DOI:10.1145/3579316

Article development led by acmqueue
queue.acm.org

Now is the time to shape what future payment flows will reveal about you.

BY RAPHAEL AUER, RAINER BÖHME, JEREMY CLARK, AND DIDEM DEMIRAG

Mapping the Privacy Landscape for Central Bank Digital Currencies

PAYMENT RECORDS PAINT a detailed picture of an individual's behavior. They reveal wealth, health, and interests, but individuals do not want the burden of deciding which are sensitive or private.[1] Central banks are exploring options to digitize cash. As of January 2023, 27 of the 38 member states of the Organization for Economic Cooperation and Development (OECD) have announced retail central bank digital currency (CBDC) research and projects.[a]

The issue of privacy needs to move center stage. Decades of work on privacy-enhancing technologies have highlighted that privacy does not come for free, it is easy to get wrong, and it is imperative to design before deployment.

CBDC has been discussed in policy reports, academic papers, and public media through lenses such as monetary policy,[6] impact on the financial system,[2] and technology.[3] Almost all of these documents flag the importance of privacy, but many lack in-depth discussion or concrete design choices. Figure 1 shows the uptake of privacy in the CBDC literature: While the question is raised, significant treatment is still rare. An exception is recent academic papers (shown in the top right corner of the figure), which are generally written by computer scientists. These papers offer specific solutions to include in the privacy design landscape.

Policymakers may shy away from papers with cryptographic equations that mention Alice and Bob. While there are exceptions,[9] the gap in concrete privacy solutions in policy reports is puzzling, as economists have argued that CBDC could make an essential difference in providing privacy in digital payments.[10] It is popular for authors of these reports to point out the tension between privacy and law enforcement; reiterate that it requires a solution; and ultimately punt to government officials, legislators, the judiciary, or public opinion to solve it. Occasionally, technical solutions are prescribed (for example, blockchains, cryptography, zero-knowledge proofs) without adequate operational details or even precision about exactly what data is protected from whom. The number of distinct stakeholders, combined with the technical challenges, has stalled progress toward deploying retail CBDC.

One step forward is understanding who the key stakeholders are and what their interests are in payment records. Knowledge of conflicting interests is helpful for developing requirements and narrowing the range of technical solutions. This article contributes to the literature by identifying three stakeholder groups—privacy-conscious users, data holders, and law enforcement—and exploring their conflicts at a high level.

A main insight is that nuanced da-

[a] See the January 2023 dataset update at https://www.bis.org/publ/work880.htm

IMAGE BY ANDRIJ BORYS ASSOCIATES, USING SHUTTERSTOCK

CBDC
DECENTRALIZED DIGITAL CURRENCY · BITCOIN

Figure 1. Uptake of privacy in the CBDC literature.

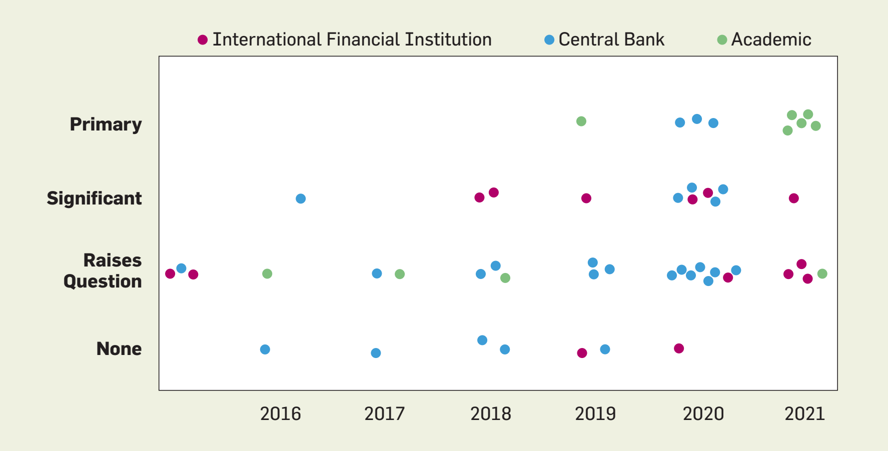

ta-access policies are best to resolve the conflicts, which in turn rule out many technical solutions that promise "hard privacy," meaning solutions relying on cryptography and user-guarded secrets without room for human discretion.[7] This observation shifts attention to a softer form of privacy-enhancing technologies, which gives authorized stakeholders the capability to access certain payment records in plaintext under defined circumstances. Such a system depends on compliance and accountability, supported with technically enforced access control, limited retention periods, and audits. This is referred to as "soft privacy."

Figure 2. Main conflicts between stakeholder categories.

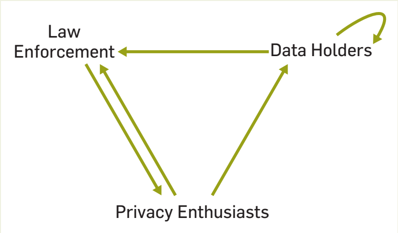

Stakeholder Analysis

Many countries exhibit a convoluted set of payment options. These have evolved from stakeholders having competing interests, including tussles over privacy. Stakeholders include the users (who might additionally be vulnerable, unbanked, undocumented, children, foreign residents, or tourists), merchants, banks and payment providers, government (central banks, financial regulators, law enforcement, and intelligence agencies), and other parties with an interest in the tension between privacy and transparency (investigative journalists and privacy advocates).

A detailed stakeholder analysis of all these parties proved unnecessary in that the key tensions are well-captured through consideration of only three stakeholder categories:

- ***Privacy enthusiasts.*** Users of a payment system with an interest in privacy.
- ***Law enforcement.*** Investigators of crimes with financial evidence.
- ***Data holders.*** Entities that record and monetize financial data, including merchants, banks, and payment processors.

Figure 2 illustrates the conflicting relationships identified among these stakeholder categories. The accompanying table shows an evaluation of payment options from the perspective of stakeholder conflicts. The table compares payment options by how well they deal with the identified conflicts using a simple ordinal scale: good, OK, and bad.

Let's start with the relationship between privacy enthusiasts and law enforcement. There is a subtype of privacy enthusiast who is law-abiding and affirms that crimes can be deterred with effective law enforcement, yet believes that errors, corruption, breaches, and overreach are present or potential future concerns. Law enforcement prefers the least amount of friction in obtaining payment information that is pertinent to their investigations.

While the privacy preferences of these two stakeholders might appear diametrically opposed, this is the case only when everyone is "treated like a criminal." Hypothetically, if criminal activity could be perfectly discriminated from benign transactions, and benign transaction data was protected unconditionally, both stakeholders would be satisfied. This is probably impossible, but privacy-affirming payment systems try to approximate it.

A common research direction, pursued in at least nine CBDC projects, is to offer privacy for transactions under a certain threshold (for example, $10,000). This is much too rigid. Investigators might not care about a $20 payment at a gas station when investigating tax evasion, but it is critical information when the payment is made by an abduction suspect on the run.

A privacy enthusiast would consider cash as largely addressing their concerns with law enforcement (denoted by P→L in the table), while law enforce-

Evaluation of payments option.

	Law enforcement	Privacy enthusiasts		Data holders	
	L → P	P → L	P → D	D → D_{new}	D → L
Cash	ok	good	good	bad	bad
Payment Network	good	bad	bad	good	ok
Crypocurrency	bad	good	good	ok	bad
Soft Privacy CBDC	good	ok	good	ok	ok
Hard Privacy CBDC	bad	good	good	bad	bad

ment would be concerned about the wide usage of cash by privacy enthusiasts (denoted by L→P). Law enforcement is not helpless at tracing cash, however, aided by serial numbers, marked bills, fingerprints, reporting of large cash transactions at regulated businesses, ATM surveillance, and the high carrying costs of transporting and protecting large holdings in low-denomination bills. Conversely, a payment network (next row in the table) recording every transaction and providing exceptional access to law enforcement (with judicial oversight) resolves their preference while leaving a privacy enthusiast concerned about its potential for abuse or breach.

Cryptocurrencies use hard privacy, where user-generated public keys define accounts, and the corresponding private keys authorize transactions using digital signatures. The design intention is to mirror the privacy provisions of cash, although success varies. Users today can choose among many variants with differing levels of anonymity (X receives $100 from someone), pseudonymity (X receives $100 from Y), confidential transactions (Alice receives $Z from Bob), or combinations thereof. The table assumes an anonymous and confidential cryptocurrency.

While cryptocurrency is sometimes considered a steppingstone for CBDC, important technical differences prevail, including with privacy. The choice exists between hard and soft privacy for CBDC. Soft-privacy CBDC uses judicial oversight to allow human discretion in balancing exceptional access to payment data with privacy, much as it is done today in payment systems. In contrast, hard-privacy CBDC eliminates human intervention by relying solely on cryptography and, perhaps, tamper-resistant hardware. With current technology, it is possible to conceive hard-privacy CBDC that is more difficult for law enforcement to trace than cash.

While the tensions between privacy enthusiasts and law enforcement get top billing in the CBCD literature, the less obvious tension between privacy enthusiasts and data holders is equally important (P→D). Payment data is personal data. It can be monetized. It is useful for profiling users through analysis techniques that improve and become cheaper over time. The difficulties in tracing cash are safeguarded, while payment systems enshrine banks and payment processors with valuable, proprietary data that could be useful for decades.[10]

Law enforcement prefers the least amount of friction in obtaining payment information that is pertinent to their investigations.

The development of soft-privacy CBDC will have to contend with a greater variety of stakeholders and a lack of transparency about how payment data can be used, attempting to avoid the thicket of rules governing today's payment systems. One option for CBDC is to shift all the payment data from today's commercial entities to the central bank, which ideally has no incentive to monetize it and can shield it from abuse for political interests. CBDC is also an opportunity to design a seamless set of soft-privacy rules from scratch.

Data holders and law enforcement have minimal conflict. If law enforcement can obtain payment information, it is not generally concerned with who it is from (for example, a commercial bank instead of a central bank). Conversely, data holders find rules for identity gathering and reporting transactions expensive and onerous (D→L) and would favor their relaxation or even elimination. Systems with a degree of traceability can lead to lighter regulation, while hard-to-track payment methods, such as cash or cryptocurrency, lead to more regulation.

In the CBDC literature, some solutions propose on-boarding users with cryptographically protected identities that can be used for selective traceability by law enforcement.[14] While better than strict anonymity for law enforcement, these systems impose greater costs on commercial banks with additional computation, procedures, and internal controls relating to the involved cryptography.

Data holders also have a tension among themselves, as any change to the payment system might increase or decrease their access to future payment data. For example, CBDC that runs on a permissioned blockchain might benefit the incumbents by creating a higher barrier to entry for newcomers. A radical change to the payment infrastructure would be direct CBDC run by a central bank where commercial banks and payment providers have no access to payment data. Commercial banks would still play a role, such as on-boarding users in a compliant way, and they are likely to continue offering traditional payment systems, perhaps with financial incentives proportional to the utility they can harvest from the payment data.

The Case for Soft Privacy

Several insights can be extracted from the stakeholder analysis. First, while payment records contain sensitive data and must be protected by default, access to payment records is vital to prevent crime in certain cases. The potential to trace money flows can deter bad actors from committing crimes in the first place. A deterrence theory of crime speaks against hard-privacy solutions where users can be assured that payment flows remain protected unconditionally, perhaps even without having to trust in much else than the integrity of one's own device and the secrecy of keys.[7]

A second insight is that no formal (that is, machine-decidable) access policy can cater to all concerns. This forces a reconsideration of middle-ground cryptographic techniques that provide conditional privacy: for example, for a defined number of payments,[4] payments of low value,[14] or payments to vetted payees.[8,15] The main problem with all formal access policies is that the conditions for revoking privacy have to be defined at the time that the data is encrypted or anonymized. This means bad actors can anticipate the rules and adapt their behavior to evade prosecution. An example is a smurfing attack, where a large payment is split into many small ones, each below the threshold used in the condition to revoke privacy. Future research is likely to provide efficient systems capable of evaluating ever-more-complex conditions and finer-grain privacy revocations, but complexity does not necessarily address the fundamental challenge. At worst, it can generate new problems.

The history of cryptography is littered with failure. Lessons such as peer review, formal analysis, and the test of time had to be learned the hard way. For example, the cornerstone of network security today is Transport Layer Security (TLS), which has endured almost 30 years of protocol flaws and implementation issues; even the latest version includes elements for which no security proof is known. TLS solves a relatively simple task of securing communication with a well-identified server using basic cryptography primitives such as encryption, hashing, and creating digital signatures. In contrast, the cryptographic building blocks necessary for CBDCs with conditional privacy are much newer, more involved, and less reviewed. Moreover, few of these techniques scale well to zillions of real-time payments.

Figure 3. Moving from primary to secondary data uses.

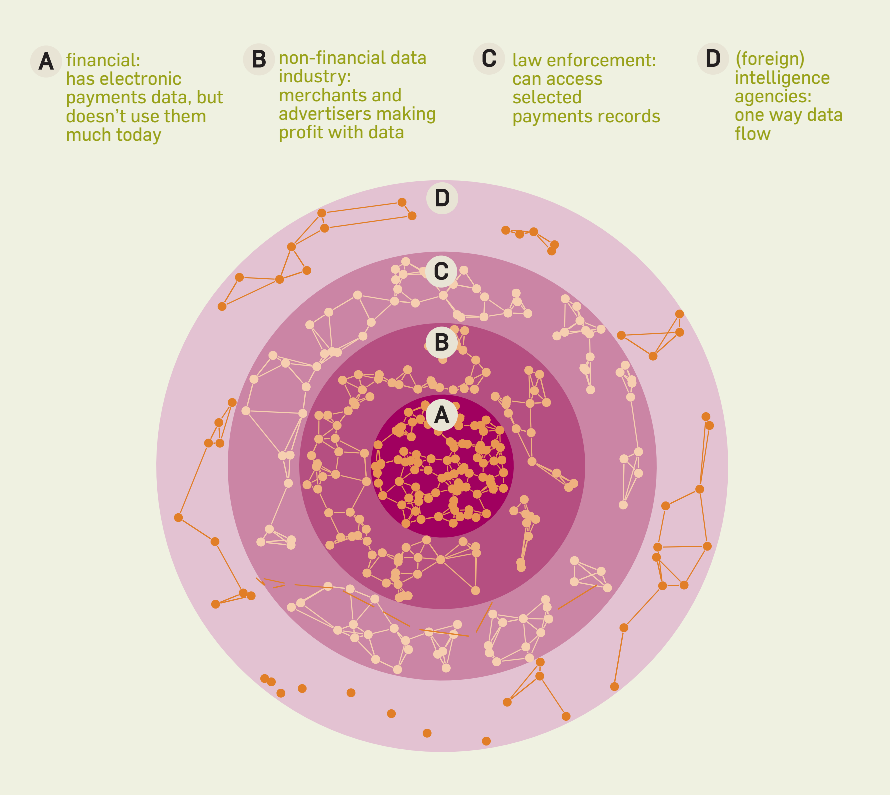
A
financial:
has electronic payments data, but doesn't use them much today
B
non-financial data industry:
merchants and advertisers making profit with data
C
law enforcement:
can access selected payments records
D
(foreign) intelligence agencies:
one way data flow
D
C
B
A

All these observations lead to the conclusion that CBDC systems should be designed in a way that protect payment records in bulk (for example, against devastating data breaches) but with plaintext access possible in justified cases. The conditions for plaintext access should be rooted in appropriate law, which by design leaves room for discretion. This is desirable—respecting checks and balances—as it allows the law to evolve and adapt to new situations while preserving its intended spirit.

As unauthorized plaintext access does not leave traces and cannot be reversed, this approach does require some trust in privileged parties, who are committed to compliance, implement technical safeguards against errors and outside attacks, and are subject to oversight and regular audits. On balance, soft privacy appears to be the most practical way forward.

Primary to Secondary Data Uses

If soft privacy is accepted as the most practical approach, the main questions to ask are: Which stakeholders are easiest to trust? How can effective checks and balances ensure that they deserve trust?

The stakeholder analysis reduced dozens of stakeholders to three with the intention of mapping conflicts. The same trick can be used in answering these questions. It is convenient to arrange the stakeholder groups in a diagram of concentric circles as shown in Figure 3, with information needs close to the payment process at the center. Secondary data uses are on the outside. The organizing principle reflects the information flows in the payment process: Inner rings have and require more information to make payments work; outer rings are less relevant for the payment process itself but might want the information for purposes *other* than payments.

Starting from the center, at ring A, today's banks and payment processors must have detailed data to make payments happen, but not all of them use it as much as they could. This data has secondary uses, such as credit scoring. The extent to which this happens varies among payment systems and regions, with the tendency to explore more secondary uses through an emerging financial technology sector.[5]

Ring B hosts the data industry outside of the financial sector. It consists primarily of merchants, who seek to commercialize the secondary uses of payment data. This ring also includes technology suppliers who seek to expand tracking and targeting across merchants and industries, and increasingly, the offline world. This is where payment data is deemed particularly valuable. Not only are payment records reliable indicators of economic activity and consumer choice, but they are also often linked to persistent identities up to real names and street addresses.

Law enforcement on ring C is involved in payments only indirectly. It can strengthen trust in CBDC by tracking stolen funds, although only a small fraction of payments is disputed. While the former is a primary use of the data, law enforcement is also interested in secondary uses. It believes that bulk access to records—disputed and undisputed alike—helps solve all kinds of crimes ranging from tax evasion to uncovering criminal networks. These are secondary uses of the data because the security of a payment system is not increased by taxes being enforced within it.

On ring D—the outermost considered here—intelligence agencies are added to the picture. A discussion of CBDC privacy cannot ignore this stakeholder. It is often brought forward as an argument for hard privacy. These organizations strive to capture data of all types, with payment data offering crucial links to real-world entities. Intelligence agencies are set up as one-way entities: They are data sinks without regular or substantial feedback to the systems they observe. As all major intelligence agencies operate across borders, it is particularly hard to ensure that soft-privacy rules are not circumvented by domestic or foreign actors.

Organizational safeguards should include transparency in the system design and oversight by an independent body equipped with resources and expertise to verify the integrity of the system.

Implications

Here, we look at privacy architectures along the diagram in Figure 3. For completeness, it briefly considers hard privacy as well.

Option 1: Hard privacy. Cypherpunks would argue for placing hard privacy everywhere, meaning that plaintext access would be reserved for end users with access to private keys only. All stakeholders in the diagram see and learn nothing more than what is necessary for the payment process.

Such a system would suffer from technical overhead on ring A because encrypted records tend to require more space. The anonymization on the network layer that all hard-privacy proposals assume adds latency and other inefficiencies because of repeated encryption, decryption, and reencryption. The benefit for the cypherpunks is excluding everyone but ring A from access. This may succeed for ring B. Ring D, in turn, will still try to subvert endpoints and devote more effort to traffic analysis and code breaking. Crucially, for ring C, law enforcement is both deprived of the capabilities needed to solve crimes and does not have the budget of an intelligence agency.

Here it is important to consider possible knock-on effects. Every architecture with hard privacy at its core will push enforcement actions to the on/off-ramps of the system, increasing the burden of record-keeping and reporting for rings A and B, as seen with cash. A shadow system of record-keeping mirroring almost all the activity in the payment system proper could result in less privacy because of poorer security and governance. For example, several US regulators routinely receive a complete picture of all cryptocurrency transactions that fall under a "travel rule." Yet, criminals find legal loopholes or technical ways to escape prosecution. This creates an inefficient dragnet that neither meaningfully prevents crime nor protects the privacy of legitimate users.

Option 2: Soft privacy. The opposite would be soft privacy everywhere. Payment data would freely flow among parties, still protected with point-to-point encryption against outside attackers. Every party involved is identified and trusted to adhere to the privacy policy. Regular audits and the threat of sanctions encourage disciplined processing.

Such a system could be very efficient. Its privacy, however, would not be any different from that offered by current payment networks, and hence worse than paying in cash. Gaining the trust of privacy enthusiasts would therefore be difficult, chiefly for the weak guarantees it offers to discourage unwanted secondary use of payment data. Another threat to privacy is that data holders on rings A and B might accept the risk of sanctions, and even bend the interpretation of law in their interests (compare with web tracking today[11]).

Moreover, preventing abuse of stolen data (for example, after data breaches) is technically impossible and a real threat given the frequency of breaches.[13] Similar concerns apply to rings C and D, which depend on the competence of law enforcement to establish data security, as well as citizens' trust in internal controls. This amount of blind trust is not ideal; hence, a better compromise is needed.

Option 3: A soft core with a hard shell. Some entities in ring A need plaintext access for efficiency; hence, soft privacy is implemented here. To respect the principle of data minimization, no other rings are given access by default. To serve justified data requests, however, particularly from ring C, ring A acts as a data custodian: It may grant plaintext access to selected records while it is ensured that all requests are authentic, justified, and proportionate, and leave an audit trail. Here, you could experiment with hard-privacy techniques to govern this data transfer. Existing techniques such as tokenization, used in current payment networks to shield customer data from merchants, can be a source of inspiration. Likewise, anonymized aggregate data could be made available to ring B, again under transparent rules.

This approach seeks to balance hard and soft privacy. Note that plaintext access on ring A does not mean data is stored or transferred in plaintext. Rather, encrypted data can be decrypted when necessary. As an additional technical safeguard, ring A could enforce strict data retention limits by using short-lived cryptographic keys and deleting them after a defined amount of time—for example, six months for transaction records and 10 years for aggregates such as account balances. The choice of these periods reflects the elevated sensitivity of detailed payment records and their metadata, as compared with account-holder information that financial intermediaries must retain today to combat financial crimes.

Preventing abuse of stolen data is technically impossible and a real threat given the frequency of breaches.

Organizational safeguards should include transparency in the system design (for example, open source, multiple vendors) and oversight by an independent body equipped with resources and expertise to verify the integrity of the system.

This leaves us with the question of which parties should run ring A, the core of the system? An ideal institution would operate as a not-for-profit organization with a clear public-interest mandate. Many countries set up their central banks as institutions that match these goals. Moreover, CBDC is a claim against the central bank's balance sheet; hence, it derives all value from the integrity of the central bank. Citizens must (and do[12]) trust their central bank anyway when using money. This suggests that from the perspective of privacy, for countries in which these conditions are immutable, ring A could be run by the central bank.

There is a broad body of literature on designing CBDC (recall Figure 1). A "direct architecture," where the central bank operates the record-keeping, is one of the options presented but never as the favorite model, chiefly because it would crowd out commercial banks from the payment system. The literature has never considered privacy from scratch, however. From the perspective of privacy, there are stakeholder-based arguments for a direct architecture. Therefore, adding privacy casts new light on the discussion of the technical architecture for CBDC.[3]

Concluding Remarks

The privacy landscape of CBDC is more complex than often appreciated. Mapping it through two perspectives—stakeholder conflicts and stakeholder proximity to the data—has not been addressed in the CBDC literature before. Both perspectives attempt to simplify as much as possible, leading to consideration of, respectively, three and four representative stakeholder groups. Future work might consider if this is too simple and leaves things out, while a more complete view might be better.

The process related in this article has led to new insights. A glaring gap in the literature up to this point is the lack of realistic privacy definitions for money, which hinders the research and development of hard-privacy solutions that address all realistic concerns when trading off between privacy and crime-fighting. Some proposals place an anti-money laundering component in the message path that can be skipped for private transactions, magically assuming that this component is all that is needed to stop financial crime. Others are designed to prevent tax evasion but cannot resolve when the taxed proceeds result from felonies such as human trafficking.

Realistic crime-fighting uses a plethora of methods, each of which conflicts with privacy goals to a varying extent. CBDC that balances privacy protections against a single method risks criminals changing behavior to evade prosecution. Proponents of hard privacy need to move past whiteboard models of crime-fighting with operable privacy definitions; without that, soft privacy at the core is the most tenable way forward.

Another challenge for operable privacy is that too much privacy can fire back. Consider tightening law enforcement's direct access to payment data. An unintended consequence might be new regulations that require increased logging and reporting of transaction details outside of the payment system. This new shadow data system might enjoy less scrutiny and public interest; hence, it becomes less secure and less transparent than building access into the core CBDC in the first place. This hints at an apparent paradox—that there are circumstances where increasing access to data can increase privacy—that researchers might consider in domains beyond CBDC.

The analysis presented in this article suggests that, from the standpoint of privacy, it could be argued that CBDC should be run by a public-sector agency. For nontechnical reasons, the literature advises against a "direct" CBDC architecture, also to maintain a status quo that in part still derives from the technology used in the 1970s. Central banks are institutions run by economists and lawyers; they do not become tech companies overnight.

Moreover, the fear of disrupting the monetary and financial system (rings A and B in Figure 3) as a side effect of changing payment technology is certainly a valid argument for keeping commercial banks in the loop. That said, from the perspective of privacy, there are benefits to limiting the number of parties with default access to detailed records.

One substantial concern with a direct CBDC architecture, however, is that it would imply the central bank hosting a significant amount of sensitive retail transaction data. Other government agencies, such as tax authorities, have a long history of dealing with the risk of data breaches and the possibility of being called as witnesses in law enforcement cases. This, however, would be a new role for central banks, one that is potentially beyond their mandates. This role could therefore also be delegated to a new type of public authority tasked with protecting retail payment data. (Commercial banks would still play a role, such as on-boarding CBDC users in a compliant way, and they would likely continue offering traditional payment systems, perhaps with financial incentives proportional to the utility they can harvest from the payment data.)

Regardless of which agency hosts the data, state-of-the-art technology must be used to reduce the likelihood and severity of data breaches. Researchers who argue that privacy is to be put first should be bold and consider technical designs and operational architectures with such reshuffled divisions of responsibility.

Acknowledgments

We thank R. Arora (Bank of Canada) for useful feedback. Clark and Demirag thank the Office of the Privacy Commissioner of Canada's Contributions Program for funding their work on this project. Böhme thanks the Anniversary Fund of the Oesterreichische Nationalbank for supporting his research project on "Privacy and the Functions of Digital Money."

References

1. Acquisti, A., Brandimarte, L., Hancock, J. How privacy's past may shape its future. *Science 375*, 6578 (2022), 270–272; https://bit.ly/3WGkkIe.
2. Andolfatto, D. Assessing the impact of central bank digital currency on private banks. *The Economic J. 131*, 634 (2021), 525–540; https://academic.oup.com/ej/article-abstract/131/634/525/5900973.
3. Auer, R., Böhme, R. The technology of retail central bank digital currency. *BIS Quarterly Rev.* (2020), 85–100; https://www.bis.org/publ/qtrpdf/r_qt2003j.htm.
4. European Central Bank. Exploring anonymity in central bank digital currencies. In Focus 4 (2019); https://bit.ly/3TlURRx
5. Boissay, F., Ehlers, T., Gambacorta, L., Shin, H.S. Big techs in finance: on the new nexus between data privacy and competition. BIS Working Papers 970, 2021; https://www.bis.org/publ/work970.htm.
6. Bordo, M.D., Levin, A.T. Central bank digital currency and the future of monetary policy. National Bureau of Economic Research. Working Paper 23711 (2017); https://www.nber.org/papers/w23711.
7. Chaum, D. Security without identification: Transaction systems to make Big Brother obsolete. *Commun. ACM 28*, 10 (Oct. 1985), 1030–1044; https://dl.acm.org/doi/10.1145/4372.4373.
8. Chaum, D., Grothoff, C., Moser, T. How to issue a central bank digital currency. 2021, arXiv 2103.00254; https://arxiv.org/abs/2103.00254.
9. Darbha, S., Arora, R. Privacy in CBDC technology. Bank of Canada. Staff Analytical Note 2020-9, 2020; http://bit.ly/3UpFPvw.
10. Garratt, R. Lee, M.J. Monetizing privacy. Federal Reserve Bank of New York. Staff Report 958, 2021; https://nyfed.org/3UIWQAt.
11. Matte, C., Bielova N., Santos, C. Do cookie banners respect my choice? Measuring legal compliance of banners from IAB Europe's Transparency and Consent Framework. In *Proceedings of IEEE Symp. Security and Privacy*, 2020, 791–809; https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9152617.
12. Official Monetary and Financial Institutions Forum. Digital currencies: a question of trust, 2020; https://bit.ly/3G0TOTY.
13. Wheatley, S., Maillart, T., Sornette, D. The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B 89*, 7 (2016); https://link.springer.com/article/10.1140/epjb/e2015-60754-4.
14. Wüst, K., Kostiainen, K., Capkun, V., Capkun, S. PRCash: Fast, private and regulated transactions for digital currencies. *Financial Cryptography and Data Security, LNCS 11598.* I. Goldberg and T. Moore, eds. Springer, 2019, 158–178; https://link.springer.com/content/pdf/bfm%3A978-3-030-32101-7%2F1.pdf.
15. Wüst, K., Kostiainen, K., Capkun, S. Platypus: A central bank digital currency with unlinkable transactions and privacy preserving regulation. ACM CCS 2022; https:dl.acm.ord/doi/abs/10.1145/3548606.3560617.

The views expressed here are those of the authors and not necessarily those of the BIS Innovation Hub.

Raphael Auer is head of the Eurosystem Centre of the BIS Innovation Hub, which has offices in Frankfurt and Paris and explores technologies to improve the functioning of the global financial system. Twitter @RaphAuer

Rainer Böhme is a professor of computer science at the University of Innsbruck, Austria, where he focuses on solving societal problems of technology policy, security, and privacy.

Jeremy Clark is an associate professor at the Concordia Institute for Information Systems Engineering in Montreal, Canada, where he holds the NSERC/Raymond Chabot Grant Thornton/Catallaxy Industrial Research Chair in Blockchain Technologies. Twitter @PulpSpy

Didem Demirag is a postdoctoral researcher at the University of Quebec in Montreal, Canada.