Enhancing the Privacy of a Digital Pound

December 2024

Bank of England

digital currency initiative

mit media lab

Report Authors

Massachusetts Institute of Technology Digital Currency Initiative

Cambridge, Massachusetts, United States of America

Gabriela Torres Vives

Dr. Madars Virza

Reuben Youngblom

F. Christopher Calabia, CAMS

Bank of England

London, United Kingdom

Nick Vaughan

Zaki Said

Cinoj Mundakkal

Shantel Mullings

December 2024

Acknowledgments and Disclaimers

Acknowledgments

The authors thank Dr. Neha Narula, Director, MIT DCI, for her suggestions for this project and the resulting paper. We would like to thank Ashley Jacobson, MIT DCI Program Manager, and Ethan Olesinski, Student Research Assistant and Engineer, MIT DCI, for their contributions to this project as well.

MIT Digital Currency Initiative expresses its thanks to the Bank of England and its staff for collaborating in this paper's research and drafting and offering financial support for this project.

MIT Digital Currency Initiative also thanks the U.S. National Science Foundation and our program officer, Anna Brady-Estevez, for additional support as part of our ongoing research theme, “Promoting Privacy in the Use of Digital Currency.” MIT Digital Currency Initiative’s work on this paper was supported by the U.S. National Science Foundation under Award Number 2338198.

Disclaimers

This paper sets out the findings of the Bank of England and Massachusetts Institute of Technology Digital Currency Initiative’s collaboration on a central bank digital currency research project.

This paper does not represent, and should not be reported as representing, the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee. This paper cannot be taken to state Bank of England policy nor does it reflect the views of the Massachusetts Institute of Technology. No decision has been taken on whether to introduce a central bank digital currency in the U.K.

The U.S. National Science Foundation provided funding to support the work of Massachusetts Institute of Technology Digital Currency Initiative Researchers. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation.

Table of Contents

Abstract

With the increase in electronic payments options, privacy considerations are becoming even more important and privacy concerns more prevalent. The Bank of England and HM Treasury's 2023 Consultation Paper on the digital pound made clear that rigorous standards of privacy would be fundamental to trust and confidence in a digital pound and that measures would be put in place to ensure the public has confidence in using any digital pound, were one to eventually be launched.

Accordingly, and acting on feedback from the digital pound consultation where respondents emphasized their concerns around privacy, the Bank of England and HM Treasury committed to a range of measures that would govern a digital pound if the decision were made to launch it. The Bank of England and HM Treasury's Consultation Response stated that the Bank and the U.K. government would not have access to users' personal data, and legislation introduced by the U.K. government for a digital pound would guarantee users' privacy. In addition, the Bank of England committed to exploring technological options that would prevent the Bank from accessing any personal data through the Bank's core infrastructure.

The Bank of England and Massachusetts Institute of Technology Digital Currency Initiative have, over the past year, studied the possible application of privacy-enhancing technologies (PETs) to a potential digital pound, with the specific aim of identifying the potential technical challenges, trade-offs, opportunities and risks of using emerging types of PETs to support privacy. This research showed that emerging types of PETs, like pseudonymization, zero-knowledge proofs, and secure multiparty computing, might feasibly be applied to digital currency systems such as the digital pound to minimize the sharing of data both with the central bank and between payment intermediaries, giving users greater control over their data and enhancing user privacy. This presents opportunities for a digital pound to be at least as private as current forms of digital money and potentially even more private, although as with any technology, there are limitations to what emerging types of PETs can achieve. Tensions may also emerge between regulations that require the disclosure of data and the latitude to deploy PETs that need to be addressed. Future technology research is likewise necessary to understand evolving risks associated with various technical limitations and potential regulatory constraints to the application of PETs.

In November 2024, the U.K. government published its National Payments Vision,[1] which confirmed the government's intention to continue the design phase for the digital pound, in partnership with the Bank of England. The National Payments Vision set out that the design of a digital pound will maintain users' privacy and control over their money, and that any future

[1] HM Treasury, "National Payments Vision," Nov. 2024, pp. 9. [Online]. Available: https://assets.publishing.service.gov.uk/media/6736385fb613efc3f182317a/National_Payments_Vision..pdf

decision to proceed with a digital pound would be accompanied by legislation that would guarantee this.

This paper aims to inform public dialogue on a digital pound, and other central bank digital currencies, and encourage further research and dialogue particularly on the application of established and emerging technology options to enhance privacy. As the economy continues to digitize, our research is intended to contribute to discussions on digital currency technologies, helping to ensure that key democratic principles such as accountability, transparency, and privacy are considered in the design of future digital currency.

Chapter 1: Introduction

All electronic financial transactions generate data. In many cases, laws and regulations require financial services providers to use that data to verify the identities of customers and understand spending patterns to help mitigate the risk of facilitating financial crime. That data can also help businesses to understand customers better and design new products and services that are innovative and support their needs.

While the generation and use of data is an intrinsic part of electronic payments and can benefit consumers and businesses, it may also present privacy concerns, particularly if there are not sufficient safeguards. For example, where data from electronic payments reflect personally identifiable information, it could potentially give insight into consumers' habits, lifestyles, beliefs, personal preferences, or even health. As a result, layers of legal, operational, and technological safeguards are essential to protect privacy.

As the growth and pace of innovation in electronic payments accelerate, privacy considerations are becoming ever more important. The Bank for International Settlements reports that 94% of surveyed central banks are conducting research on whether to issue a digital version of their fiat currencies, known as central bank digital currencies (CBDCs).[2] CBDC will need effective mechanisms to safeguard user privacy.

In February 2023, the Bank of England and His Majesty's Treasury (HM Treasury) published a Consultation Paper to seek feedback on the design of a potential U.K. CBDC, known as a 'digital pound,' for use by households and businesses for their everyday payments needs.[3] The 2023 Consultation Paper on the digital pound made clear that rigorous standards of privacy would be fundamental to trust and confidence in a digital pound and that measures would be put in place to ensure the public has confidence in using any digital pound, were one to be launched eventually.[4] For example, the Bank of England and the U.K. government would not have access to users' personal data.[5] The Bank of England also published a Technology Working Paper that set out a non-exhaustive list of privacy-enhancing technologies (PETs) that might support privacy in a digital pound and that the Bank of England would assess further.

Notwithstanding these assurances, privacy was one of the main themes received in the feedback to the consultation. In response, the Bank of England and HM Treasury underscored the commitment that the Bank of England and the U.K. government would not have access to users' personal data, stating that legislation introduced by the U.K. government for a digital

[2] A. Di Iorio, A. Kosse, and I. Mattei, "Embracing diversity, advancing together – results of the 2023 BIS survey on central bank digital currencies and crypto," BIS Papers, no. 147, June 2024. [Online]. Available: https://www.bis.org/publ/bppdf/bispap147.pdf

[3] Bank of England and HM Treasury, "The digital pound: a new form of money for households and businesses?" Consultation Paper, Feb. 2023, pp. 5. [Online]. Available: https://www.bankofengland.co.uk/-/media/boe/files/paper/2023/the-digital-pound-consultation-working-paper.pdf

[4] Digital Pound Consultation Paper, pp. 12.

[5] Digital Pound Consultation Paper, pp. 12.

pound would guarantee users' privacy. The Bank of England and HM Treasury also explained that law enforcement agencies would have access to users' personal information only in limited circumstances and where there is a fair and lawful basis – as is the case today. In addition, the Bank of England committed to exploring technological options that would prevent it from accessing any personal data through the core digital pound infrastructure.[6]

To deepen our understanding of the technologies that might help safeguard privacy in a digital pound, staff from the Bank of England have collaborated over the last year with research scientists and engineers at the Massachusetts Institute of Technology Digital Currency Initiative (MIT DCI). Together we studied how privacy-enhancing technologies (PETs) might be applied to digital pound payments and whether those technologies might – theoretically – help protect user privacy and therefore trust and confidence in a digital pound system.

Indeed, in this paper, the authors consider whether a digital pound could protect or even enhance user privacy, given the ability to set strong standards of data protection around information used by the private sector, and what potential opportunities exist to implement 'data privacy by design' using PETs.[7]

This paper offers a primer on three PETs. We consider how PETs such as pseudonymization, zero-knowledge proofs (ZKPs), and several multiparty functionalities (such as Secure Multiparty Computation) might be applied to a potential digital pound platform. In particular, we assess design options and trade-offs that may exist for each PET, how these PETs could be applied in a hypothetical digital currency system, compatible with the Bank of England's proposal for a digital pound, and whether these PETs might help safeguard users' privacy in a digital pound while adhering to regulations related to anti-money laundering (AML) and countering the financing of terrorism (CFT).

The MIT DCI collaborates with central banks and other stakeholders to explore practical ideas and design options for digital assets built from the ground up for the public good. MIT DCI's work is intended to inform public dialogue about the benefits and risks associated with various technologies and architectures. MIT DCI does not endorse any particular design for digital currency or specifically encourage its adoption. In addition to support from the Bank of England to conduct this research, MIT DCI expresses its appreciation to the U.S. National Science Foundation and our program officer, Anna Brady-Estevez, for supporting our ongoing research theme regarding "Promoting Privacy in the Use of Digital Currency."

[6] Bank of England and HM Treasury, "Response to the Bank of England and HM Treasury Consultation Paper - The digital pound: A new form of money for households and businesses?" Consultation Response, Jan. 2024, pp. 28. [Online]. Available: https://www.bankofengland.co.uk/paper/2024/responses-to-the-digital-pound-consultation-paper

[7] Bank of England, "The digital pound: Technology Working Paper," Feb. 2023, pp. 23. [Online]. Available: https://www.bankofengland.co.uk/-/media/boe/files/paper/2023/the-digital-pound-technology-working-paper.pdf

The views expressed in this paper are those of the authors and do not reflect official policy of the Bank of England or HM Treasury or the views of the Massachusetts Institute of Technology. No decision has been taken on whether to introduce a digital pound.

MIT Digital Currency Initiative’s work on this paper was supported by the U.S. National Science Foundation under Award Number 2338198. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation.

Chapter 2: Privacy in financial services

2.1. What is privacy?

“Privacy,” according to one of the most widely cited academic studies of related terminology by Pfitzmann, A., and Hansen, M. (2010), “is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."[8] The Royal Society in the U.K. has adopted a definition that similarly emphasizes consent, describing privacy as “the right of individuals to selectively express themselves or be known.” Data privacy, in turn, is defined by the Royal Society as “entail[ing] a degree of control and influence over personal data, including its use.”[9]

Based on the definitions of privacy identified above, privacy is not anonymity. Pfitzmann and Hansen explain that the “Anonymity of a subject means that the subject is not identifiable within a set of subjects…” They go on to note that anonymity “ensures that a user may use a resource or service without disclosing the user’s identity.”[10] The Bank of England and HM Treasury note in their 2023 Consultation Paper that a digital pound would not be anonymous because, “just like bank accounts, the ability to identify and verify users is necessary to prevent financial crime.”[11]

In the 2023 Consultation Paper, the Bank of England and HM Treasury further stated that a digital pound would be subject to “rigorous standards of privacy and data protection,” and a digital pound “would be at least as private as current forms of digital money, such as bank accounts.”[12]

Notwithstanding these proposals, respondents to the Consultation Paper raised concerns about the potential impact of a digital pound on privacy. Privacy concerns reflect, in part, the increasing digitalization of the economy. Cash transactions in the U.K. and many other countries globally have declined significantly over recent years and are being replaced with digital payments. Compared to using cash, it is often easier, faster, and safer to send and receive payments digitally online and offline using credit and debit cards, bank transfers, or

[8] A. Westin, “Privacy and Freedom,” Washington and Lee Law Review, vol. 25, no. 1, March 1, 1967. [Online]. Available: https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20/ as cited in A. Pfitzmann and M. Hansen, “A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management,” version 0.34, pp. 6, Aug. 10, 2010. [Online]. Available: https://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf

[9] The Royal Society, “From privacy to partnership,” Jan. 2023, pp. 22. [Online]. Available: https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/from-privacy-to-partnership.pdf

[10] A. Pfitzmann and M. Hansen, 2010, pp. 9.

[11] Digital Pound Consultation Paper, pp. 72.

[12] Digital Pound Consultation Paper, pp. 12.

mobile payment apps. However, these digital payment instruments all generate data. Moreover, financial services providers and merchants may be able to gather and analyze this data, some of which may be personal in nature. It should be noted that laws and regulations may currently require financial services providers to gather, evaluate, and store some customer data to reduce the risks of fraud and other financial crimes, terrorist financing, or sanctions evasion. Having access to this information may also enable financial services providers and merchants to design products and services that may benefit a customer.

Despite the potential benefits to consumers, and the need to reduce risks of fraud and other financial crime, that information might also be misused unless there are sufficient safeguards to protect user privacy. When details about consumers' purchases and payments are available, the data may give insight into their lives, choices, preferences, and health.

Our research suggests that CBDCs might incorporate technology to enhance user privacy, enabling financial services providers to process transactions and comply with AML/CFT regulations while minimizing the sharing or processing of personal data. Our research focuses on compliance with requirements regarding customer due diligence and sanctions screenings, though the AML/CFT framework encompasses other requirements that are beyond the scope of this project. Cryptographic techniques like pseudonymization, ZKPs, and SMPC are feasible, from a technological perspective, for near-term deployment and might help ensure that digital pound transactions protect user privacy at least as well as existing digital payments such as credit cards and potentially to an even greater degree. However, further research might be required to investigate privacy architectures minimizing data storage, study data minimization techniques for *de minimis* transactions, and explore the performance of systems that enhance privacy for smaller value and hence lower-risk transactions while enabling payment intermediaries to maintain visibility for larger value and potentially higher-risk ones.

2.2. The platform model for a digital pound

The Consultation Paper proposed that a digital pound should be designed as a platform model. Under this conceptual model, a digital pound would be a public-private partnership. On the public side, the Bank of England would (1) issue a digital pound and (2) build and operate aspects of the digital pound infrastructure – including the core ledger and an application programming interface (API) layer, which would allow private sector businesses to access the core ledger.

On the private side, private businesses – called Payment Interface Providers (PIPs) – would provide digital wallets as the interface between the Bank of England and end users of a digital pound. PIPs would need appropriate regulatory authorization and permission from the Bank of England before being granted access to the ledger. Users' holdings of digital pounds would be recorded on the core ledger; PIPs would never possess users' digital pounds. Instead, PIPs would deal with all user-facing interactions, including handling customers' information, creating wallets, managing balances, and sending payment instructions.

Consequently, users would interact only with their PIPs rather than directly with the Bank of England. The Bank of England would not have access to users' personal data. As such, PIPs would be responsible for conducting any customer due diligence plus broader checks required under AML/CFT regulations. In addition to PIPs, External Services Interface Providers (ESIPs) could access the core ledger to support budgeting, analytical, or other innovative functions. PIPs and ESIPs would consequently have access to a user's personal data, which presents both opportunities and risks to the protection of user privacy.

2.3. Privacy in non-anonymous systems

Achieving privacy in any system, in practice, especially over the long term, is challenging. Non-anonymous systems present more challenges for two primary reasons.

The first reason is structural: non-anonymous systems such as those cited in the Bank of England and HM Treasury's Consultation Paper generally require that information linking a pseudonym or username to a real-world identity exists *somewhere*. The protections offered to consumers today mainly **block access** to these links. An alternative would be to design a truly anonymous system, where such identity data does not exist. This approach has not been preferred for a digital pound given the heightened risks of financial crime posed by such anonymous systems, and the importance of some, highly protected, identity data for delivering consumer protections, supporting the resolution of disputed transactions, and helping to tackle and remediate cases of suspected fraud.

From a consumer perspective, it is important to understand the distinction between (a) data that exists but is currently off-limits to others according to system design, regulation, and law – such as that envisaged in a possible digital pound ecosystem – and (b) data about the user that simply does not exist – as in truly anonymous systems. The former relies on legal requirements, organizational measures, and appropriate hardware and software protection. The latter provides strong protection but still leaves the possibility that other data, including offline data, might enable the identification of the account holder.

The second reason is statistical: in practice, financial systems tend to be susceptible to data analysis in surprising ways. Even without gaining access to the underlying data, non-anonymous and anonymous datasets can reveal patterns that users may be unaware of and provide (often strong) statistical inferences about the entities behind the transactions. Weaker levels of anonymity usually provide stronger statistical inferences. If available, fields like timestamps, amounts, party identifiers, and other individual data points can be compiled to create a more complete picture of the movement of money and used to re-identify constituent parties inside a financial system. This is not necessarily an unfavorable characteristic, depending on the ultimate goal of the system, but it is something that should be noted.

Given the focus of this paper on enhancing users' privacy while remaining in compliance with existing financial regulations, we consider that data obfuscation and encrypted data processing tools offer promising means of strengthening the privacy of the proposed digital pound and will define these tools in the next chapter. The rest of this paper focuses on three specific PETs that we consider may, over time and as the technology matures, become relevant for the platform model of a digital pound, namely the following:

- Pseudonymization, which is one of the simplest data obfuscation technologies and is already implemented in many existing financial workflows;
- ZKPs, which represent another promising data obfuscation technology widely used in blockchains; and
- Certain multiparty functionalities that enable encrypted data processing.

Chapter 3: Privacy-enhancing technologies

We now turn to the question of how new technologies – some not yet widely deployed in existing financial systems – might give users greater control over their own data and protect users' privacy. In particular, we are interested in determining whether these technologies could potentially deliver privacy protections beyond those found in current digital payment systems while remaining compliant with policy objectives such as AML/CFT requirements. The technologies this research considers are classified as PETs.

We refer to PETs as a collection of technical solutions that permit data processing and analysis while protecting the confidentiality of data. They are intended to strengthen privacy, autonomy, and freedom by affording individuals greater control over how, when, and to what extent their personal information is shared and with whom. PETs are not stand-alone tools; they can be combined and complemented with layers of legal, regulatory, and other tools to achieve privacy and data governance objectives.[13]

Researchers have made efforts to classify PETs in various ways. The Organization for Economic Co-operation and Development (OECD) proposes a framework that simplifies the classification of PETs into the following four primary categories:[14]

(a) **Data obfuscation tools**. These tools rely on cryptography as a key enabler, altering data "by adding 'noise' or by removing identifying details."[15] For example, differential privacy makes small adjustments when answering queries about a data set (e.g., by adding randomness sampled from an appropriate distribution) to mask details of individual data points, while keeping responses to aggregate queries useful. A well-known example of differential privacy is its application to releases of statistical information from the U.S. Census, where it minimizes the risk of exposing private data that could be linked to specific individuals and households.[16]

(b) **Encrypted data processing tools.** These are emerging technologies that allow computers to conduct calculations on encrypted data without seeing the data directly,

[13] C. Adams, Introduction to Privacy Enhancing Technologies: A Classification-based Approach to Understanding PETs, 1st ed. Springer Cham, 2021. [Online]. Available: https://link.springer.com/book/10.1007/978-3-030-81043-6

[14] OECD, "Emerging Privacy Enhancing Technologies: Current Regulatory and Policy Approaches," OECD Digital Economy Papers, no. 351, March 8, 2023. [Online]. Available: https://www.oecd-ilibrary.org/science-and-technology/emerging-privacy-enhancing-technologies_bf121be4-en

[15] OECD, pp. 15.

[16] Population Reference Bureau and U.S. Census Bureau, "Why the Census Bureau Chose Differential Privacy," 2020 Census Briefs, March 2023. [Online]. Available: https://www2.census.gov/library/publications/decennial/2020/census-briefs/c2020br-03.pdf

while keeping the underlying data unmodified.[17] For example, technology could further improve privacy in cross-border payments, ensuring sensitive customer information for a cross-border payment remains private when transmitted between financial institutions from different countries.

(c) **Federated and distributed analysis.** These technologies allow the execution of analytical tasks, like training models, on data sets "that are not visible or accessible to those executing the tasks." Sensitive data can "remain under the custody of a data source while it is analyzed by third parties."[18] In this case, a beneficiary financial institution receiving a cross-border payment might send its compliance verification algorithm to the originating financial institution in another country. The originating institution would provide the algorithm with relevant information to assist with verifying compliance and conformance with privacy regulations without transmitting customer information across borders. Federated analysis could be useful when data localization regulations forbid the sharing of consumer data outside the home jurisdiction, for example.

(d) **Data accountability tools.** These tools provide mechanisms for auditing and verifying data processing activities, enhancing transparency and accountability in data handling. They are frequently associated with PETs because they provide "new ways to require and enforce regulations about how data are processed, or by providing organizations and individuals with more agency and control over their data."[19] Financial institutions and payment processors that must audit customer data handling to comply with the U.K. GDPR might benefit from these tools.

It is important to note that PETs, on their own, do not guarantee privacy. If privacy is the abstract goal, PETs are the practical tools to help achieve that goal. PETs cannot substitute for legal frameworks that define privacy rights. Their use could be combined or layered with legally enforceable obligations to promote privacy and data protection rights. These technologies, which are at different stages of development, will likely need to be part of a broader data governance framework.

Cryptography methodologies such as PETs may involve the use of cryptographic keys, which are high-entropy (difficult to predict) bit strings that are meant to be kept secret. Deploying advanced cryptography can consequently come with challenges associated with managing those keys. Some of these challenges – for example, the effects of a user losing private keys, which are necessary for the user to decrypt encrypted messages or to sign digital messages – are shared with many already deployed systems and thus have known learnings and

[17] OECD, pp. 19.

[18] OECD, pp. 22.

[19] OECD, pp. 23.

mitigations. To retain our broad focus on PETs in this chapter, we address key management in Appendix 2.

We turn now to consider three PETs that might help to safeguard privacy in digital currencies: pseudonymization, ZKPs, and SMPC. These technologies might offer an opportunity to shield a user's sensitive data, including personally identifiable information, from others in the payments ecosystem. In the following sections, we will describe how a particular PET works at a high level and how it might be applied to design more private digital currencies. After discussing these uses, we will consider the potential technological and policy trade-offs associated with each technology.

3.1 Pseudonymization

Overview

Pseudonymization replaces user information or another piece of data with a reference, or identifier, that does not reveal that information. For example, instead of using a National Insurance number in the U.K., a Social Security number in the U.S., or an email address for a user (such as Alice) in a payment system, a PIP might generate a random identifier like "87584938475" to refer to Alice. Once the pseudonymization has occurred, one could use this identifier anywhere in the system where one might want to store data referencing Alice, such as in a transaction with a "sender" field.

Pseudonymization is a technique widely used in areas where personally identifiable information should be hidden, but records and data still need to refer to specific users in some fashion, such as in electronic healthcare records. In commerce, merchants and technology providers often use pseudonyms to obscure user data in databases and interactions to protect user information and comply with regulations.[20] Cryptocurrencies use addresses (which might derive from cryptographic public keys) as pseudonyms for users.

Using pseudonyms for user data creates an indirection layer, or mapping, between identifiers and users in the system, protecting personally identifiable data that does not need to be shown to the viewer of the pseudonym. It can help manage transaction data without attributing it to a specific individual. However, if the mapping of pseudonyms to personally identifiable information is revealed later, then this data can be attributed by using the mapping. In the previous example, only Alice's PIP would store the fact that Alice mapped to "87584938475," so only the PIP could use this mapping to determine Alice's transactions.

An important question is who is responsible for this mapping. In the previous example, Alice's PIP creates and maintains the mapping of "87584938475" to "Alice." In cryptocurrencies, users

[20] Square. "Payment Tokenization Explained." The Bottom Line, Oct. 8, 2014. [Online]. Available: https://squareup.com/us/en/the-bottom-line/managing-your-finances/what-does-tokenization-actually-mean

generate addresses on their own, without the help of a third party, and one user might have many different pseudonyms. Note that in the case in which a user has a PIP, the mapping is explicit, meaning the PIP actually stores it. In the case of cryptocurrencies, the mapping is implicit, because the user created it without an intermediary. Pseudonyms can furthermore be used to create public aliases for users. For example, a PIP could enable users to generate aliases to give to other users as payment addresses.

Pseudonymization has several benefits. First, it is a lightweight PET: It is easy to generate and use pseudonyms and does not require complex cryptographic operations. Second, it does not add significant overhead in terms of data storage or transmission requirements, so it will not significantly slow down systems that use it. Third, it is compatible with other privacy-enhancing techniques and can be combined and layered in a more advanced system to provide different properties. In particular, it is compatible with the separation of roles: a user's pseudonym issuer could be a third-party KYC (know your customer) provider and not necessarily the user's PIP. Finally, it is already widely used in practice, meaning its properties (including drawbacks) are well-understood and well-tested.

Pseudonymization alone, however, is insufficient to provide strong privacy in a digital currency. We will turn to the issues with pseudonymization after first discussing a concrete example of its application to digital currency.

Example of pseudonymization in a hypothetical digital currency

To visualize how pseudonymization functions, we discuss how one might use pseudonymization in a hypothetical digital currency to obscure user information from the central bank. This example is compatible with the Bank of England's Technology Working Paper. The following figure shows the flow of a payment transaction in which Alice is paying Bob in multiple steps. It includes Alice, Bob, Alice's and Bob's PIPs, and the core ledger operator. Note that this is just an illustrative example that makes certain assumptions about how the system operates and elides many important details.

(1) First, Alice obtains Bob's payment alias from Bob, which is "ABC." This might be a one-time alias or a permanent alias Bob uses. The alias helps hide any personally identifying information about Bob from Alice.

(2) Next, Alice creates a transaction to send to her PIP, which includes her identifying information, the payment alias for the recipient, and the amount to send. Note that the PIP knows that it is interacting with Alice and has presumably established a prior relationship with Alice.

(3) Alice's PIP communicates with Bob's PIP to learn Bob's pseudonym. Alice's PIP will either need to derive how to contact Bob's PIP from Bob's alias, or Alice will need to give her PIP information about Bob's PIP.

(4) Bob's PIP looks up the alias-to-pseudonym mapping it holds and uses it to find Bob's pseudonym.

(5) Alice's PIP replaces her identity in the transaction with her pseudonym ("123"), replaces Bob's alias with his pseudonym ("456"), and sends the pseudonymized transaction to the core ledger operator. Each pseudonym should serve as an identifier used at the core ledger operator to store funds Alice (or Bob) is authorized to spend. Note that this example does not discuss whether Alice's PIP maintains multiple pseudonyms for Alice, exactly how Alice and her PIP authorize the transaction, or how Alice or Bob obtain an account balance at the ledger operator in the first place.

(6) The core ledger operator applies the transaction it received from Alice's PIP, assuming account "123" has sufficient balance to enact the payment.

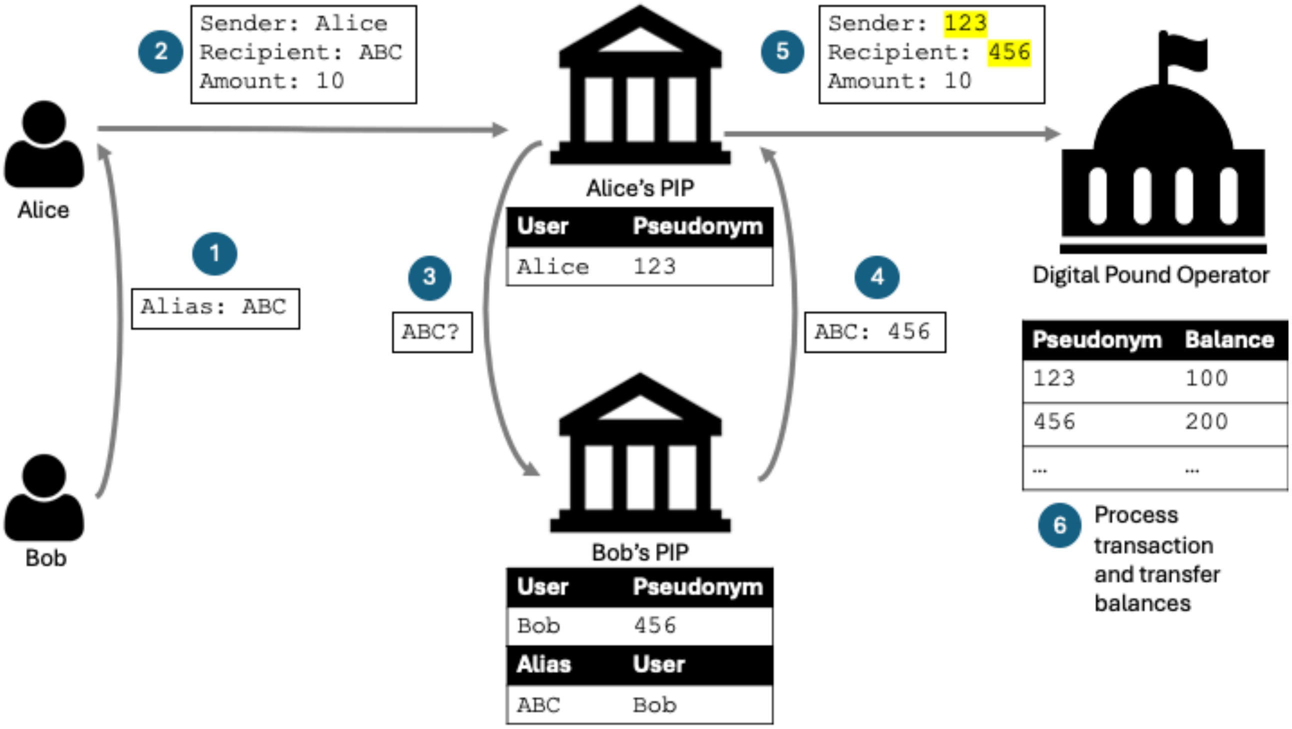

In this example, only Alice's PIP knows the mapping of "Alice" to her identifier, meaning that no one else who later sees the pseudonymized transaction (including the core ledger operator) can trace the payment back to Alice solely with the information in the transaction. However, because Alice's PIP knows its customers' identities, law enforcement agencies can still approach Alice's PIP. Provided they have the necessary legal authority and authorization, they may compel disclosure of the actual identity and resolve who is behind a pseudonym and a transaction. Consequently, in addition to the PIP knowing that the customer is Alice, law enforcement agencies can access this information about the actual customer via the PIP, where necessary and authorized by an appropriate body, such as a court. This arrangement ensures that the PIP can comply with the KYC requirements of the relevant AML regulations and that law enforcement can conduct investigations when necessary and authorized.

Limits of pseudonymization

Two major challenges must be addressed when using pseudonymization. First, pseudonymization alone is often not enough to guarantee strong levels of privacy. In the example above, though Alice's identity is kept secret from the core ledger operator, the ledger operator can still see the flows of funds and how money is spent between pseudonyms. This linkage is known as the transaction graph. If the goal is to hide the transaction graph, stronger cryptographic techniques would be required. There are some benefits to preserving the transaction graph, such as fraud prevention, but this trade-off is for policymakers to consider.

Risks remain that pseudonymized information can be compromised using different techniques. For example, there are commercial services that can combine various off-chain data sources with blockchain data to deanonymize and identify transactions on a blockchain as belonging to the same user, even when the user maintains many different addresses. Likewise, significant research exists on deanonymizing blockchain users.[21]

Second, pseudonymization is inherently at odds with restrictions enforced across pseudonyms. For example, consider a proposal that limits the amount of digital currency a user may hold, such as to a few thousand units, to reduce the systemic risk of disintermediating funds from the banking system.[22] Though Alice's PIP could enforce this limit, if Alice signed up for multiple accounts with multiple PIPs, the individual PIPs might need some mechanism for sharing enough information to enforce this holding limit across the entire system. Potential ways to address this requirement include adjusting the scope of a limit (by making it a limit per user per PIP instead of per user), by establishing information exchange functionalities, or by conducting probabilistic auditing. A concrete design is outside this paper's scope, but for suggestions on potentially useful avenues to address such a limit, refer to Appendix 3.

How pseudonyms are assigned will matter greatly to the success of a digital currency platform using pseudonymization. Policymakers should consider important trade-offs in how pseudonymization is used in an overall system design. Certain design goals might require intermediation for pseudonym generation, like enforcing holding limits. However, the intermediated setting poses challenges for offline uses, as it necessitates communication with the pseudonym issuer. Policymakers might also need to consider the risks of deanonymization at various payment points unless additional privacy techniques are applied.

[21] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, "A fistful of Bitcoins: characterizing payments among men with no names," in IMC '13, 2013, pp. 127-140. [Online]. Available: https://dl.acm.org/doi/10.1145/2504730.2504747

[22] U. Bindseil, "Tiered CBCD and the Financial System," European Central Bank, Working Paper, no. 2351, Jan. 2020. [Online]. Available: https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2351%7Ec8c18bbd60.en.pdf

3.2 Zero-knowledge proofs

Overview

A ZKP[23] is a cryptographic building block that allows one party, called the prover, to convince another party, called the verifier, that some claimed fact, called the statement, is true without revealing anything more than the fact that the statement is true. Specifically, the prover holds some secret data, called the prover's secret witness, and seeks to convince the verifier of a claim about the prover's secret.

For example, consider the scenario of secure document redaction. Here Alice (the prover) has a private document (e.g., a digitally signed bank statement) and wishes to reveal some but not all fields of the document to Bob (the verifier). To get a loan, she might want to reveal just the ending balance but not the transactions that led to it. If Alice just redacted the statement herself, the redacted statement would no longer bear a valid signature from the bank. This means that Bob might have little reason to accept Alice's redaction since there is no guarantee that Alice did not tamper with the fields she retained (e.g., by changing the ending balance).

However, using cryptography, specifically ZKPs, Alice can convince Bob that the following three facts hold:[24]

(a) that she possesses an unredacted bank statement in a digital form;[25]

(b) that this bank statement is signed by the bank (since the bank's public key is known to both Alice and Bob); and

(c) that the redacted statement is exactly the result of applying a prescribed redaction procedure (e.g., just retaining the ending balance).

Because it is cryptographically infeasible to produce “proofs” of false statements, Bob can be convinced that the original, unredacted bank statement has the claimed ending balance. Furthermore, the ZKP protocol reveals nothing more about Alice's bank statement than what is necessarily revealed by her three claims above. In particular, while Bob learns useful information from their interaction (Alice's ending balance), Alice may keep the original bank statement private.

Numerous types of ZKP protocols exist. They differ in the types of statements they can prove, their efficiency, the cryptographic assumptions they make, and details about the settings in

[23] S. Goldwasser, S. Micali, and C. Rackoff, “The Knowledge Complexity of Interactive Proof Systems,” in SIAM Journal on Computing, vol. 18, no. 1, pp. 186-208, 1989, https://www.doi.org/10.1137/0218012.

[24] This is similar to how the zkTax system proposes disclosures from signed tax returns. See, A. Berke, T. South, R. Mahari, K. Larson, and A. Pentland, “zkTax: A pragmatic way to support zero-knowledge tax disclosures,” 2023. [Online]. Available: https://arxiv.org/abs/2311.13008

[25] This embeds an important theoretical notion of proof-of-knowledge: that not only Bob is convinced that the statement is true but that Alice personally possesses the witness demonstrating it. Not all ZKPs are proofs-of-knowledge, but to simplify our discussion we will say ZKP to mean ZKP of knowledge.

which they operate.[26] However, they all realize the same common paradigm outlined above: an ability to prove claims about secrets without revealing them. In practice, most deployed ZKPs are non-interactive and publicly verifiable, meaning that Alice can produce a binary string (called the proof) that can later be verified without Alice's participation (non-interactivity) by anyone to whom this proof is presented (public verifiability). Of course, Bob is not required to reveal the proof to other parties. That said, these two properties are potentially useful in financial applications: Alice can post a proof on a public ledger or bulletin board, and then this proof can be verified by anyone who reads the bulletin board or ledger in the future (the verifier does not need to be known at the proof generation time), without any further involvement from Alice.

It is important to note that, in applications, ZKPs are rarely standalone: the statement being proved references an agreed-upon root of trust. For example, in the document redaction scenario above, it is not particularly useful to know that Alice has a document bearing someone's digital signature; Alice could have just generated a fresh signing key and signed the document herself. In contrast, it can be valuable to know that Alice possesses a document signed by a particular, specified entity (e.g., her bank). That way the integrity property of ZKPs provides a digital chain of custody for the document up to the root of trust (bank's public key). In contrast, the confidentiality property lets Alice maintain her privacy and choose what details about the document she wishes to disclose. Similarly, in a digital ledger system, ZKPs about transactions require a reference to a system state; otherwise anyone could invent a transaction that assigns themselves a large amount of money. The root of trust is the entire ledger itself: the statements are typically proved relative to a short summary of the ledger, e.g., a block hash that recursively commits to the entire system state up to that block.

ZKPs for financial applications

We now turn to several applications of ZKPs in the financial ecosystem where ZKPs might serve a role in strengthening users' privacy while enhancing system integrity or other relevant goals. As we will see, multiple applications could be relevant to a digital pound ecosystem.

Private transactions. ZKPs are widely used to achieve transaction privacy in distributed ledger systems. Deployed L1 systems[27, 28] use ZKPs to keep information private about a transaction's sender, recipient, and value and achieve strong, provable guarantees. Similar techniques are used in prominent shielded pools on smart contract platforms. Generally, transactions consist of encrypted payment details (e.g., sender, recipient, value), together with a ZKP establishing that the following statements are true:

[26] ZKProof, "ZKProof Community Reference," zkproof.org, version 0.3, D. Benarroch, L. Brandão, M. Maller, and E. Tromer, Eds., July 2022. [Online]. Available: https://docs.zkproof.org/pages/reference/versions/ZkpComRef-0-3.pdf

[27] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized Anonymous Payments from Bitcoin," in Proc. 2014 IEEE Symp. on Security and Privacy, Oakland, CA, 2014, pp. 459-474. [Online]. Available: http://zerocash-project.org/media/pdf/zerocash-oakland2014.pdf

[28] Zcash. (2024). [Online]. Available: https://z.cash/

(1) the transaction spends authentic (valid, relative to the current system state) funds;

(2) the spending has been authorized by the sender; and

(3) the transaction preserves balance.

ZKPs are also used in some variants of Chaumian eCash[29, 30] and parts of systems like CryptoNote[31]/Monero,[32] as well as others.

Transaction screening. Multiple proposals[33] suggest using ZKPs to mitigate privacy concerns in systems with strong privacy by either selectively revealing specific information about a transaction or constraining valid transaction types. For example, in the context of screening the recipient, ZKPs might be used to implement both inclusion (the recipient belongs to a particular set of recipients, e.g., persons who have been through a KYC process) and exclusion (the recipient does not belong to one of the sanctioned recipients) while otherwise keeping the recipient's identity private.

ZKPs for identity verification. Most financial systems require access authorization for at least two reasons. First, users must show they have a balance or funds to transfer and authorize the transfer. Second, in many digital payment systems, a transaction can be submitted only by an explicitly identified user to comply with KYC regulations. Though identification and authorization are needed to comply, at the same time, some users may seek greater privacy for small or *de minimis* transactions.[34] We can use ZKPs to align these two desires as follows: Alice can first authenticate herself with her PIP or ESIP and receive a digital certificate establishing her identity. Alice would use parts of this certificate with her payment later, as we describe. To address the first issue, the certificate would give Alice the ability to authorize a payment. Second, when Alice makes a payment, she could reveal or prove or prove various properties encoded in her certificate, without revealing the underlying data or her identity.[35] For

[29] Chaum's original proposal assumed a common denomination. See, D. Chaum, "Blind Signatures for Untraceable Payments," in Advances in Cryptology: Proc. of Crypto 82, 1983, pp. 199-203, doi: 10.1007/978–1-4757-0602-4_18.

[30] Follow-up works like S. Canard and A. Gouget (2007) uses ZKPs to provide additional features, such as, divisibility. See, S. Canard and A. Gouget, "Divisible E-Cash Systems Can Be Truly Anonymous," in EUROCRYPT 2007, 2007, pp. 482-497. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-540-72540-4_28

[31] N. van Saberhagan, "CryptoNote v 2.0," CryptoNote, Oct. 17, 2013. [Online]. Available: https://web.archive.org/web/20201028121818/https://cryptonote.org/whitepaper.pdf

[32] KOE, K. M. Alonso, and S. Noether, "Zero to Monero: Second Edition, a Technical Guide to a Private Digital Currency; for Beginners, Amateurs, and Experts," Apr. 4, 2020, ver. 2.0.0. [Online]. Available: https://www.getmonero.org/library/Zero-to-Monero-2-0-0.pdf

[33] J. Burleson, M. Korver, and D. Boneh, "Privacy-Protecting Regulatory Solutions Using Zero-Knowledge Proofs," a16z crypto, R. Hackett, Ed. Nov. 16, 2022. [Online]. Available: https://api.a16zcrypto.com/wp-content/uploads/2022/11/ZKPs-and-Regulatory-Compliant-Privacy.pdf

[34] An ESCB EUROchain proposal explores achieving this using "anonymity vouchers". An AML authority would issue these vouchers to every CBDC user at regular intervals, and one such time-limited voucher would permit one anonymous transfer of a small amount of money. See, European Central Bank, "Exploring anonymity in central bank digital currencies," IN FOCUS, no. 4, Dec. 2019. [Online]. Available: https://www.ecb.europa.eu/press/intro/publications/pdf/ecb.mipinfocus191217.en.pdf

[35] Technically this approach is called *anonymous credentials. See, D. Chaum, "Security without identification: transaction systems to make big brother obsolete," Communications of the ACM, vol. 28, no. 10, pp. 1030-1044, https://dl.acm.org/doi/10.1145/4372.4373, and J. Camenish and A. Lysyanskaya, "Signature Schemes and Anonymous Credentials from Bilinear Maps," CRYPTO 2004, vol. 3152, 2004, pp. 56-72, https://iacr.org/archive/crypto2004/31520055/cl04.pdf.* We use certificate and PKI language, as it is more common in financial applications.

example, some properties she may wish to prove could include, “I am a client of a particular bank,” “I am at least 18 years old” (for age-restricted purchases, if such verification is performed by payment intermediaries), or a combination of such. Each of those claims, while protecting Alice’s privacy, could enable her to make different kinds of payments in a tiered way: a *de minimis* payment could be made by any authorized user of the system, while a large value payment might require revealing more about her identity.[36]

Proof-of-reserves and privacy-preserving audits. Use cases described above have highlighted various ways in which a system's users can employ ZKPs to establish claims to a system while protecting their privacy. As we will see now, this dynamic can be reversed: a system can also use ZKPs to prove correct operation to its users or auditors. For example, systems like Provisions[37] have proposed that cryptocurrency exchanges could use ZKPs to prove their solvency (i.e., that their assets exceed their liabilities) to their users. The zkLedger system[38] proposes and prototypes an inter-bank settlement system simultaneously achieving (1) privacy for trading banks – their transactions hide sender, recipient, asset and value; and (2) the possibility for banks to make selective disclosures about their trading history. For example, a bank can produce a ZKP attesting to the fact that this bank is not overly exposed to a certain asset, all while keeping its actual exposure (and other financial details) private.

Lawful exceptional access. ZKPs have also been used in proposals for lawful exceptional access through data escrow.[39] Here, whenever Alice uses her identity to access a system, she could also escrow this identity (i.e., encrypt to an auditor's public key) and use a ZKP to prove that the escrowed identity truly matches the identity used. Then, at a later point, the auditor (e.g., law enforcement), when investigating a particular transaction and following an appropriate policy and process, could recover this identity. However, having all transactions have their identities stored in a centralized form (even in an encrypted form) increases the potential for a number of risks, including unintentional data leakage and intentional security breaches.[40, 41]

[36] Identity is a complex topic. We encourage readers who wish to learn more about ZKP and identity to explore zupass, a system for proving identity assertions in zero knowledge, and UTT, a system in which ZKPs and anonymous credentials are used to assign an anonymity budget for tiered KYC, and related works. See, zupass. [Online]. Available: https://github.com/proofcarryingdata/zupass and A. Tomescu, A. Bhat, B. Applebaum, I. Abraham, G. Gueta, B. Pinkas, and A. Yani, “UTT: Decentralized Ecash with Accountable Privacy,” VMware Research, April 9, 2022. [Online]. Available: https://eprint.iacr.org/2022/452.pdf

[37] G. G. Dagher, B. Bünz, J. Bonneau, J. Clark, and D. Boneh, “Provisions: Privacy-preserving proofs for solvency for Bitcoin exchanges,” in CCS ‘15, 2015, pp. 702-731, doi: 10.1145/2810103.2813674.

[38] N. Narula, W. Vasquez, and M. Virza, “zkLedger: Privacy-Preserving auditing for Distributed Ledgers,” in 15th USENIX Symp. on NSDI ‘18, 2018, pp. 65-80. [Online]. Available: https://www.usenix.org/system/files/conference/nsdi18/nsdi18-narula.pdf

[39] M. Kohlweiss, A. Lysyanskaya, and A. Nguyen, “Privacy-Preserving Blueprints,” in EUROCRYPT 2023: 42nd Annual Int. Conf. on the Theory and Application of Cryptographic Techniques, Apr. 2023, pp. 594-625, doi: 10.1007/978-3-031-30617-4_20.

[40] S. Landau and W. Diffie, “Keys Under Doormats: Mandating insecurity by requiring government access to all data and communications,” MIT Press, July 2015. [Online]. Available: https://mitpress.mit.edu/keys-under-doormats-security-report/

[41] *See also* Appendix 2

Challenges and limitations

ZKP cryptography is a rapidly advancing and maturing field. This technology has both been implemented and deployed in privacy-preserving cryptocurrency projects and shows relevance to privacy-preserving CBDC proposals. However, ZKPs are a powerful and complex cryptographic primitive, and, as such, their deployments are associated with a number of challenges.

Performance. While ZKP functionality is generic, the specific performance characteristics (time and computational resources needed to prove, verify, store, and transmit proofs) depend on the complexity of the statement being proved. Some statements have extremely efficient ZKP implementations, others pose a performance challenge, and still others are infeasible with today's technology. That said, ZKPs form an active ecosystem with rapid, ongoing development. Their tooling, performance, and hardware support are improving year by year. CBDC system designers could involve cryptography researchers and engineers early on to gain an understanding of potential performance and ways that a system design could be improved.

Design, code, and audit complexity. Specific domain expertise to ensure correctness, performance, and security, is required when designing systems that incorporate ZKP components. Similarly, domain expertise is required to audit these systems.

Standardization and recognition. Standardization of cryptographic primitives is important both for security assurance and interoperability. While ZKP standardization efforts are underway,[42] this is a work in progress. Similarly, from a legal standpoint, certain workflows might currently mandate some data to be provided and collected explicitly. Incorporating ZKP in such systems (i.e., replacing explicitly provided and checked data with a ZKP that such data exists and has passed certain checks) might require recognition of ZKPs as a valid alternative to explicitly provided data,[43] similar to how digital signatures gained recognition as a valid alternative to a physical (wet) signature.

3.3 Multiparty functionalities

ZKPs, described in the previous section, are a generic primitive: setting efficiency concerns aside, any computation about a particular secret can be proven in ZKP. However, by the very definition of ZKPs, this must be known to a single party, namely, the prover who is producing the said ZKP. Therefore, it is natural to ask whether it is possible to derive knowledge about secrets held by multiple parties, in such a way that results are trustworthy even though no participant is made privy to others' secrets. Secure multiparty computation (SMPC), introduced

[42] ZKProof, “ZKProof Standards,” [Online]. Available: https://zkproof.org/

[43] K. A. Bamberger, R. Canetti, S. Goldwasser, R. Wexler, and E. J. Zimmerman, “Verification Dilemmas in Law and the Promise of Zero-Knowledge Proofs,” Berkeley Technology Law Journal, vol. 37, no. 1, 2022. [Online]. Available: https://btlj.org/wp-content/uploads/2023/04/0001-37-1-Wexler.pdf

in seminal works of Yao,[44] and Goldreich, Micali, and Widgerson,[45] makes it possible to synthesize results securely based on several parties' information, without disclosing everything to all parties. A long line of research has culminated in practical implementations, and specific computations have been shown to admit particularly efficient protocols and have also been deployed in practice.

In this section we briefly outline a number of techniques that relate to computing data held by multiple parties, namely homomorphic commitments and private information retrieval (PIR), as well as applications of general-purpose multiparty computation.

Homomorphic commitments

In cryptography, a commitment scheme describes a process by which a party (called committer) can encode a piece of information to be opened up later. The requisite properties of a commitment scheme are for it to be hiding and binding. Here, hiding means that the commitment does not reveal anything about data committed to. Binding means that no one, including the committer itself, can later convincingly open the commitment to a different value than originally committed to. Commitment schemes can be useful in financial applications. For example, a commitment scheme can be part of a sealed-bid auction process. Here, the hiding property guarantees a bid's secrecy from other bidders, and the binding property means that a participant cannot change a bid after submitting it.

A particularly useful class of commitment schemes is called homomorphic commitments, which, in addition to being hiding and binding, also support carrying out certain computations on the committed values, without needing to know the values themselves. Here, anyone who holds two commitments, Com(a) committing to value a and Com(b) committing to value b, can produce a third commitment Com(c) to value c that is derived from a and b, without being privy to a and b. For example, in an additively homomorphic commitment scheme, any system participant can obtain a commitment Com(c), where c=a+b, given just Com(a) and Com(b).[46]

In a digital currency system, one might desire to shield transaction amounts from the ledger operator, while permitting the ledger operator to perform the integral function of transaction validation. Such validation can involve ensuring the user has the funds to spend and ensuring that the integrity of the money supply is maintained. One useful technique to achieve this is to store homomorphic commitments to transaction amounts instead of cleartext transaction amounts in the core ledger. Using this technique, in combination with ZKPs, one can construct digital currency systems that hide transaction amounts but support verification. This is

[44] A. C.-C. Yao, "How to generate and exchange secrets," in 27th Annual SFCS, 1986, pp. 162-167, doi: 10.1109/SFCS.1986.25.

[45] O. Goldreich, S. Micali, and A. Wigderson, "How to play ANY mental game," in STOC '87, Jan. 1987, pp. 218-229, doi: 10.1145/28395.284204.

[46] Pedersen commitments are a popular additively-homomorphic commitment scheme, where commitment to message m is value g^m h^r for a random value r.

demonstrated in Confidential Transactions and UTT.[47] One challenge is computing the amount of outstanding supply; one technique to mitigate this is to store encrypted transaction amounts under a separate key. In addition, other systems (zkLedger as well as ongoing research at MIT DCI on OpenCBDC privacy-preserving audits[48]) leverage the homomorphism to go beyond transaction validation and support audits over the outstanding supply. Note that this is an example of a multiparty protocol where one party is the auditor and the other party (most likely the user) engages in an exchange of information to satisfy an auditor's requirement for a certain property.

Private Information Retrieval

PIR encompasses a suite of cryptographic protocols that enable querying a remote database server and receiving a response, without revealing the exact query or specific response to the database operator.[49] This is a more advanced technique that is less widely used in practice in the financial context and in some cases can pose large communication or computational overheads on operation, though research is under way to make these techniques significantly more practical.[50, 51, 52]

Intuitively, one could implement a basic PIR scheme as follows: the person issuing the query (the client) sends the database operator (the server) a message that says “everything,” and the server sends back the entire database to the client. The client has not revealed their specific query to the server, and the server sends the entire database, so it did not learn anything unique to the client from the response. The client can query the database locally themselves once they've received it to find the answer to their real query. But this is undesirable and inefficient; the database might be large, and the server might not want to support queries that reveal everything to every client.[53] PIR schemes are designed to achieve this basic idea more efficiently, and, in some extensions, reveal less information to the client.

Two of the most widely-used implementations of closely related problems, Private Set Intersection (PSI) and Private Set Members (PSM), are used in Google Chrome; the first checks

[47] A. Tomescu, A. Bhat, B. Applebaum, I. Abraham, G. Gueta, B. Pinkas, and A. Yanai, “UTT: Decentralized Ecash with Accountable Privacy,” VMware Research, April 2022. [Online]. Available: https://eprint.iacr.org/2022/452.pdf

[48] See, for example, HalosGhost, “Enable supply auditing by storing cryptographic commitments,” GitHub, Issue #101, May 11, 2022. [Online]. Available: https://github.com/mit-dci/opencbdc-tx/issues/101

[49] B. Chord, E. Kushilevitz, O. Goldreich, and M. Sudan, “Private information retrieval,” Journal of the ACM, vol. 45, no. 6, pp. 965-981, Nov. 1998. [Online] Available: https://dl.acm.org/doi/10.1145/293347.293350 .

[50] C. Aguilar-Melchor, J. Barrier, L. Fousee, and M.-O. Killijian, “XPIR: Private Information Retrieval for Everyone,” in Proc. Privacy Enhancing Technologies Symp., 2016, pp. 155-174, https://petsymposium.org/popets/2016/popets-2016-0010.

[51] A. Henzinger, M. M. Hong, H. Corrigan-Gibbs, S. Meiklejohn, and V. Vaikuntanathan, “One server for the price of two: simple and fast single-server private information retrieval,” in Proc. 32nd USENIX Conf. on Security Symp., Aug. 2023, no. 218, pp. 3889-3905. Henzinger, Alexandra, Matthew M. Hong, Henry Corrigan-Gibbs, Sarah Meiklejohn, and Vinod Vaikuntanathan. "One server for the price of two: Simple and fast single-server private information retrieval." In Usenix Security, vol. 23. 2023.

[52] M. Zhou, A. Park, E. Shi, and W. Zheng, “PIANO: Extremely Simple, Single-Server PIR with Sublinear Server Computation,” in Proc. 2024 IEEE Symp. on SP, 2024, pp. 58, doi: 10.1109/SP54263.2024.00055.

[53] In particular, the PIR notion that provides privacy for both query and the database is called *symmetric PIR*.

whether a user's passwords around the web have been compromised without revealing those passwords to Google.[54] The second helps Chrome operating system (OS) devices enroll in a privacy-preserving way.[55] In this setting, a user's sensitive device information must be checked against encrypted Google databases to determine if a device is enrolled correctly or has a license.

Examples of how PIR might be useful in a financial context

An example is enhanced due diligence. Imagine a PIP needs to check with another party to determine if a transaction participant is on a list requiring additional data collection, in accordance with relevant AML regulations (see Section 4.1). It may not be desirable to send the identity of every recipient to the party, as this could reveal transaction information. Similarly, it may not be prudent to have every PIP download a copy of the list of users who require higher scrutiny, the regulatory watchlist, or other institutions' list of politically exposed persons,[56] since in some cases these lists may need to be kept private. Potentially a design using PIR or PSM could help address this. Another example in a similar spirit would be sanctions screening.

Challenges with secure multiparty computation

One challenge with SMPC such as PIR is that many schemes require multiple servers to achieve efficiency, along with an assumption that the servers do not collude.[57] If this assumption is broken, privacy can be lost. It is crucial to have a well-defined regulatory framework to address the main challenge of maintaining trust among a number of autonomous authorities and guarantee that they do not collude.

Another challenge with protocols like PIR, which broadly applies to many more complex PETs, is the challenge of carefully describing the problem PIR could be used to address, and determining whether it is a good fit for the exact privacy or security goals.

General-purpose multiparty computation

We can also consider applications of general-purpose multiparty computation to strengthen privacy and security of digital payments. For example, in a cross-border payment scenario, a PIP might engage in a multiparty computation with its foreign counterpart to perform certain types of screening. These might involve private information known only to the PIP's

[54] J. Pullman, K. Thomas, and E. Bursztein, "Protect your account from data breaches with Password Checkup," Google Security Blog, Feb. 15, 2019. [Online]. Available: https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html

[55] K. Yeo and S. Patel, "Protecting your device information with Private Set Membership," Google Security Blog, Oct. 28, 2021. [Online]. Available: https://security.googleblog.com/2021/10/

[56] Tookitaki, "Global Watchlist Screening - Types, Importance, and Solutions," March 8, 2024. [Online]. Available: https://www.tookitaki.com/glossary/global-watchlist#:~:text=Some%20common%20 types%20 include%20 sanctions,financial%20crimes%20and%20 illicit%20 activities

[57] "Collude" is the technical term for this sharing of information; it does not imply any prejudice against lawful collaboration.

counterparty but would better protect individual user's privacy than the current cross-border payment coordination which involves sending over private user data.

Finally, we can use multiparty computation and related techniques to distribute trust and authority. This includes systems that deliberately distribute information so that it is never available in a centralized form, and then use multiparty computation to reconstruct it. A notable concrete example here is using multiparty computation to protect cryptographic key material. Here, a party can independently generate many pieces (called shares) of a private key in different locations, potentially protected by different security mechanisms, and under control of different organizational units. Crucially, this share generation is done without ever materializing the combined secret key in a single location. Afterwards, to use this secret (e.g., to sign a transaction), different organizational units can cooperate to produce the signature jointly, again without ever holding a fully assembled key in a single place. The system can be set up in such a way that an attacker who obtains only a small number of these shares cannot recover the (implicit) secret key or perform the cryptographic operations that it enables, thereby strengthening the system against a wide range of possible breaches.

Before concluding this overview of the three PETs considered in this research, we close with thoughts on how designers of CBDC systems could consider the potentially widespread adoption of quantum computing in the future. This emerging technology could have significant implications for cryptographic technologies that are based on conventional computing platforms.

3.4 Post-quantum security

Any assessment of PETs in the context of CBDCs like a digital pound, specifically the cryptographic primitives on which they depend, should also consider compatibility with post-quantum computing schemes. Although quantum computing does not yet pose a threat to the classical asymmetric primitives commonly employed in key exchanges, algorithms based on integer factorization such as RSA, or the discrete logarithm problem such as Diffie-Hellman, it would be prudent to prepare for this threat. It is vital that the use of such technologies for CBDC does not introduce critical dependencies on cryptographic algorithms that are vulnerable to exploitation in a post-quantum world. As the U.K.'s National Cyber Security Centre recommends, the best mitigation against the threat of quantum computing in the future is the adoption of quantum-safe or quantum-resistant cryptography now.[58]

It will also be important to determine where vulnerabilities may still exist in the third-party provided infrastructure that is not managed by a central bank, but upon which the CBDC system may still rely. As the National Institute of Standards and Technology has stressed, it took nearly two decades to develop and establish the public key cryptography infrastructure

[58] National Cyber Security Center, "Next steps in preparing for a post-quantum cryptography," Nov. 3, 2023. [Online]. Available: https://www.ncsc.gov.uk/whitepaper/next-steps-preparing-for-post-quantum-cryptography

that we depend on today, so we should prepare for the future and abate a rushed, complex, and costly transition.[59]

The Bank of England's Technology Working Paper recognized these future risks, noting that the quantum computing threat is "an additional layer of risk that the Bank must factor into its CBDC design thinking." The Technology Working Paper also stated that crypto agility would be a design goal for a digital pound ecosystem.[60]

[59] National Institute of Standards and Technology, "Post-Quantum Cryptography," Jan. 3, 2017. [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography

[60] Digital Pound Technology Working Paper, pp. 28.

Chapter 4: The potential of PETs to improve privacy and compliance

We turn now to consider how PETs might be applied to a CBDC like a digital pound to improve privacy, including possible considerations around compatibility with the existing AML/CFT framework. We specifically consider whether PETs could give users greater control over who may see their personal information when making payments while enabling payment intermediaries to comply with regulatory requirements such as the customer due diligence and sanctions screening requirements that exist today. As noted earlier, PETs are not stand-alone tools. Additional mitigants might be needed to safeguard user privacy and support compliance with existing regulations.

4.1 Overview of U.K. regulations for anti-money laundering and countering the financing of terrorism

The core AML/CFT requirements for financial companies in the U.K. are set out in the Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations (the ‘Money Laundering Regulations’).[61] These regulations are informed by international standards, namely the 40 recommendations of the Financial Action Task Force (FATF).[62] Where an entity is defined as being a “relevant person,”[63] they must establish policies, controls, and procedures to mitigate the risks of money laundering and terrorist financing. This includes the following responsibilities:

- **Customer Due Diligence (CDD).** Where a business relationship is established, they must identify their users and verify their identities, including undertaking an assessment of the purpose and nature of the business relationship or occasional transaction.
- **Enhanced Due Diligence (EDD).** This is triggered where a higher risk of money laundering or terrorist financing is identified or prevalent. Examples include obtaining additional information on the customer and their source of wealth/funds.

[61] “The Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017,” U.K. Statutory Instruments, 2017, no. 692. [Online]. Available: https://www.legislation.gov.uk/uksi/2017/692/contents/made

[62] Financial Action Task Force, “International Standards on Combating Money Laundering and the Financing of Terrorism and Proliferation,” The FATF Recommendations, Feb. 16, 2012. [Online]. Available: https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html

[63] Examples of ‘relevant persons’ in the MLRs include financial institutions, credit institutions and auditors. It is not yet clear whether PIPs would fall under this definition in a CBDC ecosystem.

- **Sanctions screening.** They will need to screen customers against lists of sanctioned individuals prior to providing services, and potentially thereafter.[64]
- **Ongoing transaction monitoring.** This will be required throughout the business relationship and includes scrutinizing transactions.
- **Suspicious Activity Reports (SARs).** SARs must be made regarding information where they know or suspect that a person is engaging in money laundering or terrorist financing.

The AML/CFT framework as set out by FATF requires the collection and retention of specific personally identifiable information and transactional data to meet existing regulations during customer onboarding and transaction analysis. FATF standards and national implementation of those standards into AML/CFT regulations, as noted above, require financial institutions to identify the customers while on-boarding them and to monitor their activity continuously to detect and report suspicious transactions.[65]

In the case of commercial banking, this responsibility to conduct customer due diligence has traditionally fallen on the bank that opens an account for the consumer, making it necessary that the bank has access to that customer's information. For a digital pound, a PIP would likely be responsible for assessing any new customer seeking to open a digital pound account or wallet and for ongoing monitoring of the customer and related transactions to comply with the U.K.'s AML framework. After a digital pound wallet or account has been opened, a PET might be applied to transactions to ensure no transmission of user data, which could protect the privacy of users and enhance user control over their personal data. An additional benefit of using a PET could be to safeguard the user's identity from others involved in the clearing and settlement of a digital pound payment.

4.2 Potential applications of PETs for AML and CFT compliance

In summary, our research suggests that PETs are relevant to, and may offer benefits for, a digital pound design when considering both privacy and the core AML/CFT regulatory requirements for financial companies.

[64] Note that sanctions screening requirements derive from distinct legal and regulatory regimes relating to sanctions (rather than the Money Laundering Regulations); however, in practice it is common for sanctions screening to be conducted alongside AML/CFT checks.

[65] See, for example, Money Laundering, Terrorist Financing and Transfer of Funds (Information on Payer) Regulations 2017 (as amended), Article 29(6): the relevant person must not set up an anonymous account, an anonymous passbook or an anonymous safe-deposit box for any new or existing customer. The global standard is found in Financial Action Task Force, International Standards on Combating Money-Laundering and the Financing of Terrorism and Proliferation: Customer Due Diligence and Record Keeping, and in particular, Recommendation 10: Financial institutions should be prohibited from keeping anonymous accounts or accounts in obviously fictitious names.

Pseudonymization

Pseudonymization might enable financial services providers, especially PIPs, to shield their customers' data from other entities in a digital pound ecosystem when transactions are made by substituting a random identifier for the user's actual identifying information. As noted in Chapter 3, this substitution reduces the exposure of the consumer's identity to the central bank and potentially to others involved in the clearing and settling of a digital pound payment, thereby safeguarding the user's privacy. If questions arise about the transaction or the user later, law enforcement could still, in principle, access this information where necessary and authorized by law, such as by asking a court for an order that the PIP unveil the user's identity behind the pseudonym. This approach might assist a PIP in complying with the KYC requirements of the relevant AML regulations and in ensuring that law enforcement can conduct investigations when necessary and authorized.

A downside to the pseudonymization approach is that the mapping of identities to pseudonyms exists at the PIP and might, in remote circumstances, be revealed erroneously or inadvertently as part of an operational failure or cyberattack. Moreover, were the pseudonym to be revealed, it may be possible to identify a wide range of transactions that have taken place since the pseudonym's first use. Finally, as with other efforts to preserve the privacy of individual users, it may still be possible for third parties to discern the identity of the user without access to the mapping, especially if the third party can draw on other data sources in its analysis to narrow down the list of possible identities associated with a particular pseudonym.

Zero-knowledge proofs

In theory, ZKPs might allow PIPs to gather and process the information they need from users to fulfill their legal responsibilities to perform checks necessary for AML/CFT requirements, but still shield a user's personal data from others. For example, a user or enterprise seeking to make a digital pound payment could use ZKPs to prove to the central bank, ESIP, or other authorized entities that they have successfully passed through a recognized KYC process. As ZKPs allow the proving of a subset of information, by providing users with specially formatted digital identity certificates, PIPs could allow their customers to prove properties about their identity to a third party, without revealing their underlying identity information. Doing this would provide flexibility so that services could vary the amount of data a user discloses to suit their privacy needs based on the particular transaction they wish to conduct. A KYC proof could also help to shield a user's private information from others, including shielding the information from the central bank and from commercial enterprises involved in clearing and settling the payment, while ensuring that the transaction complies with the legal and regulatory customer due diligence and sanctions-screening requirements.

Users could moreover use zero-knowledge proofs to keep certain transaction details private in offline and other self-custody scenarios: a user's wallet could produce a zero-knowledge proof that (1) the transaction preserves balance and spends valid funds and (2) its value is below a threshold (and thus is not subject to certain checks), without explicitly revealing the amount transferred. This can afford greater privacy for *de minimis* transactions, while keeping large value transactions subject to the traditional AML/CFT process.

As noted in Chapter 3, a benefit of ZKPs is that most deployed ZKPs are non-interactive and publicly verifiable. This means that a ZKP could be generated earlier and then be later verified without any further participation by the user (non-interactivity) and by anyone who can see the proof (public verifiability). For example, a PIP could later present user's proofs to the regulatory authorities during regulatory inspections or audits.

While ZKPs have been deployed in production environments, they remain an emerging technology. Important questions concerning their impact on the performance of a digital pound will require research and experimentation to ensure the development of efficient proofs and test their impact on transaction speeds. Ensuring that the technology remains auditable will likewise require expertise and research.

Notwithstanding the assessment above, existing regulations might limit or otherwise impact the application of ZKPs for these purposes. For example, the U.K.'s Money Laundering Regulations require relevant persons (e.g., PIPs) to retain responsibility for conducting customer due diligence, notwithstanding they might use an agent to perform the relevant actions.[66] As such, it might be important to consider how the application of these regulations impacts responsibilities/liabilities of parties in utilizing ZKPs.

Multiparty functionalities

Cryptographic techniques can also aid secure collaboration between multiple parties, each of which holds a certain piece of sensitive information. Specifically, in a payment system, these might involve using homomorphic commitments to hide transaction values, maintaining overall system auditability. Here, a ledger operation could be convinced of total outstanding funds in the system, without knowing individual account or UTXO token balances.

Similarly, through either specialized or general multiparty computation, stakeholders not privy to some information held by others could synthesize useful results based on secret information held by different parties. For example, a financial services provider making a digital pound payment to a recipient could confirm that the sender is not on a list of unacceptable counterparties at the recipient institution without sharing which user wishes to make a payment and without seeing which counterparties are subject to heightened scrutiny at the recipient organization. This could be applicable in cross-border payment scenarios, where the sending bank and the beneficiary bank may currently both be required to screen the sender and the

[66] MLR reg 39(7)

recipient of a payment transaction to ensure that neither represents a sanctioned entity. Using PIR or multiparty computation, a bank intending to make a cross-border transaction could query the recipient institution's watchlist or sanctions list to ensure that the payer is permitted to send a payment to that institution without disclosing all of that customer's information to the recipient's bank.

Finally, multiparty computation and related techniques can also be used to enhance the security posture of cryptographic systems, through use of threshold cryptography and related key-management techniques specifically. These cryptographic techniques can help eliminate a single point of failure (when protecting integrity) or help aid a system of checks and balances (when protecting privacy).

Similar to the discussion on ZKPs above, the Money Laundering Regulations require that relevant persons retain responsibility for conducting customer due diligence. In particular, where anonymity is provided for, specific policies, controls and procedures are required under those regulations.[67] It might therefore be important to consider how the application of regulations impacts the responsibilities or liabilities of parties utilizing these multiparty functionality approaches.

4.3 Considerations for the collection, storage and processing of data for AML and CFT compliance

In light of the AML/CFT requirements set out above, policymakers might consider the following questions to analyze design options for a digital pound and trade-offs that may exist for each PET in preserving users' privacy while ensuring compliance with AML/CFT regulations:

(1) Which data should be stored for a digital pound transaction, in what format, at which stage of the payment journey, and by whom?

Data that is subject to AML/CFT regulations could be grouped into two categories, namely (i) customer or user data and (ii) transaction/payment data.[68] The user would be required to provide this data to the PIP, such as the user's own name or other identifier, a password or other means to validate the user's identity, the account that will fund the payment, the kind of transaction and its value, and information on the recipient of the payment, etc. Both categories of data contain private, personally identifiable information.

(2) Which data is being received or processed and by whom?

While privacy protections may be stronger when personal data is not received or processed at all, applying PETs might help to reduce the ability of others to see or

[67] MLR regs 19(4)(b) and 19A(4)(b)
[68] Based on the authors' analysis of existing messaging formats. See Appendix I.

process personal data. As one example, under existing regulations, PIPs would collect the originator's and beneficiary's names, addresses, and account details. The U.K.'s implementation of FATF's travel rule (FATF 16) requires PIPs to collect such information to identify the originator and beneficiary and transmit it to the receiving party's financial institution. If the transactions involve digital wallets, the sender must provide a wallet address instead of the account details. The sender's and the receiver's financial institutions or wallet providers use this information to assess the risks of a transaction facilitating money laundering or terrorist financing and to screen for payments to sanctioned individuals or politically exposed persons (or PEPs).[69]

By exploring various options for addressing these questions and assessing their likely impact throughout the payment process, policymakers might develop more informed insights on whether and how PETs may help to achieve the best balance in trade-offs and the most effective implementation methods as they mature and evolve.

[69] FATF defines a PEP as "an individual entrusted with a prominent public function." See, FATF, "Politically Exposed Persons (Recommendations 12 and 22)," FATF Guidance, June, 2013. [Online]. Available: https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Guidance-PEP-Rec12-22.pdf.coredownload.pdf

Chapter 5: Conclusions and Next Steps

Cryptographic techniques like pseudonymization, ZKPs, and multiparty functionalities may be increasingly feasible for possible deployment. These technologies have been deployed in many other contexts, showing their increasing maturity and practicality. A CBDC like the proposed digital pound might, in the fullness of time and as the technology matures, make use of these emerging PETs. PETs might help enhance user privacy and thereby support trust and confidence in a digital pound. In particular, PETs could, in theory, be used to help discharge responsibilities around KYC, validate transactions, and comply with core AML/CFT regulations, while minimizing the sharing of personal data. As such, PETs might help to align both the privacy objectives and the compliance requirements for a digital pound, presenting opportunities to use these technologies to provide stronger privacy protection than what exists today.

Nonetheless, additional work is required. Emerging technologies provide impetus for central banks to investigate alternative architectural approaches that could improve security and data management, presenting compelling opportunities for the future of money and its use in society.

In addition, imperfections will always exist in any technological solution, even if the best available cryptography and technology are deployed. For example, pseudonyms can create strong privacy for users by obscuring the link between a user's identity and the user's accounts. Yet pseudonyms might, in some circumstances, be mapped to users' real identities and could be inadvertently disclosed. Likewise, regardless of which PETs might be deployed, third parties may be able to discern who is behind a transaction by gathering and studying a wide range of information, such as combining publicly available transaction data (even if encrypted) with offline data and metadata, as seen with the deanonymizing of blockchain data. Moreover, there are limits to what can be accomplished with various cryptographic and security-related technologies. Whenever data exists and is stored somewhere, it might be susceptible to disclosure, whether by legitimate access, misuse by third parties, or exposure through leaks or hacks of encryption key systems. Furthermore, although it might be feasible to apply PETs to support regulatory compliance, regulators might need to recognize the application of PETs as valid compliance mechanisms.

Finally, certain trade-offs may exist between implementing privacy protections and complying with regulations that require the collection, storage, and disclosure of data. Policymakers, central bankers, and other stakeholders will need to consider these trade-offs.

Future technology research is necessary to address evolving risks associated with potential system vulnerabilities, encompassing ongoing monitoring, security considerations, novel approaches in the new architecture/interoperability, and additional architectures that may compromise the platform. We recommend the following steps for further study:

- investigating privacy architectures that minimize data storage (e.g., MIT DCI's OpenCBDC transaction processor developed with the Federal Reserve Bank of Boston[70]);
- examining architectures designed to enhance privacy for *de minimis* transactions;
- exploring complementary systems that would need adaptation in the event of a change to the current architecture;
- exploring how a higher degree of privacy could be provided to smaller (*de minimis*) transactions, while maintaining a higher degree of visibility in large transactions; and
- exploring whether and how existing laws and regulations limit opportunities for using PETs.

Dialogue is critical

Staff at the Bank of England and the Massachusetts Institute of Technology Digital Currency Initiative worked together on this project to help inform public dialogue on a digital pound in the U.K. as well as in other countries studying the potential issuance of a CBDC. As the economy continues to digitize, our research is intended to contribute to discussions on digital currency technologies, helping to ensure that key democratic principles such as accountability, transparency, and privacy, are considered in the design of future digital currency.

To achieve these goals, we welcome continued dialogue. It is important for central bankers, policymakers, financial services providers, technologists, innovators, and consumer advocacy groups to stay informed about innovations and their implications. It is beneficial for stakeholders to ask questions and share insights into recent cryptography and system design advancements. These and other financial innovations may present novel possibilities to benefit consumers and businesses.

While PETs, on their own, do not guarantee privacy, the approaches we explored in this paper seek to safeguard consumers' private information, enable compliance with existing regulations, and strengthen trust and confidence in a digital pound, should one be launched in the future.

[70] Massachusetts Institute of Technology Digital Currency Initiative and Federal Reserve Bank of Boston. OpenCBDC. Feb. 2023. [Online]. Available: https://github.com/mit-dci/opencbdc-tx

Appendices

Appendix 1: Payment data required in transactions in the U.K. ecosystem[71]

P2P transaction

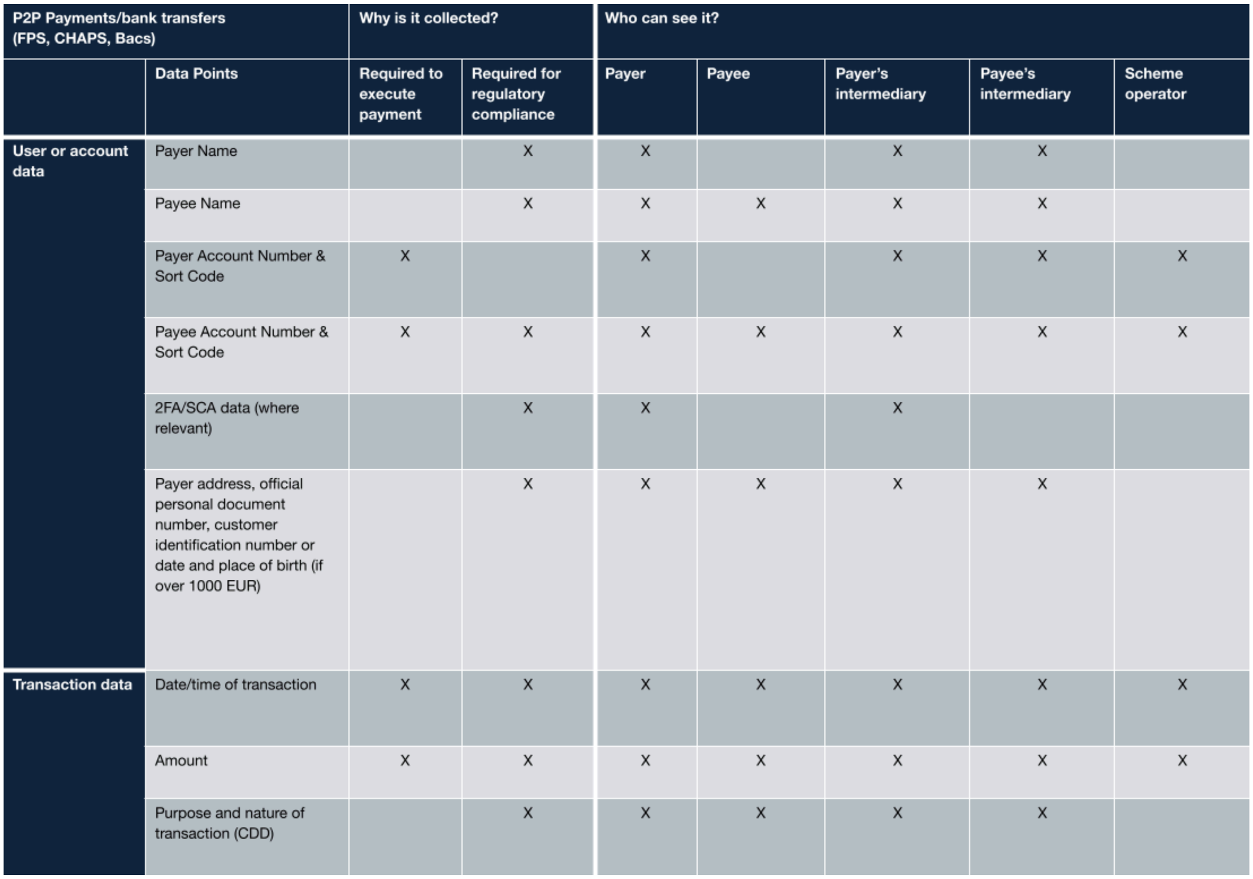

P2P Payments/bank transfers (FPS, CHAPS, Bacs)		Why is it collected?		Who can see it?				
	Data Points	Required to execute payment	Required for regulatory compliance	Payer	Payee	Payer's intermediary	Payee's intermediary	Scheme operator
User or account data	Payer Name		X	X		X	X	
	Payee Name		X	X	X	X	X	
	Payer Account Number & Sort Code	X		X		X	X	X
	Payee Account Number & Sort Code	X	X	X	X	X	X	X
	2FA/SCA data (where relevant)		X	X		X		
	Payer address, official personal document number, customer identification number or date and place of birth (if over 1000 EUR)		X	X	X	X	X	
Transaction data	Date/time of transaction	X	X	X	X	X	X	X
	Amount	X	X	X	X	X	X	X
	Purpose and nature of transaction (CDD)		X	X	X	X	X	

[71] To develop these tables, the authors analyzed requirements in the Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017 (as amended), U.K. Statutory Instruments 2017, No. 692. [Online]. Available: https://www.legislation.gov.uk/uksi/2017/692

Retail transaction

Card Payments		Why is it collected?		Who can see it?				
	Data Points	Required to execute payment	Required for regulatory compliance	Payer	Payee	Payer's intermediary (Card Issuer)	Payee's intermediary (Merchant Acquirer)	Scheme operator
User or account data	Payer Name (i.e. cardholder)		X	X	X	X	X	X
	Payee Name (i.e. merchant)		X	X	X	X	X	X
	Payer Card Number	X	X	X	X	X	X	X
	Payee Card Number	X	X	X	X	X	X	X
	Authentication data (SCA)	X	X	X		X		
	Payer address, official personal document number, customer identification number or date and place of birth (if over 1000 EUR)		X	X		X	X	
	Other payer card information (PIN, CVC code)	X	X	X	X	X	X	
Transaction data	Date/time of transaction	X	X	X	X	X	X	X
	Amount	X	X	X	X	X	X	X
	Merchant location		X	X	X	X	X	X
	Institutional identifier for issuer and merchant acquirer	X	X		X	X	X	
	Purpose and nature of transaction (CDD)		X	X	X	X	X	

Appendix 2: Key Management

Many privacy-enhancing technologies rely on cryptographic keys and the ability to safeguard them. Deploying these PETs can consequently come with challenges associated with managing those keys. Some of these challenges – for example, the effects of losing private keys – are shared with many already deployed systems and thus have known learnings and mitigations. However, others are unique to recent innovations in digital assets. We emphasize that these challenges are not insurmountable and can be mitigated and thus do not prejudice any particular approach or design. However, they should be taken into account when evaluating cryptographic systems.

Key loss. Cryptographic systems use possession of secret key material as an implicit delineation between authorized and unauthorized users of the system. In a reasonably private system, these keys are, for security purposes, entirely under the user's control—meaning that the responsibility of safeguarding them also falls on the user in a way for which some users may not be ready. In the cryptocurrency space, for example, users who lose their secret keys also lose access to their funds, often permanently. This property of self-custody is a security feature of cryptocurrency in which a third party cannot redirect, censor, or reverse one's payments. To guard against loss of funds due to key loss, users could seek the help of a third-party intermediary who could, for example, help users backup their key material.

Key theft. Stolen cryptographic keys can cause damage in multiple ways. The most obvious is the potential loss of funds and the associated challenge of a payment's reversal (if part of the system). Beyond that, a question remains regarding reputational attacks. Because secret keys are digital objects, their copies are bit-perfect equivalents of the originals. Digital signatures arising from copied keys are indistinguishable from those arising from the originals. Consequently, stolen keys can be used to engage in unauthorized and potentially illegal finance, making attribution and forensic analysis challenging.

Integrity risks. Certain digital asset proposals rely on the secrecy of key material to preserve core financial system invariants, such as the total supply of funds. For example, the original eCash proposal[72] relies on a trusted mint that produces signed serial numbers in exchange for one unit's worth of deposits in the system. The signing process relies on a special signature scheme (blind signatures) so that the mint does not learn the serial number it is signing, instead learning only that it has signed one serial number. The serial number together with the signature (also private from the issuer) is later revealed to make one unit's worth withdrawal from the system.

Cash has attractive privacy and performance properties. However, if the mint's signing key is compromised, this can lead to covert and untraceable issuance of unauthorized money as

[72] D. Chaum, "Blind Signatures for Untraceable Payments," in Advances in Cryptology: Proc. of Crypto 82, 1983, pp. 199-203, doi: 10.1007/978–1-4757-0602-4_18.

unauthorized signatures are indistinguishable from real deposits. Such attacks can be mitigated through threshold cryptography that splits the mint's signing key into multiple parts, called shares, such that a certain number (say, a majority) of these shares is required to perform a cryptographic operation. Such threshold issuance variants of eCash are, for example, used in the cryptocurrency world.[73] Furthermore, these key shares themselves could be stored in secure hardware modules for increased security. Another approach could be per-account holding limits: that way any potential excess emission is limited by the attacker's ability to obtain and use authorized user accounts.

Such challenges are not exclusive to eCash. For example, many ZKP systems rely on one-time trusted system parameter generation and mitigate this trust assumption through a multiparty setup ceremony.[74,75]

Privacy risks. In encrypted systems, the safety of the underlying data is only as secure as the key(s) protecting it. This risk is heightened as larger and larger stores of data are secured by keys with fewer and fewer points of failure (either due to fewer keys, or more keys under centralized control). Consider, for example, the scenario of lawful exceptional access. A hypothetical system could offer cash-like privacy (hiding sender, recipient, and amount) with the following caveat: each transaction also carries encrypted payment details (i.e., identities of sender, recipient, and amount) to be placed in escrow. To enhance data security, such escrow could even use multiple keys: e.g., requiring cooperation between a law enforcement agency, system operator, and an ombuds office to decrypt the payment details. Such a hypothetical system is a clear privacy improvement over traditional electronic payments: plaintext transaction details are not retained. However, in a key compromise scenario, the damage is much greater than in a traditional payment system: one can now retroactively decrypt every single payment and do so in a centralized escrow database. In comparison, traditionally the plaintext data would only be kept by the corresponding PIPs.

This highlights the need to focus on proper key management, data retention, and data separation. For example, if new escrow keys are used each month, then a compromise of February's keys would not help decrypt January's transactions. Similarly, if last year's ciphertexts are beyond statutory retention periods, they could be promptly expunged together with the keys that could be used to decrypt them. Likewise, to prevent centralization, the escrowed data could be kept at individual PIPs. Of course, if certain transactions (e.g., low value payments) are not subject to data escrow, one should verify that their data flow does not produce such escrowed data.

[73] Fedimint. "An overview of the Fedimint System, Core Technology Components." [Online]. Available: https://fedimint.org/docs/GettingStarted/TechCompontents

[74] M. E. Peck. "The Crazy Security Behind the Birth of Zcash, the Inside Story." IEEE Spectrum, Dec. 2, 2016. [Online]. Available: https://spectrum.ieee.org/the-crazy-security-behind-the-birth-of-zcash

[75] Ethereum Community. "KZG Summoning Ceremony." [Online]. Available: https://ceremony.ethereum.org/

Appendix 3: Techniques for addressing challenges in pseudonymization

One challenge with pseudonymization is enforcing properties across PIPs, such as overall user holding limits (see Section 3.1: Limits of Pseudonymization for discussion). This appendix offers a few techniques that might be helpful to address such limits in a system design.

Probabilistic auditing. A user's transaction could be selected at random for a transaction volume check. When selected,[76] a user's wallet could interact with all of that user's PIPs and produce a proof that daily transactions for a particular day are indeed below the mandated threshold. Such enforcement could probabilistically forgive an occasional excess above a limit, but could detect users who routinely exceed daily volume limits.

Delayed linkage. The tallying step to determine if a user is over a global holding limit could be done once every enforcement period rather than once for every transaction. For example, to verify that daily limits are honored, intermediaries could prepare a daily report that lists each user's stable identifier and the user's total daily CBDC flows. When centrally aggregated, such reports could reveal users whose total payments for that day are above the daily limit. To limit the privacy impact of such aggregation, one could employ a number of other privacy-enhancing techniques.[77]

Introducing pseudorandom identifiers. Using cryptographic techniques, one can replace stable pseudonym identifiers with verifiably pseudorandom values that incorporate both a stable identity and an auditing period number. This way, the same user identity yields a different pseudorandom identifier when considered for a different auditing period. In more detail, when a user with a stable identity "***id***" sends a transaction on a day "***d,***" the wallet could compute an auxiliary hash value of "***h = VRFsk (id, d)***" and include this string in the transaction. Such hash values "***h***" could be the same per-user, per-day, regardless of the PIP used, but could be different (and unlinkable) between different days. In such a hypothetical design, the central transaction processor could learn that all transactions derived from the same stable identity on the same day share that stable identity.[78] It could enforce per-user, per-day limits by tallying together transactions that share the same pseudorandom identifier "***h.***"

[76] Probabilistic enforcement raises questions about deliberately *selective* enforcement, e.g., what if a rogue actor selects a certain demographic for increased scrutiny, instead of selecting one in 10,000 transactions at random? Such scenarios can be prevented through the use of cryptographic techniques. For example, audit decisions can be made verifiably pseudo-random (yet unpredictable to the transaction originator). Furthermore, each decision of a selection or non-selection can carry a receipt that users can later aggregate to look for selection anomalies.

[77] To guard against transaction-identity linkage by value, a flow could be reported in increments of £10. To guard against tracking users across different auditing periods, one could replace stable identifiers with auditing-period-dependent pseudorandom identities as discussed above. Additionally, instead of centrally aggregating such reports, we could use secure multiparty computation among the PIPs.

[78] It has to be noted that such additional linkage can be harmful to privacy (e.g., by signaling that two relationship pseudonyms are the same) and must be managed in an appropriate way. For example, the pseudorandom identifiers should be deleted after their usage period has passed (e.g., every night) and would ideally be checked only by a dedicated part of the system that does not get access to other transaction details besides the value; in particular, that system component does not see pseudonyms.