data."[71] The Biden administration's recent executive order on digital assets,[72] which outlined the US government's goal to take a more active role in global standard-setting bodies for CBDCs and encouraged US participation in cross-border CBDC pilots, is a welcome step forward. US policy makers should explore a transatlantic CBDC cross-border wholesale trial with an explicit focus on standards development and mitigation of cyber threats. By involving FATF and FSB in such a CBDC pilot, regulators could ascertain where current international standards provide sufficient protections, in what areas new rules are necessary, and what new regulations might look like.

Regulators should also study ongoing and completed cross-border CBDC trials, including Project Dunbar and mCBDC Bridge, to build on these projects' cybersecurity findings for future tests. Based on our research, we understand that several countries are interested in collaborating with the United States on cross-border pilot projects using both wholesale and retail CBDCs. Through its innovation hub and linkages with the banking industry, the Federal Reserve Bank of New York may be particularly well placed to lead on wholesale testing. Given the Federal Reserve Bank of Boston's continued work on a retail-based CBDC, it could facilitate retail testing with other central banks.

The cyberattack on Bangladesh Bank, as detailed in the appendix, illustrates the risk of attackers using cross-border financial infrastructure, in this case SWIFT, to infiltrate a central

70 *The Promises and Perils of Central Bank Digital Currencies*, US House Committee on Financial Services, 117th Cong. (2021) (statement of Julia Friedlander, Atlantic Council's C. Boyden Gray senior fellow and GeoEconomics Center deputy director), https://financialservices.house.gov/uploadedfiles/hhrg-117-ba10-wstate-friedlanderj-20210727.pdf.

71 Ibid., 9.

72 "Executive Order 14067 of March 9, 2022."

bank. While cross-border payments via CBDCs will be settled differently, the case of Bangladesh Bank underscored the importance of incorporating cybersecurity considerations into payment verification mechanisms from the outset. A question of central importance is how to handle incoming international transactions that are validated and confirmed using different, possibly weaker, security standards. Accepting such transactions (and building upon them) can have cascading effects at a faster timescale than in the traditional financial system.

It is important to note that the United States does not need to reach a final decision on issuing a CBDC to have enormous influence on the design of CBDCs around the world. If Congress were to authorize a limited cross-border testing project with the goal of determining cybersecurity vulnerabilities and protecting user privacy, this alone would send a strong signal to central banks that are further along in the CBDC process.

PRINCIPLE 6: WHEN NEW LEGISLATION IS APPROPRIATE, MAKE IT TECHNOLOGY NEUTRAL

In the United States, Congress has considered a sizable number of bills related to cryptocurrency, including several directly about CBDCs. For example, the bipartisan Responsible Financial Innovation Act introduced by Senators Luumis and Gillibrand requires an interagency report on cybersecurity standards and guidance on all digital assets including CBDCs.

Few of these draft bills have moved out of committee or gotten to a successful floor vote in Congress, so it is difficult to make nuanced recommendations about granular legislative changes or comparisons at this point.

Still, two overarching points are worth highlighting:

First, Congress is still in a prime position to study and oversee the application of federal cybersecurity laws to a potential CBDC. Past or pending legislation scarcely mentions cybersecurity in any depth. One of the more detailed provisions is in H.R. 1030, titled the "Automatic Boost to Communities Act," introduced by US Rep. Rashida Tlaib (D-MI),[73] which states that:

> "(i) (1) (G) Digital dollar account wallets shall comply with the relevant portions of the Bank Secrecy Act in establishing and maintaining digital dollar account wallets and shall impose privacy obligations on providers under the Privacy Act of 1974 that mirror those applicable to Federal tax returns under sections 6103, 7213(a)(1), 7213A, and 7431 of the Internal Revenue Code of 1986…

Table 3: Summary of US Congressional Activity Related to Cryptocurrency/CBDCs

PENDING LEGISLATION:

Since 2021, members of the US Congress introduced more than 35 crypto-related bills, including:

TAX REPORTING

- Infrastructure Investment and Jobs Act (H.R. 3684)
- Keep Innovation in America Act (H.R. 6006)
- Cryptocurrency Tax Clarity Act (H.R. 5082)
- Cryptocurrency Tax Reform Act (H.R. 5083)

BLOCKCHAIN

- Blockchain Innovation Act (H.R. 3639)
- Blockchain Promotion Act of 2021 (S. 1869)
- Blockchain Technology Coordination Act of 2021 (H.R. 3543)
- Blockchain Regulatory Certainty Act (H.R. 5045)

TAXONOMIES

- Digital Taxonomy Act (H.R. 3638)
- Securities Clarity Act (H.R. 4451)
- Token Taxonomy Act of 2021 (H.R. 1628)

RANSOMWARE

- Ransom Disclosure Act (S. 2943 and H.R. 5501)
- Sanctions and Stop Ransomware Act of 2021 (S. 2666)

CBDCs

- Central Bank Digital Currency Study Act of 2021 (H.R. 2211)
- 21st Century Dollar Act (H.R. 3506)
- A Bill to Require a Study of the National Security Implications of the People's Republic of China's Efforts to Create an Official Digital Currency (S. 2543)
- Automatic Boost to Communities Act (H.R. 1030)
- A Bill to Provide for Responsible Financial Innovation and to Bring Digital Assets within the Regulatory Perimeter (S. 4356)

COMPETITIVENESS

- A Bill to Require the Secretary of the Treasury to Submit to Congress a Report on Virtual Currencies and Global Competitiveness (S. 2864)
- US Virtual Currency Market and Regulatory Competitiveness Act of 2021 (H.R. 5101)

73 Automatic Boost to Communities Act, H.R.1030, 117th Cong., 1st Session (2021), https://www.congress.gov/bill/117th-congress/house-bill/1030/text.

> "(i) (3) (C) a Digital Financial Privacy Board shall be— (i) established by the Secretary to oversee, monitor, and report on the design and implementation of the digital dollar cash wallet system; (ii) maintained thereafter to provide ongoing oversight over its administration; and (iii) designed in such a way as to replicate the privacy and anonymity-respecting features of physical currency transactions as closely as possible, including prohibition of surveillance or censorship-enabling backdoor features."[74]

Also relevant is H.R. 2211, the "Central Bank Digital Currency Study Act of 2021," introduced by US Rep. Bill Foster (D-IL), which commissions a study including:

> "(1) consumers and small businesses, including with respect to financial *inclusion, accessibility, safety, privacy, convenience, speed, and price considerations* (emphasis added);
>
> "(7) data privacy and security issues (emphasis added) related to CBDC, including transaction record anonymity and digital identity authentication;
>
> "(8) the international technical infrastructure and implementation of such a system, *including with respect to interoperability, cybersecurity, resilience, offline transaction capability, and programmability* (emphasis added)."[75]

However, this bill, in particular, may have been overcome by the Biden administration's issuance of the Executive Order on Ensuring Responsible Development of Digital Assets on March 9, 2022.[76] That executive order commissions upwards of nine separate reports and repeatedly emphasizes the importance of privacy and developing a CBDC that comports with democratic values.[77] Of particular relevance to cybersecurity are the portions of the executive order that ask "the Director of the Office of Science and Technology Policy and the Chief Technology Officer of the United States, in consultation with the Secretary of the Treasury, the Chairman of the Federal Reserve, and the heads of other relevant agencies" to study "how the inclusion of digital assets in Federal processes may affect the work of the United States Government and the provision of Government services, including risks and benefits to cybersecurity."[78]

Second, Congress should keep in mind the overarching principle of technology neutrality, which augurs toward developing laws that apply evenhandedly to different technologies over time—as opposed to a specific technological product or feature that may exist today (but be upgraded or overtaken by other innovations tomorrow).[79] In the context of CBDCs, that may mean using incentives and accountability (described above), rather than setting a precise numerical threshold (for an acceptable number of cyber incidents per year, or precise NIST standards that are applied). Alternatively, Congress may consider setting CBDC security requirements at a fairly high level of abstraction and empowering a federal agency or private consortium to utilize their expertise to develop and periodically update the details.

74 Ibid.

75 Central Bank Digital Currency Study Act of 2021, H.R.2211, 117th Cong., 1st Session (2021), https://www.congress.gov/bill/117th-congress/house-bill/2211/text?format=txt.

76 "Ensuring Responsible Development of Digital Assets."

77 "What does Biden's executive order on crypto actually mean? We gave it a close read," *New Atlanticist* (Atlantic Council), March 11, 2022, https://www.atlanticcouncil.org/blogs/new-atlanticist/what-does-bidens-executive-order-on-crypto-actually-mean-we-gave-it-a-close-read/.

78 Ibid.

79 See, for example, Rajab Ali, *Technological Neutrality*, Lex Electronica, 14 (2) (Fall 2009), https://www.lex-electronica.org/files/sites/103/14-2_ali.pdf.

Conclusion

This report seeks to shine light on the novel cybersecurity risks for governments, the private sector, and consumers of introducing CBDCs. Our research demonstrates, however, that the design space for CBDCs is large, and offers policy makers and regulators ample options to choose a technological design that is both reasonably secure and leverages the unique benefits a CBDC can provide.

According to recent surveys about using CBDCs, privacy is consumers' number one concern.[80] Our analysis shows that privacy-preserving CBDC designs are not only possible, but also come with inherent security advantages that reduce the risks of cyberattacks. At the same time, the report explains that CBDCs can offer authorities regulatory oversight while providing strong user privacy. In short, cybersecurity concerns alone need not halt the development of a CBDC. It is up to policy makers to make the appropriate foundational design choices that will enable central banks and PSPs to develop safe CBDCs.

To address other, cross-border cybersecurity risks of introducing a CBDC, policy makers should promote global interoperability between CBDCs through international coordination on standard setting. This applies to all governments irrespective of whether they decide to develop a digital fiat. Imbuing the process to craft global CBDC regulations with democratic values is in the United States' national security interest. With more than 100 countries actively researching, developing, or piloting CBDCs, it is time to act to ensure domestic and international systems are prepared for the rapidly evolving digital currency ecosystems. The United States can and should play a leading role in shaping standards around the future of money.

80 *Eurosystem Report on the Public Consultation on a Digital Euro*, European Central Bank, April 2021, https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf.

Appendix:
Lessons from the Federal Reserve's Current Cybersecurity Measures for Deploying CBDCs

For understandable security reasons, the Federal Reserve (the Fed) has shared little detail about the vulnerabilities of its current systems and of the broader payments landscape. While this makes an exact evaluation of current dangers difficult, this report uses public information to outline cyber risks across the financial and payment systems. We focus on public and private wholesale layers and especially on Fed services since the central bank would presumably be the issuer of a digital dollar.

Fedwire, operated by the Fed, is the dominant domestic funds transfer system, handling both messaging and settlement. The Clearing House Inter-Payments System (CHIPS), privately operated and run by its member banks, fills a similar role for dollar-denominated international funds transfers.[81] The Society for Worldwide Interbank Telecommunications (SWIFT), operated as a consortium by member financial institutions, is a global messaging system that interfaces with Fedwire and CHIPS for the actual settlement of payments.[82] On the horizon, the Fed's FedNow promises instant, around-the-clock settlement and service, with a full rollout over the next two years.

While Chapter 1 assesses CBDC cybersecurity from a global perspective, this appendix focuses on the US payment system given the dollar's reserve and vehicle currency status, the Fed's centrality to the wholesale payment system, and the diversity of layers. Studying the Fed's cybersecurity system also sheds light on other countries' approaches as the Fed's payment cybersecurity practices are largely analogous, and often the model, to those of other central banks considering the deployment of a CBDC.

PUBLIC WHOLESALE LAYERS

Fedwire: The Fedwire Funds Service is a real-time, gross settlement (RTGS) system that enables "financial institutions and businesses to send and receive same-day payments."[83] RTGS means that payments immediately process and are irrevocable and that payments are not netted out over a longer time period. It operates twenty-two hours a day every business day and has thousands of participants who use it for "large-value, time-critical payments."[84] To make a transfer, the master account of the sending institution is debited by its Federal Reserve Bank, and the master account of the recipient institution is credited.[85] Payments are final, which makes it difficult to fix mistakes. In 2021, Fedwire handled more than 204 million transfers with a total value greater than $991 trillion, a sum more than forty times the United States' 2021 GDP.[86] This translates into an average value of $4.57 million, which is reportedly skewed by a small number of high-value payments.[87]

In assessing Fedwire's cybersecurity, the Fed aims for the core principle that it should possess "a high degree of security and operational reliability and should have contingency arrangements for timely completion of daily processing."[88] On the reliability front, Fedwire has an availability standard of 99.9 percent. In 2013, it exceeded this standard for all forms of access.[89] Any wholesale CBDC must achieve similar results to underpin the financial system. To preserve continuity of operations, the Fed focuses on both its own systems and those of Fedwire participants. The Fed requires high-volume and high-value Fedwire participants (core nodes) to participate in multiple contingency tests each year, including for their backup

81 Ibid.

82 Financial Crimes Enforcement Network, *Feasibility of a Cross-Border Electronic Funds Transfer*.

83 "Fedwire Funds Service," Federal Reserve Bank Services, accessed January 30, 2022, https://www.frbservices.org/binaries/content/assets/crsocms/financial-services/wires/funds.pdf.

84 "Fedwire Funds Services," Board of Governors of the Federal Reserve System, last updated May 7, 2021, https://www.federalreserve.gov/paymentsystems/fedfunds_about.htm.

85 "Fedwire Funds Service."

86 "Fedwire Funds Service - Annual Statistics," Federal Reserve Bank Services, last updated February 15, 2022, https://www.frbservices.org/resources/financial-services/wires/volume-value-stats/annual-stats.html; and Bureau of Economic Analysis, Gross Domestic Product, Fourth Quarter and Year 2021 (Advance Estimate), news release, January 27, 2022, https://www.bea.gov/news/2022/gross-domestic-product-fourth-quarter-and-year-2021-advance-estimate.

87 Anton Badev et al., "Fedwire Funds Service: Payments, Balances, and Available Liquidity," Finance and Economics Discussion Series (Washington, DC: Board of Governors of the Federal Reserve System, October 5, 2021), 12, https://www.federalreserve.gov/econres/feds/files/2021070pap.pdf.

88 Board of Governors of the Federal Reserve System, *The Fedwire Funds Service: Assessment of Compliance with the Core Principles for Systemically Important Payment Systems*, revised July 2014, 26, https://www.federalreserve.gov/paymentsystems/files/fedfunds_coreprinciples.pdf.

89 Ibid., 27.

sites.[90] To preserve the functionality of the core Fedwire service, the Federal Reserve Banks "maintain multiple out-of-region backup data centers and redundant out-of-region staffs for the data centers."[91] Thus, Fedwire's availability is secured both through redundant systems and endpoint security.

The Fedwire network has a "core-periphery" structure: the top five banks are responsible for around half of the payment volume, and the most important banks have a far greater number of network connections.[92] The concentration makes it a scale-free network: one with "most nodes having few connections but with highly connected hub nodes."[93] As a scale-free network, Fedwire has "significant tolerance for random failures but [is] highly vulnerable to targeted attacks." A random failure is likely to happen at a small institution, while a targeted attack on a core node could impact large amounts of transfers and severely reduce liquidity.[94] In this case, the Fedwire network could become "a coupled system where payments cannot be initiated until other payments complete," causing the entire system to grind to a halt.[95]

The Federal Reserve Bank of New York conducted a "pre-mortem" assessing how cyberattacks could disrupt Fedwire, specifically focusing on the type of targeted attack the network is vulnerable to.[96] The researchers assessed how a cyberattack impacting the availability or integrity (core elements of the CIA triad) of a top-five financial institution ripples through the wholesale payments network. They found that, excluding the target bank, "6 percent of institutions breach their end-of-day reserves threshold." When weighted by assets, this is equivalent to 38 percent of bank assets.[97] Breaching the reserves threshold means that reserves fall significantly below a bank's average level, impairing its liquidity and thereby its financial stability. The seizing up of the payments network is partly due to Fedwire's structure, which enables the receiving of payments even if an institution cannot send or observe them, which means the impacted institution could become a "liquidity black hole."[98] Such an institution would receive payments, and, therefore, liquidity, from the rest of the financial system but not send any payments to other institutions, thereby draining liquidity from other institutions. This spillover is magnified further if banks strategically hoard liquidity in response to the disruption. If the attack lasts for several days, liquidity shortfalls could grow to reach $1 trillion by the fifth day, requiring a massive intervention from the Fed.[99] Additionally, any attack on Fedwire could harm liquidity in financial market utilities (FMUs) like CHIPS and CLS, which are crucial to wholesale payments and foreign exchange markets, respectively.[100] Since Fedwire operates as the plumbing of these other forms of infrastructure (meaning it handles the final settlement of payments), any compromise of Fedwire would impact them.

Eisenbach, Kovner, and Lee document how escalating levels of private information about network interconnectedness (breaches of confidentiality) and days with large payment volumes allow attackers to maximize damage and systemic risk.[101] For example, an attacker who lingers in the network of a financial institution for months can observe payment patterns and choose the day when maximum damage will be inflicted.[102] One additional vulnerability of Fedwire is that third-party service providers are often shared across institutions, making them attractive targets for attackers looking to take down the network.[103]

The history of disruptions to Fedwire paints a mixed picture of its resilience. In the aftermath of the September 11 attacks, payment volumes rebounded despite financial infrastructure failing in Lower Manhattan and core nodes essentially ceasing to function.[104]

Perhaps no incident better captures the vulnerability of Fedwire, and the broader public-private wholesale payment system, than the attempted heist of Bangladesh Bank in 2016. Hackers infiltrated the network of Bangladesh Bank, which lacked a firewall and was poorly secured. The attackers used the bank's SWIFT messaging system to send fraudulent payment orders to the Federal Reserve Bank of New York. Despite issues with the messages that led them to be returned and

90 *The Fedwire Funds Service: Assessment of Compliance*, 27.

91 Ibid.

92 Thomas M. Eisenbach, Anna Kovner, and Michael Junho Lee, *Cyber Risk and the U.S. Financial System: A Pre-Mortem Analysis*, Federal Reserve Bank of New York, No. 909, January 2020, revised May 2021, https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr909.pdf.

93 Mark J. Bilger, "Cyber-Security Risks of Fedwire," *Journal of Digital Forensics, Security, and Law* 14 (4) (April 2020): 4, https://doi.org/10.15394/jdfsl.2019.1590.

94 Ibid., 4–5.

95 Ibid., 5.

96 Eisenbach, Kovner, and Lee, Cyber Risk.

97 Ibid., 2–3.

98 Ibid., 14.

99 Ibid., 41.

100 Ibid., 37.

101 Ibid., 24–26.

102 Ibid., 25.

103 Ibid., 32–33.

104 Bilger, "Cyber-Security Risks," 5.

resent, the differing time zones, work schedules, and absence of a communications channel between the two banks prevented Bangladesh Bank from being able to stop the New York Fed from transferring funds.[105] It took four days after the attack for communication to be established, and the Fed had already sent $101 million of funds through Fedwire via correspondent banks.[106] Nearly $1 billion could have been lost if not for the "total fluke" that the address of one recipient bank had the word "Jupiter," which was the name of an oil tanker and a sanctioned Athens-based shipping company, triggering further scrutiny.[107]

While not a direct attack on the Fedwire network, the Bangladesh Bank incident illustrates how the integrity of the current wholesale payment system is dependent on the practices of individual nodes. While the 2016 attack aimed at monetary gain and not explicitly at systemic disruption, the successful theft of $1 billion could have easily shaken confidence in the entire system. Additionally, the attack revealed the shocking reliance of the payment system on outdated technology. To its credit, the Federal Reserve Bank of New York set up a "24-hour hotline for emergency calls from some 250 account holders, mostly central banks" to prevent future miscommunication.[108] As discussed in Chapter 1, the ledger technology of CBDCs could enable innovations to reduce the likelihood of unauthenticated, fraudulent payments and enable faster communication. That said, quicker final settlement could add risks, since there is less time available for catching mistakes.

The cybersecurity challenges and approach of Fedwire are similar to those of other major central bank payment systems. For example, the European Central Bank's (ECB's) TARGET2 RTGS system relies on SWIFT for payment messages, exposing it to the vulnerabilities of that system.[109] Similar to the Fed, the ECB has focused on risks from TARGET2 participants via self-certification of information security and implementation of SWIFT's Customer Security Programme.[110]

FedNow: After a long series of delays, the Fed is planning to launch FedNow in 2023 or 2024. This is an RTGS system that, unlike FedWire, will operate twenty-four hours a day and three hundred and sixty-five days a year and offer instant payments that are irrevocable. As discussed earlier, instant payments, while convenient, limit chances to retract fraudulent payments. The service will be available to financial institutions with accounts at Federal Reserve Banks through the FedLine network, meaning it will not be available to nonbanks.[111] End users will encompass both individuals and businesses.[112]

While details are still limited, the Fed has promised to include fraud prevention tools to protect integrity, including transaction value limits (with a maximum set by the Federal Reserve Banks), conditions for rejecting transactions, and reporting features. Future features that may be implemented include aggregate transaction limits and centralized monitoring.[113]

PRIVATE WHOLESALE LAYERS

SWIFT: The Society for Worldwide Interbank Telecommunications (SWIFT) system is a messaging system used for international payments and run by a consortium of member banks. While FedWire and CHIPS handle both messaging and settlement, SWIFT only acts as a uniform messaging service for funds transfer instructions. Financial institutions can then "map" the SWIFT message into a FedWire or CHIPS message for the actual transfer of funds.

From January 2015 to January 2018, at least ten hacks were based on SWIFT, leading to initial losses of $336 million and actual losses of around $87 million.[114] As highlighted in the section on FedWire, one of these attacks was on Bangladesh Bank and relied on infiltrating its SWIFT messaging system. This is the chief vulnerability of the SWIFT system: attackers will access the messaging capability of a member bank, observe payment patterns, and then begin sending payment messages. Since the Bangladesh Bank hack, SWIFT has taken

105 Krishna N. Das and Jonathan Spicer, "How the New York Fed Fumbled over the Bangladesh Bank Cyber-Heist," Reuters, July 21, 2016, https://www.reuters.com/investigates/special-report/cyber-heist-federal/#:~:text=When%20hackers%20broke%20into%20the,into%20paying%20out%20%24101%20million.

106 Joshua Hammer, "The Billion-Dollar Bank Job," New York Times, May 3, 2018, https://www.nytimes.com/interactive/2018/05/03/magazine/money-issue-bangladesh-billion-dollar-bank-heist.html.

107 Das and Spicer, "How the New York Fed."

108 Ibid.

109 "Factbox: How Do Bank Payments Work in the Euro Zone?" Reuters, May 20, 2016, https://www.reuters.com/article/us-cyber-heist-ecb/factbox-how-do-bank-payments-work-in-the-euro-zone-idUSKCN0YB29H.

110 Jere Virtanen, "Endpoint Security in TARGET2," European Central Bank, Frankfurt, December 4, 2019, https://www.ecb.europa.eu/paym/groups/shared/docs/01eec-ami-pay-2019-12-04-item-5.2-endpoint-security-in-target2.pdf.

111 Margaret Tahyar, Jai Massari, and Andrew Samuel, "FedNow: The Federal Reserve's Planned Instant Payments Service," Harvard Law School Forum on Corporate Governance, August 31, 2020, https://corpgov.law.harvard.edu/2020/08/31/fednow-the-federal-reserves-planned-instant-payments-service/.

112 "Use Case Series: Unlock Instant Payment Use Cases with the FedNow Service," Federal Reserve Bank Services, accessed January 31, 2022, https://www.frbservices.org/binaries/content/assets/crsocms/financial-services/fednow/general-use-case.pdf.

113 Tahyar, Massari, and Samuel, "FedNow: The Federal Reserve's."

114 Antoine Bouveret, "Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment," IMF Working Paper WP/18/143, International Monetary Fund, June 2018, 13, https://www.imf.org/-/media/Files/Publications/WP/2018/wp18143.ashx.

several steps to shore up its defenses, focusing on stronger security standards and quicker response.[115]

This means that attacks are stopped during the preparation period before fraudulent transaction instructions are sent out. However, banks further down the payment chain can also stop transactions.[116] During the Bangladesh Bank hack, this was possible due to the lag in actual settlement. In a cybersecurity report, SWIFT notes that the role of other institutions will become even more important "as the speed of cash pay-outs increases."[117] Wholesale CBDCs could offer even faster payments, decreasing the time to retract a payment and requiring quick action to stop fraud by banks involved in settlement.

Following the Bangladesh Bank attack, SWIFT introduced the Customer Security Programme (CSP) with three pillars: "(1) securing your local environment, (2) preventing and detecting fraud in your commercial relationships, and (3) continuously sharing information and preparing to defend against future cyber threats."[118] Most recently, in 2019, SWIFT introduced the Customer Security Controls Framework (CSCF) as part of CSP. These require member banks to implement certain levels of security standards. The CSP has been successful in reducing successful attacks and securing SWIFT's integrity.[119]

While wholesale CBDCs will reshape the messaging and settlement functions of international payments, the SWIFT network's vulnerabilities illustrate the vital role of banks in securing their own systems.

RETAIL PAYMENTS

Physical cash: The most basic form of retail payments, and the only current public layer, is paper money. It is worth noting that cash also has security risks, even if these risks are not in the cyber realm. While the confidentiality and availability of cash is not of concern, cash can be counterfeited or physically stolen, damaging its integrity. The Treasury Department devotes technical effort to develop anti-counterfeiting features, such as holograms, paper selection, ink formulation, and artistic design, as well as other security choices, including serial numbers and storage at regional Federal Reserve Banks.[120] Federal Reserve Banks screen currency to identify possible counterfeits and send these to the Secret Service for investigation.[121] Additionally, physical cash provides perspective on the privacy trade-offs of CBDCs. While cash is largely anonymous, any cash transaction over $10,000 must be reported to the Internal Revenue Service (IRS) on Form 8300 to assist in combatting money laundering.[122]

Payment cards: Credit and debit cards are highly targeted by cybercriminals. In 2018, nearly $25 billion was lost to payment card fraud worldwide.[123] Such fraud, which often is part of identity theft, increased by more than 40 percent in 2020.[124] With new data breaches emerging more often than consumers can keep track of, enormous amounts of credit card information are floating around for purchase. In the past, payment card fraud often occurred in person, with criminals using "skimmers" to collect data at ATMs or gas stations and then replicating cards for use at point-of-sale terminals. Recently, online fraud has become more prevalent due to chip cards and the movement to e-commerce. Hackers now use digital skimmers, which entail installing malware in a merchant's website, to collect data for use in online purchases.[125] While credit card companies often offer fraud protection tools to protect consumers from losses, the prevalence and annoyance of credit card and identity theft shows the current retail payment system is far from risk free in terms of confidentiality and integrity.

The industry has taken steps to address this problem, with several major companies founding the Payment Card Industry Data Security Standard (PCI DSS) in 2006. PCI DSS is a set of twelve requirements, with penalties for noncompliance, to protect payment card data, and it applies to anyone storing,

115 *Three Years on from Bangladesh: Tackling the Adversaries*, SWIFT, April 10, 2019, https://www.swift.com/news-events/news/swift-report-shares-insights-evolving-cyber-threats.

116 Ibid., 2.

117 Ibid.

118 "SWIFT Customer Security Program," KPMG, 2021, https://assets.kpmg/content/dam/kpmg/qa/pdf/2021/04/swift-customer-security-program.pdf.

119 Adrian Nish, Saher Naumann, and James Muir, *Enduring Cyber Threats and Emerging Challenges to the Financial Sector, Carnegie Endowment for International Peace*, November 18, 2020, https://carnegieendowment.org/2020/11/18/enduring-cyber-threats-and-emerging-challenges-to-financial-sector-pub-83239.

120 "The Latest in U.S. Currency Design," U.S. Currency Education Program, accessed January 31, 2022, https://www.uscurrency.gov/sites/default/files/downloadable-materials/files/en/multinote-booklet-en.pdf.

121 Allison Chase, "Fed's Counterfeiting Experts Fight Flow of Fake Money," Federal Reserve Bank of Boston, October 15, 2019, https://www.bostonfed.org/news-and-events/news/2019/10/counterfeiting-experts-at-boston-fed-fight-flow-of-fake-money.aspx.

122 "Form 8300 and Reporting Cash Payments of Over $10,000," Internal Revenue Service, accessed February 23, 2022, https://www.irs.gov/businesses/small-businesses-self-employed/form-8300-and-reporting-cash-payments-of-over-10000.

123 "Credit Card Fraud Statistics," SHIFT Credit Card Processing, last updated September 2021, https://shiftprocessing.com/credit-card-fraud-statistics/.

124 "25 Credit Card Fraud Statistics to Know in 2021 + 5 Steps for Reporting Fraud," Intuit Mint, last modified December 17, 2021, https://mint.intuit.com/blog/planning/credit-card-fraud-statistics/.

125 *2021 Banking and Financial Services Industry Cyber Threat Landscape Report*, Intsights, accessed March 31, 2022, https://intsights.com/resources/2021-banking-and-financial-services-industry-cyber-threat-landscape-report.

processing, or transmitting this data. While PCI DSS is proven to reduce cyber risk, compliance is declining.[126] PCI standards can help payment providers work toward the CIA triad: the standards focus heavily on insecure protocols with the aim of protecting cardholder confidentiality.[127]

ACH: The Automated Clearing House (ACH) is a network operated by the National Automated Clearing House Association (Nacha) that aggregates transactions for processing and enables bank-to-bank money transfers.[128] In 2020, ACH handled more than $60 trillion in payments.[129] ACH is used for direct deposit of paychecks, paying bills, transferring money between banking and brokerage accounts, and paying vendors, and also underpins apps like Venmo.[130] This makes it a competitor to functions that a retail CBDC could fulfill, such as direct deposits of Social Security payments or tax refunds to individuals.

ACH is subject to fraud risks, though there are safeguards in place. Users must register with a username, password, bank details, and routing number. While these steps are similar to payment cards, ACH payments are not subject to the same PCI standards. That said, merchants can take additional steps like micro validation, tokenization, and encryption, and secure vault payments.[131] As with any retail payment system, many risks also stem from user behavior, such as falling prey to phishing scams. Overall, ACH payment fraud is relatively rare, accounting for only .08 basis points of all funds transferred.[132]

Digital payments: Payment services play a major role in facilitating online payments, and services like Stripe and Circle enable merchants to easily accept payments. While it is impossible to cover the cybersecurity risks of all these services, each has undergone security challenges and adaptations. For example, researchers recently found that attackers could target Apple Pay and bypass iPhone security through contactless messages that would drain the user of funds.[133] Two-tiered retail CBDCs would likely operate through many of the same current digital payments platforms, so security vulnerabilities and fraud opportunities could impact the rollout of a CBDC.

As discussed in this appendix, current wholesale and retail payment systems face a complex cybersecurity landscape and represent a major point of attack for both criminals and geopolitically motivated actors. Cybersecurity risks posed by CBDCs must be assessed relative to this landscape and how the technology could remedy existing vulnerabilities.

126 Leonard Wills, "The Payment Card Industry Data Security Standard," American Bar Association, January 3, 2019, https://www.americanbar.org/groups/litigation/committees/minority-trial-lawyer/practice/2019/the-payment-card-industry-data-security-standard/.

127 Alara Basul, "How PCI Compliance Is the First Step in Achieving the 'CIA Triad,'" Payment Eye, June 21, 2017, https://www.paymenteye.com/2017/06/21/how-pci-compliance-is-the-first-step-in-achieving-the-cia-triad/.

128 Rebecca Lake, "ACH Transfers: What Are They and How Do They Work?" Investopedia, April 30, 2021, https://www.investopedia.com/ach-transfers-what-are-they-and-how-do-they-work-4590120.

129 "ACH Network Volume and Value Statistics," Nacha, accessed January 30, 2022, https://www.nacha.org/content/ach-network-volume-and-value-statistics.

130 Lake, "ACH Transfers."

131 "Is ACH Secure?" Clover, accessed January 31, 2022, https://blog.clover.com/is-ach-secure/.

132 "Understanding the Basics of ACH Fraud," Sila, October 23, 2020, https://silamoney.com/ach/understanding-the-basics-of-ach-fraud.

133 Pieter Arntz, "Apple Pay Vulnerable to Wireless Pickpockets," Malwarebytes Labs, October 1, 2021, https://blog.malwarebytes.com/exploits-and-vulnerabilities/2021/10/apple-pay-vulnerable-to-wireless-pickpockets/.

About the Authors

Giulia Fanti is an Assistant Professor of Electrical and Computer Engineering at Carnegie Mellon University. Her research interests span the security, privacy, and efficiency of distributed systems. She is a two-time fellow of the World Economic Forum's Global Future Council on Cybersecurity and a member of NIST's Information Security and Privacy Advisory Board. Her work has been recognized with best paper awards, a Sloan Research Fellowship, an Intel Rising Star Faculty Research Award, a U.S. Air Force Research Laboratory Young Investigator Grant, and faculty research awards from Google and JP Morgan Chase.

Kari Kostiainen is Senior Scientist at ETH Zurich and Director of Zurich Information Security Center (ZISC). Before joining ETH, Kari was a researcher at Nokia. He has a PhD in computer science from Aalto. Kari's research focuses on system security. Recent topics include trusted computing, blockchain security, and human factors of security.

William Howlett is currently a senior at Stanford University, where he is studying economics and international relations and writing an honors thesis on US financial diplomacy towards China's current account surplus from 2009-13. On campus, he also conducts research for LTG H.R. McMaster on national security and economics. After graduation, William will be joining the Treasury Department as a junior fellow in the Office of International Monetary Policy, where he will work on IMF and G7/20 issues. He previously worked at the Atlantic Council's GeoEconomics Center and on the legislative team in California Governor Newsom's office.

Josh Lipsky is the senior director of the Atlantic Council's GeoEconomics Center. He previously served as an advisor at the International Monetary Fund (IMF) and Speechwriter to Christine Lagarde. Prior to joining the IMF, Josh was an appointee at the State Department, serving as Special Advisor to the Under Secretary of State for Public Diplomacy. Before joining the State Department, Josh worked in the White House and was tasked with helping plan President Obama's participation at the G-20 and other global summits. He is a term-member at the Council on Foreign Relations and an Economic Diplomacy Fellow at Harvard University's Belfer Center for Science and International Affairs.

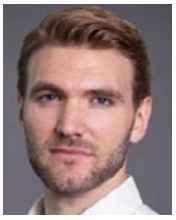

Ole Moehr is a nonresident fellow and consultant with the Atlantic Council's GeoEconomics Center. Previously, he served as the GeoEconomics Center's associate director. In Ole's current capacity, he contributes to the Center's future of money work and conducts research on global finance, growth, and trade. Ole's project portfolio includes work on global monetary policy, central bank digital currencies, global value chains, the EU's economic architecture, and economic sanctions. Prior to joining the Council, Ole served as a Brent Scowcroft Award Fellow at the Aspen Institute.

John Paul Schnapper-Casteras is a nonresident senior fellow with the GeoEconomics Center, focusing on financial technology, central bank digital currency, and cryptocurrency. JP is the founder and managing partner of Schnapper-Casteras, PLLC, a boutique law firm that advises technology companies, non-profits, and individuals about cutting-edge regulatory issues, litigation, and compliance. Previously, he worked on a broad array of constitutional and civil cases as Special Counsel for Appellate and Supreme Court Advocacy to the NAACP Legal Defense Fund and in the appellate practice of Sidley Austin LLP.

Josephine Wolff is a nonresident fellow with the Atlantic Council's Cyber Statecraft Initiative, an associate professor of cybersecurity policy at the Tufts University Fletcher School of Law and Diplomacy, and a contributing opinion writer for the New York Times. Her research interests include the social and economic costs of cybersecurity incidents, cyber-insurance, internet regulation, and security responsibilities and liability of online intermediaries. Her book "You'll See This Message When It Is Too Late: The Legal and Economic Aftermath of Cybersecurity Breaches" was published by MIT Press in 2018.

Acknowledgments

This report was made possible by the generous support of PayPal.

The GeoEconomics Center would like to thank Erinmichelle Perri, Ali Javaheri, Eli Clemens, Victoria (Hsiang Ning) Lin, Nathaniel Low, Claire (Ning) Yan, Thomas Rowland, and Jerry (Xinyu) Zhao for their important contributions to this report.

The GeoEconomics Center would also like to express their gratitude to Trey Herr and Safa Shahwan Edwards from the Scowcroft Center's Cyberstatecraft Initiative for their close collaboration in developing this report.

Atlantic Council Board of Directors

CHAIRMAN
*John F.W. Rogers

EXECUTIVE CHAIRMAN EMERITUS
*James L. Jones

PRESIDENT AND CEO
*Frederick Kempe

EXECUTIVE VICE CHAIRS
*Adrienne Arsht
*Stephen J. Hadley

VICE CHAIRS
*Robert J. Abernethy
*C. Boyden Gray
*Alexander V. Mirtchev

TREASURER
*George Lund

DIRECTORS
Stéphane Abrial
Todd Achilles
Timothy D. Adams
*Michael Andersson
David D. Aufhauser
Barbara Barrett
Colleen Bell
Stephen Biegun
Linden P. Blue
Adam Boehler
John Bonsell
Philip M. Breedlove
Myron Brilliant
*Esther Brimmer
Richard R. Burt
*Teresa Carlson
*James E. Cartwright
John E. Chapoton
Ahmed Charai
Melanie Chen
Michael Chertoff
*George Chopivsky
Wesley K. Clark
*Helima Croft
*Ankit N. Desai
Dario Deste
*Paula J. Dobriansky
Joseph F. Dunford, Jr.
Richard Edelman
Thomas J. Egan, Jr.
Stuart E. Eizenstat
Mark T. Esper
*Michael Fisch
*Alan H. Fleischmann
Jendayi E. Frazer
Meg Gentle
Thomas H. Glocer
John B. Goodman
*Sherri W. Goodman
Murathan Günal
Frank Haun
Michael V. Hayden
Tim Holt
*Karl V. Hopkins
Ian Ihnatowycz
Mark Isakowitz
Wolfgang F. Ischinger
Deborah Lee James
*Joia M. Johnson
*Maria Pica Karp
Andre Kelleners
Brian L. Kelly
Henry A. Kissinger
John E. Klein
*C. Jeffrey Knittel
Franklin D. Kramer
Laura Lane
Yann Le Pallec
Jan M. Lodal
Douglas Lute
Jane Holl Lute
William J. Lynn
Mark Machin
Mian M. Mansha
Marco Margheri
Michael Margolis
Chris Marlin
William Marron
Christian Marrone
Gerardo Mato
Timothy McBride
Erin McGrain
John M. McHugh
Eric D.K. Melby
*Judith A. Miller
Dariusz Mioduski
Michael J. Morell
*Richard Morningstar
Georgette Mosbacher
Dambisa F. Moyo
Virginia A. Mulberger
Mary Claire Murphy
Edward J. Newberry
Franco Nuschese
Joseph S. Nye
Ahmet M. Ören
Sally A. Painter
Ana I. Palacio
*Kostas Pantazopoulos
Alan Pellegrini
David H. Petraeus
*Lisa Pollina
Daniel B. Poneman
*Dina H. Powell McCormick
Michael Punke
Ashraf Qazi
Thomas J. Ridge
Gary Rieschel
Lawrence Di Rita
Michael J. Rogers
Charles O. Rossotti
Harry Sachinis
C. Michael Scaparrotti
Ivan A. Schlager
Rajiv Shah
Gregg Sherrill
Ali Jehangir Siddiqui
Kris Singh
Walter Slocombe
Christopher Smith
Clifford M. Sobel
James G. Stavridis
Michael S. Steele
Richard J.A. Steele
Mary Streett
Gil Tenzer
*Frances M. Townsend
Clyde C. Tuggle
Melanne Verveer
Charles F. Wald
Michael F. Walsh
Ronald Weiser
Maciej Witucki
Neal S. Wolin
*Jenny Wood
Guang Yang
Mary C. Yates
Dov S. Zakheim

HONORARY DIRECTORS
James A. Baker, III
Ashton B. Carter
Robert M. Gates
James N. Mattis
Michael G. Mullen
Leon E. Panetta
William J. Perry
Condoleezza Rice
Horst Teltschik
William H. Webster

**Executive Committee Members*

List as of May 2 2022

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

© 2022 The Atlantic Council of the United States. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Atlantic Council, except in the case of brief quotations in news articles, critical articles, or reviews. Please direct inquiries to:

Atlantic Council

1030 15th Street, NW, 12th Floor,
Washington, DC 20005

(202) 463-7226, www.AtlanticCouncil.org