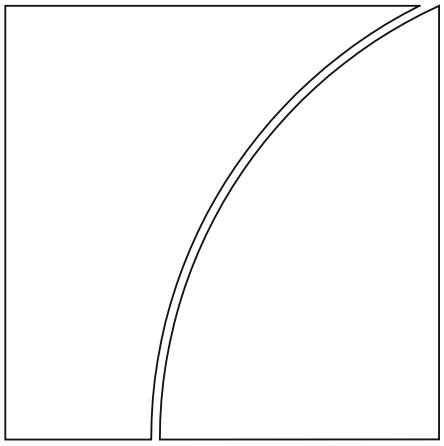

BIS Working Papers

No 1242

Privacy-enhancing technologies for digital payments: mapping the landscape

by Raphael Auer, Rainer Böhme, Jeremy Clark and Didem Demirag

Monetary and Economic Department

January 2025

JEL classification: E42, G23, G28, O32.

Keywords: Privacy, privacy-enhancing technology, payments, BigTech, fintech, regulation, smart contracts, zero-knowledge proofs, applied cryptography, digital money, digital currency, stablecoins.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Privacy-Enhancing Technologies for Digital Payments: Mapping the Landscape*

Raphael Auer† Rainer Böhme‡ Jeremy Clark§ Didem Demirag¶

23 January 2025

Abstract

How can technology enhance privacy in digital payment systems? This paper presents a systematic evaluation of the interests of privacy-conscious users, commercial data holders, and law enforcement. We classify privacy-enhancing designs along the dimensions of privacy versus auditability, as well as soft institution-based versus hard technology-based solutions. We map existing technologies into this taxonomy and assess them. Sophisticated techniques allow having both hard privacy and limited transparency by employing hard-coded rules that dictate which data remains inaccessible. On balance, there is promise in novel concepts like modern zero-knowledge-proofs, but current technologies also suffer from limitations in terms of security and computational capacity. More technological development is needed in this area. Additionally, efforts could focus on technological development that augments such hard privacy with technologically-enforced access control and systems minimizing the amount of data that is being stored, render abuse transparent and make data holders accountable.

Keywords: Privacy, Privacy-enhancing technology, Payments, BigTech, Fintech, Regulation, Smart Contracts, Zero-Knowledge Proofs, Applied Cryptography, Digital Money, Digital Currencies, Stablecoins

JEL: E42, G23, G28, O32

*Full version. A shorter version of this article has appeared in *Communications of the ACM* (via *ACM Queue*), see [Auer et al., 2023]. The authors thank the editor, Jon Frost, Martin Hood, Fabian Schär, Andras Valko, and William Zhang for comments and suggestions. Clark and Demirag acknowledge funding from the Office of the Privacy Commissioner of Canada's Contributions Program, Autorité des Marchés Financiers (AMF), and NSERC. Böhme acknowledges funding from the Anniversary Fund of the Oesterreichische Nationalbank for a project on "Privacy and the Functions of Digital Money." The views expressed here are those of the authors and not necessarily those of the Bank for International Settlements.

†Bank for International Settlements

‡University of Innsbruck

§Concordia University

¶Université du Québec à Montréal

1 Taking back privacy in digital payments

Payment records contain highly sensitive information. With the widespread adoption of digital payments and the proliferation of smartphones and data-driven technologies, users' personal information — including wealth, interests, and other sensitive data — is at risk of being collected, analyzed, and used for price discrimination, advertising, and other purposes.[1]

Without advocating for a specific solution, we argue that the issue of privacy for digital money needs to move center stage in the public policy discussion. Economists have long argued that the privacy that physical cash offers is valuable to society (see [Kahn et al., 2005]). And privacy is more than just a convenience. This characteristic is important, not necessarily for illicit purposes, but for ensuring freedom from data abuse and helping to maintain the boundaries of private lives.

Decades of work on privacy-enhancing technologies have highlighted that privacy does not come for free in electronic systems; it is easy to get wrong, and it is imperative to design and test privacy protections before deployment. Across a wide range of public and private institutions, data breaches are becoming increasingly widespread and costly. Emerging technologies such as artificial intelligence and quantum computing may soon post novel risks to cryptographic encryptions and data integrity.[2] In the area of smartphones, individuals often fail to optimally protect their privacy because long-winded terms and conditions make it impossible to determine which data are sensitive or private (see [Acquisti et al., 2022]).[3]

Well-designed payment systems present an opportunity to enhance consumer welfare by offering a level of digital privacy that currently does not exist (see e.g. [Agur et al., 2022, Tinn, 2024]). One area where the discourse on privacy has intensified is retail central bank digital currencies (CBDCs).[4] For example, initiatives like the GNU Taler and Project Tourbillon have showcased potential technological solutions such as payer anonymity.[5] Project Polaris ([BIS Innovation Hub, 2023a]), on the other hand, focuses on offline functionality; since fully offline device-based funds can be exchanged physically without a digital trace, such designs also can enhance privacy. Learnings from these experiments could also be considered by the private sector to improve consumer privacy.

Privacy-enhancing technologies should remain a significant consideration for all digital payment options in the private sector. Traditionally, confidentiality is one of the key services offered by the financial industry. However, the use of payments data in the financial industry is evolving (see, e.g., [Visa, 2024]) and regulators are discussing the need for robust data protection measures, user rights, and transparency to safeguard personal data within financial services (see e.g. [European Data Protection Supervisor, 2023, UK Payment Systems Regulator, 2018]).

In this paper, we thus examine the role that technology could play to enhance privacy in payments (see also [Darbha and Arora, 2020]). To have a meaningful discussion of the technology that novel payment systems could run on, it is necessary to analyze the underlying trade-offs. To this end, we start by assessing existing privacy-enhancing technologies from the perspective of different stakeholders, such as privacy-conscious users, commercial data holders, and law enforcement.

The underlying tension is between privacy goals, law enforcement (auditability goals), and technical performance goals. Figure 1 presents a schematic depiction of the conflicting objectives encountered when engineering privacy-enhanced digital payment systems. Many conventional technologies may scale to good

[1] In this context, the forays of fintech and big tech companies into finance, given their vast access to data and ability to process it, have come under scrutiny for the handling of user data, including financial information, see e.g. [Frost et al., 2020].

[2] see [Auer et al., 2024, BIS, 2023] for a discussion of the threat of quantum computing to data confidentiality and potential solutions.

[3] As another case in point, even cryptocurrencies specifically designed to provide privacy have been shown to be partly traceable (see [Möser et al., 2018] and [Kappos et al., 2018]) or have displayed technological vulnerabilities (see [Ruffing et al., 2018]).

[4] A retail CBDC is a cash-like direct claim on the central bank's balance sheet available to the general public. See [Auer and Böhme, 2021] for a discussion of the underlying technology.

[5] See Project Tourbillon ([BIS Innovation Hub, 2023b]) for an implementation of payer anonymity, as well as [Chaum et al., 2021, Chaum and Moser, 2022]. See also [Gross et al., 2021] for an alternative approach.

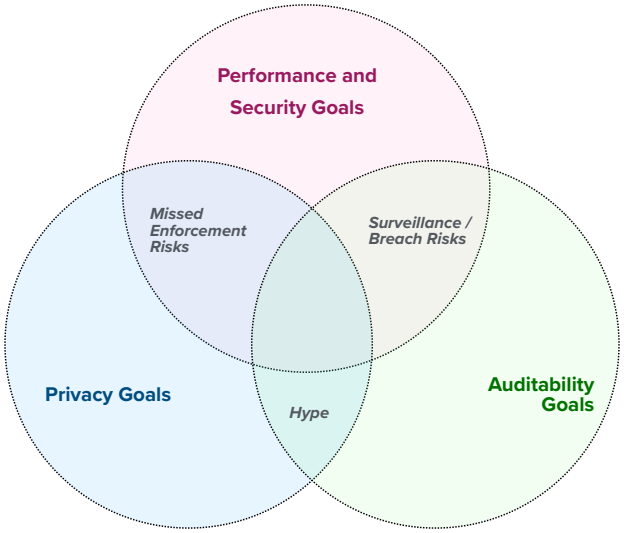

Figure 1: Objectives when engineering privacy-enhanced digital payments

performance and be easily auditable due to a lack of privacy protections, but these open the door to surveillance and data breach risks. On the other end of the spectrum, a complete lack of auditability may mean that there are missed opportunities to fight crime. Lastly, as we discuss below, many privacy-enhancing technologies satisfy both privacy and auditability goals, but they do not scale to the computing needs of a large-scale retail payment system that may need to process thousands of transactions per second.

We next define what it means to guard data against unauthorized access and to provide authorized access to it, and map out a taxonomy of the technological and institutional means by which this can be achieved. In academic and policy discussions, privacy is often simplistically portrayed as a trade-off along a single spectrum between privacy and transparency. This uni-dimensional thinking misses sophisticated cryptographic techniques that allow having both privacy and transparency, by giving access to some payment records under hard-coded rules that cannot be circumvented. For example, this could mean that the payer always remains anonymous, or that transactions are anonymous except if the owner of the funds authorizes partial use of the data (see [Chaum et al., 2021, Chaum and Moser, 2022, Buterin et al., 2023]). We hence start by classifying privacy-enhancing designs into two categories: privacy versus auditability and institution-based versus technology-based solutions.

Thereafter, we describe and evaluate current technological proposals for the design of privacy in digital payments. We first investigate "hard privacy" solutions, i.e., those relying exclusively on cryptography and user-guarded secrets without room for human discretion. We find that some of today's technical solutions that promise such hard privacy suffer from technical performance issues that make it infeasible to operate large-scale payments systems using these technologies. Further, even once the technology is ready, the set of rules might not be. There are various types of crime with different types of payments, and criminals can adapt to rules.[6] This observation shifts our attention to the need for further development in the area of hard privacy.

Technology also has an important role to play for soft, institution-based, privacy. For soft privacy, data protection policies and technologies can curb access to payments data, authorizing only entities that require

[6]Modern cryptography can essentially implement any type of private computations, as well as establishing any type of rule that governs what data should be minimally disclosed and what data should be protected unconditionally. The main challenge here is not always the technical one; often, it is a policy one. For example, rules that provide auditability only when transactions exceed a certain limit are too simplistic. Another challenge is timing. The rules must always be set before the transaction takes place, sometimes even at the design time of the system, and cannot be changed afterwards.

it to facilitate the mechanics of a payment and minimizing secondary uses (e.g., monetization of the data), as well as establishing strictly enforced retention periods.

On balance, we hence argue that better technological solutions offering hard privacy need to emerge, and that development efforts should also focus on building payment systems that combine hard and soft approaches. This means limiting the amount of data that is accessible in the first place via hard privacy, as well as making data abuse harder, via the use of technologies such as privacy-enhanced identity management, tokenization, technology-enforced limits on retention periods, and audits of data holders. Such a system would depend on compliance and accountability, but it is supported with technically enforced access control, limited retention periods, and audits. Finally, jurisdictions mindful that consumers are concerned by the limitations of current privacy-preserving methods in digital payments could also consider policies that promote the general acceptability of physical cash. This ensures a universally trusted, private payment option that is also resilient to emerging threats such as the one posed by quantum computing.

The remainder of this paper is structured as follows. Section 2 maps out the interests of various stakeholder groups. Section 3 presents a taxonomy of privacy and data access, and then discusses privacy-enhancing technologies. Section 4 discusses architectures for the design of digital money, while Section 5 looks more closely into lessons for payment system design. Section 6 concludes.

2 What's at stake: mapping the interests of users, law enforcement, and commercial firms

One step towards a sensible privacy design is understanding who the key stakeholders are and what their interests are in payment records. Knowledge of potentially conflicting interests helps develop requirements and narrow down the range of technical solutions. Stakeholders include the users (who might additionally be unbanked, undocumented, children, foreign residents, or tourists), the merchants, the banks and payment providers, the government (regulators, law enforcement, and intelligence agencies), and other parties with an interest in the tension between privacy and transparency (e.g., privacy advocates).

We started with a detailed stakeholder analysis of all these parties, but found that the key tensions are well-captured through consideration of only three stakeholder categories:

1. privacy enthusiasts: users of a payment system with an interest in privacy;
2. law enforcement: investigators of crimes with financial evidence; and
3. data holders: entities that record and monetize financial data including merchants, banks, and payments processors.

Figure 2 visualizes the conflicting relationships we identified between these stakeholder categories. Table 1 compares different payments options by how well they deal with the identified conflicts using a simple ordinal scale: good, OK, and bad.

We start with the relationship between privacy enthusiasts and law enforcement. Here, we consider a type of privacy enthusiast who is law-abiding and affirms that crimes can be deterred with effective law enforcement, yet believes that errors, breaches, and overreach are a potential future concern. Law enforcement prefers the least friction in obtaining payments information pertinent to their investigations.

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

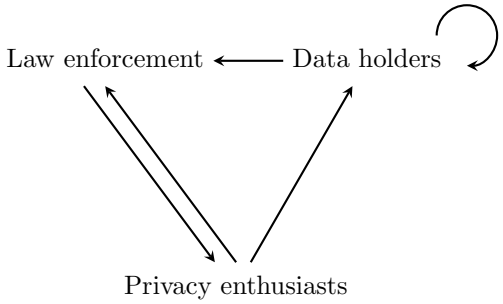

Figure 2: The main conflicts between stakeholder categories.

Table 1: Evaluation of payments options from the perspective of stakeholder conflicts.

	Law enforcement	Privacy enthusiasts		Data holders	
	L → P	P → L	P → D	D → D_{new}	D → L
Cash	Fair	Good	Good	Bad	Bad
Cryptocurrency	Bad	Depends	Depends	Fair	Bad
Soft privacy payments	Good	Bad	Bad	Good	Fair
Hard privacy payments	Bad	Good	Good	Bad	Bad

systems try to approximate this.[7] One research direction is to offer privacy for transactions under a certain threshold (e.g., \$10,000). This might not always serve the purposes of law enforcement, for example due to "smurfing", where large payments are split into multiple smaller transactions that fall below the threshold.

A privacy enthusiast would consider cash as largely addressing their concerns with law enforcement (denoted P → L in Table 1), while law enforcement would be concerned about the wide usage of cash by privacy enthusiasts (denoted L → P). However, law enforcement is not helpless at tracing cash, 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.

The level of privacy cryptocurrencies offer (next row in Table 1) is heterogeneous. The major ones, such as Bitcoin and Ethereum, are pseudonymous and can be de-anonymized (see [Meiklejohn et al., 2013] and [Nadler and Schär, 2023]). Other cryptocurrencies offer stronger 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, but the success varies.[8] Users today have the choice between many variants with differing levels of anonymity (e.g., X receives \$100 from someone), pseudonymity (e.g., X receives \$100 from Y), confidential transactions (e.g., Alice receives \$Z from Bob), or combinations thereof.

All payment systems can rely on soft and/or hard privacy. Soft privacy uses judicial oversight to allow human discretion in balancing exceptional access to payment data with privacy in current payment networks. Providing exceptional access to law enforcement (with judicial oversight) is favorable to law enforcement but can leave privacy enthusiasts concerned about its potential for abuse or breach.

[7] Another step in this direction is [Keller et al., 2021], who offer a design in which users can, without disclosing information on their identity, voluntarily disclose some information on the origin of their funds. For example, they may want to do so in order to prove that their funds do not originate from a protocol exploit. Proposals for implementations include [Baranski et al., 2023] and [Buterin et al., 2023].

[8] See [Ruffing et al., 2018] and [Möser et al., 2018] for a discussion of some weaknesses and vulnerabilities.

By contrast, hard privacy eliminates human intervention by relying solely on cryptography and perhaps tamper-resistant hardware. With current technology, it is possible to conceive hard privacy digital payments that are more difficult for law enforcement to trace than cash today (see the below discussion in the next sections). Similar to cash, even a completely anonymous payment system would not rule out other investigative techniques.[9]

While the tensions between privacy enthusiasts and law enforcement get top billing in the payments literature, the less obvious tension between privacy enthusiasts and data holders is equally important (P $\rightarrow$ D). Payments data is personal data; it can hence be monetized. It is used for profiling users through analysis techniques that improve and become cheaper over time. The difficulties in tracing cash are safeguarded, while payments systems enshrine banks and payments processors with valuable, proprietary data that could be useful for decades.[10] The development of a soft privacy digital payment systems would have to contend with a greater variety of stakeholders and a lack of transparency into how payments data can be used. This however also is also an opportunity to design a seamless set of soft-privacy rules from scratch, with the goal to protect abuse from today's commercial entities.

Data holders and law enforcement have minimal conflict. As long as law enforcement can obtain payments information, it is not generally concerned with who it is from. Inversely, data holders do find rules around identity gathering and reporting transactions expensive and onerous (D $\rightarrow$ 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 lead to more. Some solutions propose on-boarding users with cryptographically protected identities that can be used for selective traceability by law-enforcement (*cf.* [Wüst et al., 2019]). 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.

3 Privacy in payments: definition and technologies

The stakeholder analysis uncovers a nuanced picture of the underlying trade-offs, which requires moving beyond a rudimentary conceptualization of privacy as existing along a singular continuum between anonymity and full transparency. Instead, we next develop a taxonomy of privacy-enhancing designs along two dimensions: privacy vs auditability, as well as soft institution-based versus hard technology-based solutions. Distinguishing between soft and hard forms of auditability allows painting a much more nuanced picture of the tensions between various stakeholders. Many countries exhibit a convoluted set of payments options. These have evolved from stakeholders having competing interests, including tussles over privacy.

3.1 A taxonomy

We start this section by defining what it means to guard data against unauthorized access and give authorized access to it, and by mapping out a taxonomy of the technological and institutional means by which this can be achieved.

In both the academic literature and policy discussions, readers are often presented with a vague mental model of privacy on a single line, with one end marked privacy and the other end marked transparency, and told to think of the situation as a trade-off. Recently, however, cryptographers have advocated sophisticated

[9] We note that one challenge for operable privacy is that too much privacy can backfire. 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, for example web traffic. 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 payment system in the first place. This hints at an apparent paradox—that there are circumstances where increasing access to data can increase privacy.

[10] Payment data can hence lead to the emergence of monopolies and foster their entrenchment (see [Garratt and Lee, 2021] and [Bank for International Settlements, 2021]).

techniques that purport to allow having both privacy and transparency [Chatzigiannis et al., 2021]. For payments specifically, this might involve granting access to some payment records under certain specified conditions. We evaluate a number of these proposals in Section 4.3. Examples include access to the identity of the payee but not the payer [BIS Innovation Hub, 2023b, Chaum et al., 2021, Chaum and Moser, 2022];[11] access to transactions over a certain amount [Wüst et al., 2019]; or access to all transactions except those chosen to be protected (up to a certain amount in a time period) by the payer [Tomescu et al., 2022, Zafrani et al., 2022].

Properly studying how a payment system can offer both privacy and access to data, as well as characterizing the underlying technologies requires an expansion of our mental model to a quadrant system like Figure 3. The two dimensions of this figure regard the way in which the payment system protects the data of its users, as well as how it grants access to third parties should such access be authorized. More specifically, the two dimensions, in turn, are defined as:

- **Privacy.** Privacy is the freedom from unauthorized disclosure of one's personal payment information to any other party.
- **Auditability.** Auditability refers to the ability of specified third persons to obtain personal payment information once such access has been lawfully authorized.

We note that payment privacy can be achieved in two distinct ways. In today's financial system, holders of digital payment records — banks, credit card schemes, payment processors, financial intermediaries — implement internal controls and deploy cybersecurity techniques to protect user privacy. They might also minimize data retention to the minimum required under law. We consider this "soft" privacy as the data is still accessible.[12] Soft privacy is preferred by enforcement agencies, who often seek court-authorized access to this data to investigate crimes.

By contrast, "hard" privacy ensures the data will never be accessed, locking out enforcement and intelligence agencies from this data intentionally or not. Some forms of payment, like banknotes, can be accomplished without needing a data trail. Others, less time-tested but seen in some existing cryptocurrencies, can utilize cryptographic techniques to offer anonymous payments (X pays Y $100), confidential transactions (Alice pays Bob $Z), or both.

In sum, soft and hard privacy are defined as follows for the domain of payments:

- **Soft Privacy.** Payment records are protected by access control measures implemented by the data holder and can be accessed in plaintext when authorized.
- **Hard Privacy.** Payment records (i) are not created in the transaction process, or (ii) are protected by default through cryptography and user-guarded secrets and cannot be accessed by the data holder.

Researchers and policy makers often argue that a "middle ground" between soft and hard privacy needs to be found, for example using strong cryptographic protection that can be decrypted under pre-defined conditions. We argue that a better conceptualization of these techniques is to consider a second dimension: auditability. Once access is granted, auditability defines how broad the access is, how verbose the data records are, and how easy it is to pull in new data at the discretion of the investigators. Auditability can also be "soft" and "hard":

- **Soft Auditability.** Any stored payment data can be shared with authorized stakeholders, assuming a reasonable case for access can be built under the law.

[11] See also the GNU Taler for an earlier implementation: `http://taler.net/`

[12] These techniques are not foolproof, as the history of data breaches and insider attacks demonstrate.

- **Hard Auditability.** All data is decrypted and hence inaccessible by default. Computer programmable rules that leave no room for discretion can make select data available to authorized stakeholders.

Payment data should be accessible to authorized entities (auditability) while being inaccessible to unauthorized entities (privacy). Prioritizing auditability could result in devastating data breaches or abuses, while prioritizing privacy could result in missed enforcement. If enforcement agencies could reduce their investigative techniques to a set of rules that accurately predict if payment data is likely to be investigated in the future or not, hard privacy could be applied conditionally. The cryptography for conditional privacy can be built using many of the tools presented in the next section.

3.2 A description of hard privacy technologies

What are the various types of technologies that enhance hard privacy, what do they set out to achieve, how do they differ, and what are their limitations? We summarize these in Table 2. The technologies are varied in terms of maturity with some techniques dating back to the 1980s while others were invented in the past 5 years. It is important to understand the limitations of each technology and to consider if the assumptions they make are realistic. While the techniques can be combined or layered to some extent, it is often a non-trivial undertaking to ensure the strengths of each are preserved and new weaknesses are mitigated.

A number of hard privacy technologies have emerged for digital transaction and data handling, each with distinct advantages and disadvantages.

- Zero-knowledge proofs (ZKPs): ZKPs enable the proof of information possession without revealing the information itself. For instance, a user can prove they have sufficient funds for a purchase without disclosing their actual bank balance. This enhances privacy in transactions. Advanced versions like SNARKs and STARKs improve efficiency for those verifying the proofs but increase costs for the prover, and may require using a set of entities, trusted not to collude, during the initial system setup. Generally ZKPs face the disadvantage of high computational costs, making them resource-intensive. For example, a payment created with ZKPs, circulated amongst 10 banks, and verified by all parties might take 800ms [Narula et al., 2018].

- Homomorphic encryption (HE): HE allows computations on encrypted data, enabling actions like updating account balances in payment processing without exposing actual figures. An advantage of HE is its ability to perform data analysis while preserving privacy, crucial for handling sensitive financial data. On the downside, HE is computationally demanding, especially for complex operations, which limits its practical use on consumer-grade hardware. Viand *et al.* write, "as a very rough heuristic, computations that take more than a few hundred milliseconds without FHE [fully HE] are unlikely to be practical once translated into FHE as of today" (see [Viand et al., 2021]). Managing decryption keys is also a challenge, adding complexity to its implementation.

- Secret sharing (SC) or Multi-party computation: SC involves multiple parties collaboratively computing functions without revealing individual inputs. This is advantageous in scenarios like joint decision-making in financial transactions, where privacy of individual inputs is maintained. However, SC requires continuous communication among all parties, which can be logistically challenging. One SC experiment, using real-world trading data in a simulation, found that 41% of trades incurred a wait time of longer than 5 seconds with the longest trade taking 101 seconds [Li et al., 2023]. Also, the necessity for all parties to be constantly online can limit its applicability in environments with intermittent connectivity.

- Anonymity-enhanced signatures: Technologies such as ring and blind signatures significantly enhance privacy in digital transactions by enabling users to verify attributes or sign messages anonymously. A ring signature permits a member of a group to sign a message in a manner that confirms the signature

originates from a group member, while making it computationally infeasible to identify the specific individual responsible. This ensures the anonymity of the signer, while still providing the verifier with assurance of the group's authenticity. Conversely, a blind signature involves a process whereby the signer endorses a message without having visibility of its content, akin to signing a document enclosed in an envelope without seeing the contents. This ensures the message remains concealed from the signer, yet the signature can subsequently be verified as genuine. However, these technologies face interoperability challenges with other cryptographic systems. Additionally, revoking or invalidating signed data in anonymous contexts can be difficult, posing a challenge in maintaining data integrity.

- Tamper-resistant hardware: Hardware Security Modules (HSMs), smart cards, and Physically Unclonable Functions (PUFs) provide strong physical security solutions, safeguarding digital keys and sensitive operations. This is crucial in environments like banking, where data security is paramount. However, the effectiveness of such hardware is reliant on the integrity of the manufacturing and supply chain, posing a risk if this trust is compromised. Additionally, adapting these technologies to emerging threats requires ongoing updates and certifications, which can be resource-intensive.

- Trusted computing: Technologies like Trusted Platform Modules (TPMs) and Trusted Execution Environments (TEEs) secure general-purpose devices against software attacks, enhancing the security of payment processors and sensitive data. A notable advantage is their ability to safeguard devices in a broad range of consumer environments. The primary disadvantage is that they offer a lower level of security compared to specialized, dedicated hardware [van Schaik et al., 2022]. Moreover, their reliance on the control of operating system and platform manufacturers can limit their flexibility and accessibility.

- Privacy metrics: k-anonymity and differential privacy are methods to protect individual privacy in datasets, guiding the removal or obscuring of information. These metrics are beneficial for generating statistical analyses on payment data without compromising individual privacy. However, the balance between data utility and privacy can be delicate, with the potential to render data less useful if excessive privacy controls are applied. Additionally, the effectiveness of these metrics depends on the accuracy of assumptions about data distribution, which can be challenging in dynamic or evolving datasets. Domingo-Ferrer *et al.* write, "it is impossible to collect [differential privacy]-protected data from a community of respondents an indefinite number of times with a meaningful privacy guarantee" (see [Domingo-Ferrer et al., 2021]).

While these technologies offer a range of solutions for enhancing privacy and security in digital transactions and data handling, each comes with its own set of advantages and disadvantages. The choice of technology depends on the specific requirements and constraints of the application, including the need for computational efficiency, the level of privacy required, and the practicality of implementation in the intended environment. As technology evolves, so do the methods to safeguard privacy, necessitating a continuous assessment and adaptation of these technologies to ensure effective and robust privacy protection in our digital world.

Zero-knowledge proofs (ZKP)

SNARKs, STARKs, Σ-protocols

A software client based on ZKP produces a sequence of numbers that can be shown to others to convince them that asserted features about their private data are true. For example, for a payee to know that a payment will be executed, it suffices to know that the payer's account balance exceeds the value of the payment amount. ZKPs are often used in combination with other approaches described below.

Limitations. Very high computation cost – depending on the implementation by a factor 100–1000 compared to unencrypted text (plaintext) for general features. The security of ZKPs often depends on mathematical assumptions that are not yet time-tested. The computationally more efficient types of ZKPs often have vulnerabilities when it comes to collusion during the initial system setup. With many users proving different features in ZKP, privacy can be lost if users can be tracked by what they prove.

Homomorphic encryption (HE)

Partially, fully (FHE), commitments

A payment processor can perform computations on encrypted data without having access to the plaintext. For example, a payment amount can be subtracted from an account balance without learning either the payment amount or the balance.

Limitations. Extremely high computational costs. On consumer grade computers, only a small set of predefined computations can be performed. For example, while subtraction or addition may work on a consumer grade computer, multiplication or division is well out of bounds. No existing payments architecture offers a convincing solution for handling access to decryption keys.

Secret sharing

Two/multi-party computation, threshold signatures, multi-sig

Private data is stored simultaneously with multiple data holders in such a way that no individual (or subset below a definable threshold, for example 4 out of 7) can reconstruct the information. By extension, not only can data be stored, but also pre-defined computations can be performed on the data with the same confidentiality. Some examples include multi-sig, key recovery, or distributed auctions.

Limitations. Hypothetically applicable, but there do not exist workable technological design proposals. Computation requires substantial synchronous communication between the data holders. Data holders always need to be online, even if they can be trusted to never deviate from the protocol.

Anonymity-enhanced signatures

Ring/blind signatures, anonymous credentials

A user can present a message signed by a payment processor in such a way that no one can link it back to the user who asked for the signature (blind signatures). This is applicable to verifying attributes (*e.g.,* age, unspent coin) of a user's identity (anonymous credentials). Anonymous e-cash can also be built such that showing the same attribute twice (*i.e.,* double-spending a coin) revokes the user's anonymity. A final variation (ring signature) allows a user to sign a message in such a way that no one can tell who exactly signed the message, only that it was one person within a group of people chosen by the user.

Limitations. Best examples do not easily interoperate (cryptographically) with other primitives. Revocation of signed data is hard. Payment systems built with this often require interactive spending protocols and double-spending is detectable (but not directly preventable). Ring signatures become computationally inefficient as the group grows large, which limits anonymity.

Tamper-resistant hardware

Hardware security modules (HSMs), smart cards, physically uncloneable functions (PUFs)

A manufacturer designs very efficient, special-purpose hardware robust against tampering. It can be provisioned to a processor (*e.g.,* HSM) to handle payments while preventing insiders from accessing or changing customer data. Alternatively, it can be given to users and merchants (*e.g.,* smart cards) to execute payment protocols while protecting PINs, balances, and other data. PUFs are a more futuristic building block for keeping pieces of hardware unique, assuming certain physical properties hold. Can be used in combination with cryptography to secure end points and keys.

Limitations. Resistance to tampering or cloning is assumed based on the current state of art and science, and reasonable limits on the adversary's budget. Attacks get better and cheaper over time, requiring regular replacements and costly certification processes. Untested at scale for fully offline payments with large (*e.g.,* over $1000) account balances. Manufacturer and supply chain of the specialized hardware must be fully trusted, akin to banknotes.

Trusted computing

TPMs, TEEs, Intel SGX, ARM TrustZone

Trusted computing uses hardware-level cryptography to provide a special mode for general purpose devices (*e.g.,* a computer or phone) that is secure against software attacks (*e.g.,* malware) and consumer-grade modification (*e.g.,* jailbreaking on phones). Payment processors can use trusted computing to deploy services on untrusted cloud services and/or untrusted user devices.

Limitations. Security appears to be lower than dedicated tamper-resistant hardware—many attacks are known and discovered each year that work in laboratory conditions. Currently, most operating system and platform manufactures take exclusive control over these features. Manufacturer and supply chain of general purpose devices must be fully trusted.

Privacy metrics

k-anonymity, differential privacy

One approach to privacy is to remove information before data is shared. This primarily protects individuals in the data set. Privacy metrics guide the selection and amount of information to remove. Common forms include understandable heuristics (*e.g.,* k-anonymity) for making records more similar (*e.g.,* distinct birth years, 1972 and 1978, are replaced by the attribute 1970s); and mathematically principled frameworks that add random noise (*e.g.,* differential privacy), resembling randomized response techniques known from population surveys. This approach is often combined with cryptography to blind data instead of fully removing it.

Limitations. The entity applying the metric needs to have a copy of the raw data and must be trusted. These metrics do not tell users how to remove information; they only specify when to stop. Applications are best-suited for generating statistics on payments (secondary use of the data) rather than protecting the primary data. Naturally, heuristic metrics are based on non-trivial assumptions about the data distribution, and, if inaccurate, privacy may be compromised. Specifically, if data is added over time (in live systems), it becomes harder and harder to continually assure the assumptions are met. Strict mathematical frameworks are very conservative, protecting the most extreme cases (*i.e.,* the most exposed individual in the worst case). This can result in adding so much noise that the data becomes useless. Claims regarding the use of differential privacy in industry often prioritize the utility of the data over privacy hence being more to the effect of marketing purposes rather than protecting users.

Table 2: Summary of technology that enables hard privacy.

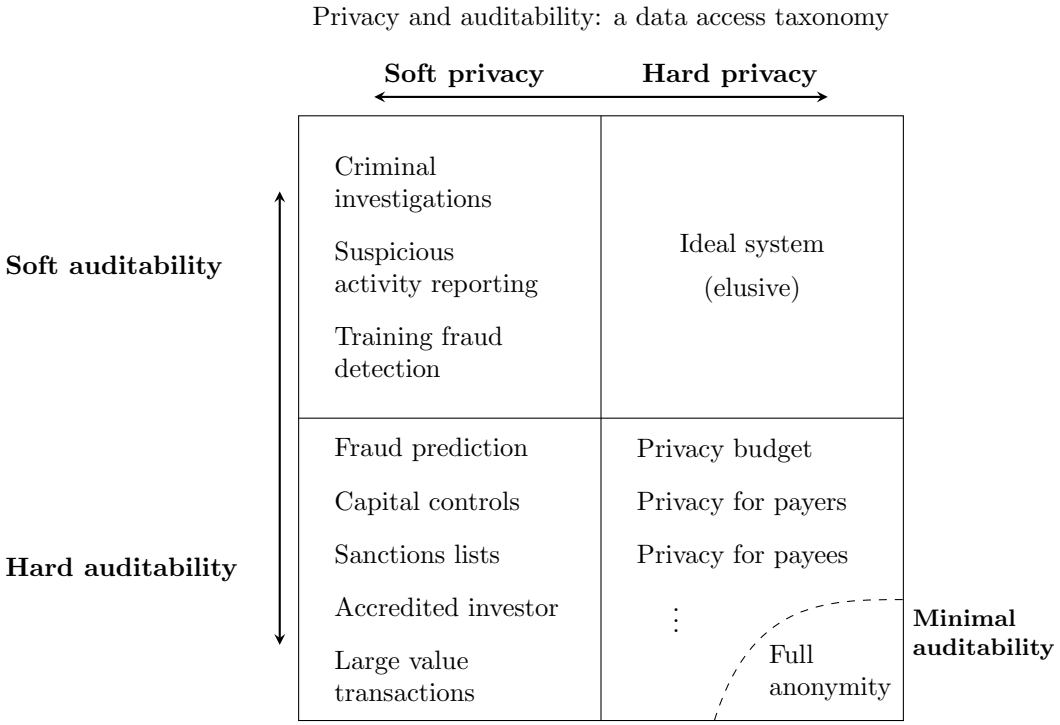

Figure 3: Privacy is the freedom from unauthorized disclosure of one's personal payment information to any other party, while auditability refers to the ability of specified third persons to obtain personal information once such access has been authorized. Soft solutions are institution-based, while hard solutions are technology-based without room for human discretion.

4 Evaluating architectures for privacy in digital payments

While payments records contain sensitive data and must be protected by default, in certain cases, access to payments records is vital to prevent crime. We next illustrate how the combination of hard and soft privacy and accessibility can deal with the various conflicts of interests between the stakeholders. A summary is given in Figure 3 and the following text walks through the quadrants counter-clockwise.

4.1 Soft Privacy and Soft Auditability

Electronic retail payments systems operate under a default of soft privacy and soft auditability in most advanced economies. This regime enables exceptional access to payment records by law enforcement and other government investigators when authorized by the judicial system. The advantage from an enforcement perspective is that investigations can be flexible and look for payment patterns that are "suspicious" but difficult to exactly specify. It also can cover a wide variety of crimes, essentially anything that is not based on using cash.

There are also disadvantages and concerns. On the privacy side, if records can be produced for enforcement teams, they are in reach of insiders and could be breached by an outside adversary. They also enable secondary uses for the data. On the auditability side, discretionary access requires trust in the legal system, especially since many investigations are kept confidential from the public. Distrust of this system leads to concerns of surveillance, dragnet investigations, or the potential of such a regime to creep in this direction

over time.

Finally, soft auditability is expensive. It is common for situations in this quadrant to involve two parties: the holder of the financial data and the enforcement team who obtains and analyzes the data. However for specific kinds of financial crime—money laundering, financing terrorists, evading sanctions—the onus of who conducts a preliminary analysis has shifted onto the data holder, who must examine their own data and pro-actively report suspicious activity. Suspicious activity reports (SARs) are soft auditability because there is no hard (computer-decidable algorithm) definition of "suspicious." The financial burden of SARs on data holders is characterized by the D→L tension in our stakeholder analysis. Also the possibility of data breaches mean it is easier to discard data than to retain it and maintain its confidentiality.

4.2 Soft Privacy and Hard Auditability

At first, this quadrant might seem pointless. If privacy is soft, law enforcement can seek authorization to any data at its discretion. However this quadrant serves as a complementary. In today's payment system, by establishing hard rules for auditability, law enforcement has low friction access to data it considers the most pertinent, while having soft auditability as a fall-back mechanism. To illustrate this, consider the current policy in many countries that banks and financial services must automatically report large value transactions (*e.g.,* over $10K USD) to an enforcement agency. This does not mean law enforcement is locked out from accessing smaller value transactions. In contrast, law enforcement gains automatic access to larger payments without the friction from obtaining special authorization through the courts. If however an investigation leads to data outside these disclosures, law enforcement can take a higher friction path to getting the data via the courts. Machine-decidable rules can also apply to cross-border money transport (*e.g.,* an upper limit on cash) or to halting payments (*e.g.,* a sanctions or terrorist organization list) although banks are expected to also apply softer discretion and have an onus to recognize attempts at obfuscating such payments.

However since hard auditability rules are most effective when clearcut (*e.g.,* transactions between two entities that exceed $10K USD within 24 hours), it has been questioned if these are actually the most useful to law enforcement or if they are just the easiest to codify. Another challenge with hard auditability is that publicly known rules can be evaded (*e.g.,* breaking large transactions into smaller ones executed between a variety of entities to obfuscate the fact that the payment ultimately flows to a single entity, which is called smurfing). Even if rules can be evaded, the increased costs, delays, and friction of obfuscating payments can be a deterrent.

4.3 Hard Privacy and Hard Auditability

In this quadrant, a machine-decidable rule that cannot be manipulated is established to decide if a payment is accessible by law enforcement or if it will be given cryptographic protection that precludes access (at least, without the involvement of a party to the transaction). Examples from the literature develop the appropriate cryptography (typically employing advanced privacy enhancing technologies like zero-knowledge proofs, homomorphic encryption, or blind signatures from Table 2) for applying privacy in a conditional way to payments.

The main takeaway from the previous two soft privacy quadrants is that hard auditability is not a regime onto itself, but works in concert with soft auditability. This is often misunderstood by those who advance hard privacy with hard auditability. Without soft auditability as a fallback, the audit rules need to be more precise than the rules from the previous subsection 4.2, or they will impede law enforcement investigations.

The far-end corner of this quadrant affords full anonymity to every transaction, leaving no auditability of the payment records. While it might sound extreme, it is in fact a potential location to consider for a CBDC. It would be the embodiment of true digital cash, since transacting with cash (banknotes) does not create a payment record by default. A CBDC could emulate this with either tamperproof hardware that can

exchange value offline (hypothetically at least, as sophisticated hardware attacks are feasible if the gains are high enough) or with strong encryption.

The remainder of the quadrant develops a hard-coded 'rule' for whether a payment, or part of a payment, should be protected with strong cryptography that makes it inaccessible for law enforcement.

Examples of rules are given in the rows of Table 3. The first rule is payer privacy, where the operator of the payment system cannot tell who made a payment, while recipients could be investigated [Chaum et al., 2021, Chaum and Moser, 2022].[13]

Privacy for payers: The operator of the payment system cannot tell who made a payment, while recipients could be investigated [Chaum et al., 2021, Chaum and Moser, 2022].
Privacy for payees: No one learns who the recipient of a payment is, while senders could be investigated [Todd, 2014].
Privacy threshold: Payments below a defined threshold is guaranteed full anonymity [Wüst et al., 2019].
Privacy budget: Payers have a secondary account with full anonymity but can only deposit into it a set amount of currency per time period [Tomescu et al., 2022, Zafrani et al., 2022].
Privacy with aggregate disclosure: Payments have full anonymity, but account holders can elect to prove statistics about their accounts [Narula et al., 2018].
Privacy with alibi: Payments have some anonymity, but payers can elect to prove that a given payment is not theirs [Buterin et al., 2023].

Table 3: Possible definitions of hard auditability and hard privacy.

Conversely, a system with payee privacy ensures no one learns who the recipient of a payment is, while senders could be investigated [Todd, 2014]. Some systems focus on the transactional value of the payment, provided full privacy for all payments under a system-determined threshold while payments above it could be investigated [Wüst et al., 2019]. Slightly different, payers could have a secondary account with full anonymity but can only deposit into a limited amount of funds based on a system-determined budget per time period [Tomescu et al., 2022, Zafrani et al., 2022].

Individual payments could enjoy full anonymity but account holders can elect to prove statistics about their accounts [Narula et al., 2018]. Finally, payments might be anonymous by default, but the payers retain sufficient data to disavow any payment they did not make—an alibi if investigated. While law enforcement does have access to payer information directly, it could use a process of elimination to identify payers.

In all this, one must keep in mind that, aside from the performance and security issues discussed above (cf Table 2), the set of rules might not be ready, even if the technology is. For example, in digital settings, it is trivial to smurf a large payment so that each sub-payment is below a threshold. This suggests that simply adopting hard auditability rules from the soft privacy quadrant and applying them to hard privacy is not appropriate. These soft privacy rules were designed with the assumption that soft auditability would be available as a fallback mechanism.

An interesting avenue for future research is hence to examine, from both a law enforcement and a privacy perspective, how select hard privacy rules could be combined to facilitate effective crime enforcement without

[13]GNU Taler: `http://taler.net`

compromising too much privacy. There are many different types of crime involving various types of payments, and some hard privacy regimes might be suitable for certain crimes but not others. For example, privacy thresholds or privacy budgets might be well-suited to preventing large-scale money laundering or asset smuggling. Additionally, privacy for payers (but not payees) might make activities like extortion, terrorist financing, and receiving bribes more transparent.

4.4 Hard Privacy and Soft Auditability

The final quadrant represents the ideal world in some sense—all payment records are cryptographically protected when they are not of interest to law enforcement, and accessible if a crime has been committed, even if law enforcement cannot articulate a clear-cut, machine-decidable rule for when it is of interest. As stated, this is elusive, and so might be trials to approximate it.

5 Implications for payment system design

We next discuss different privacy architectures for payment systems in the light of our stakeholder analysis and technological framework. The stakeholder analysis reduced dozens of stakeholders to three with the intention to map conflicts. In an attempt to answer these questions, it is convenient to classify stakeholders by their involvement in the payments process. Some stakeholders have and require more information simply to make payments work. We refer to this a primary use of payments data. Other stakeholders are less relevant for the payments process itself, but might want the information for *other* purposes than payments. We refer to this as secondary use of data.

Starting from primary use, today's banks and payments processors must have the underlying data to process payments. This data however has secondary uses, such as credit scoring. The extent to which this happens varies between payments systems and regions, with the tendency to explore more secondary uses through an emerging financial technology sector [Boissay et al., 2021]. There also exists a data industry outside of the financial sector, it is comprised of merchants, who seek to commercialize the secondary uses of payments data. It also includes technology suppliers who seek to expand tracking and targeting across merchants, industries, and increasingly to the offline world. This is where payments data is deemed particularly valuable. Not only are payment records reliable indicators of economic activity and consumer choice, Fintechs may aim to link them back to persistent identities up to real names and street addresses.

Law enforcement is involved in payments only indirectly. For example, they can strengthen trust in a payment system by tracking stolen funds based on the authorization of the rightful owner. While the former is a primary use of the data, law enforcement is also interested in secondary uses. They believe that bulk access to records—disputed and undisputed alike—helps them to solve all kinds of crime.

Option 1: Hard privacy

Cypherpunks would argue to place hard privacy everywhere, meaning that plaintext access is reserved to end users with access to private keys only. However, such a system would suffer from technical overhead because encrypted records tend to require much 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 re-encryption.

The benefit of this effort is excluding everyone but primary data users (i.e. those who pay or receive funds) from access. However, law enforcement is both deprived of the capabilities needed to solve crimes. 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 operators and users, as is the case with cash. A shadow system of record-keeping mirroring almost all of the activity in the payments system proper could result in less privacy due to poorer security and governance. For example, some regulators receive a picture of all cryptocurrency transactions that fall under a "travel rule." This may create an inefficient dragnet which neither meaningfully prevents crime nor protects the privacy of legitimate users.

We also note that the history of cryptography is littered with failure. Lessons like 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 where no security proof is known. TLS solves a relatively simple task of securing communication with a well-identified server using basic cryptography primitives like encryption, hashing, and creating digital signatures. By contrast, the cryptographic building blocks with conditional privacy are much newer, more involved, and less reviewed. Moreover, few of these techniques scale well to billions of real-time payments (cf. Table 2).

Option 2: Soft privacy

The opposite would be soft privacy everywhere, with auditability rules as in regulated payment systems today. Payments data would freely flow between 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 from a computational perspective. However, its privacy would not be any different than that offered by current payments networks, and hence worse than paying in cash. It will therefore be difficult to gain the trust of privacy enthusiasts, chiefly for the weak guarantees it offers to discourage unwanted secondary use of payments data. Another threat to privacy is that commercial data holders bend the interpretation of law in their interests (compare with web tracking today [Matte et al., 2020]).[14] Moreover, preventing abuse of stolen data, e.g., after data breaches, is technically impossible and a real threat given the frequency of breaches [Wheatley et al., 2016]. Similar concerns apply to state actors 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 implying that a different compromise needs to be found.

Option 3: A soft core with a hard shell

A system built around a soft core with a hard shell would start with data minimization techniques. All data is fully encrypted by default. The operator of the system enforces 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 ten years for aggregates like account balances. The choice of these periods reflects the elevated sensitivity of detailed payments records and their metadata, as compared to account holder information that financial intermediaries must retain today to combat financial crimes.

Some entities in the payments process need some plaintext access for computational efficiency, hence implementing soft privacy here around this core. To respect the principle of data minimization, no other party is given access by default. However to serve justified data requests from law enforcement, the operator of the payments system acts as a data custodian for data that has not been deleted: it may grant plaintext access to selected records while it is ensured that all requests are authentic, justified, proportionate, and leave an audit trail—soft auditability with stronger accountability and transparency.

[14] Another issue speaking against soft privacy regards foreign intelligence agencies. These organizations strive to capture data of all types, with payments data offering crucial links to real-world entities. It is particularly hard to ensure that soft privacy rules are not circumvented by foreign actors.

Again, hard privacy techniques could be used, for example to govern this data transfer. Existing techniques like tokenization, used in current payments networks to shield customer data from merchants, can be a source of inspiration. Likewise, anonymized aggregate data could be made available to the non-financial sector, again under transparent rules.

This approach seeks to balance hard and soft privacy while enabling soft auditability. Systems can further be designed in a way that payments records are protected in bulk—e.g., against devastating data breaches—but plaintext access is possible in justified cases. The conditions for plaintext access must 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.

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

6 Concluding Remarks

The privacy landscape of digital money is more complex than often appreciated. We start mapping it through different perspectives—soft/hard privacy and soft/hard auditability, (ii) stakeholder conflicts and (iii) stakeholder proximity to the data—which we have not seen in the literature before. Both of these 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 a more complete view offers aspects we missed.

In this process we have arrived at new insights. For one, a glaring gap in the literature is the lack of realistic privacy definitions for money, which hinders research and development of hard privacy solutions that address all realistic concerns when trading off privacy and crime fighting. On balance, it is essential to recognize that while advanced concepts like zero-knowledge proofs offer promising avenues for enhancing privacy in payment systems, their full potential remains somewhat elusive in today's practical applications that may need to process thousands of payments per second. These technologies, though theoretically sound, face real-world challenges such as high computational demands and complex implementation requirements.

Therefore, we conclude that there is a pressing need for the development of more robust technological solutions that offer hard privacy — a level of privacy protection that is unyielding and comprehensive. In addition, there is also a need to focus on enhancing soft privacy, which can be effectively supported through technically enforced access control. This approach involves designing payment systems that make any abuse of data transparent and ensure accountability of data holders. The underlying principle here is not just to restrict access to sensitive data but also to make any unauthorized access or misuse evident and traceable. 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.

These dual approaches — pursuing hard privacy through emerging technologies and reinforcing soft privacy through access control and accountability mechanisms — represent a balanced strategy for future developments in digital payment systems.

Additionally, our analysis shows that, while privacy-enhancing technologies in digital payments offer promising avenues, some open questions remain at present. Jurisdictions mindful of this dimension could consider policies that ensure the continued acceptability of physical cash, as it offers an immediate, universally trusted, privacy-preserving option. For example, central banks and the public sector can facilitate the availability of cash and maintaining its wholesale distribution networks, thereby preserving its role as a physical payment option alongside emerging digital alternatives.[15]

[15] [Auer et al., 2020] discuss the measures central banks have taken to support the continued acceptance of cash during the

References

[BIS, 2023] (2023). Project leap: Quantum-proofing the financial system. Technical report, Bank for International Settlements.

[Acquisti et al., 2022] Acquisti, A., Brandimarte, L., and Hancock, J. (2022). How privacy's past may shape its future. *Science*, 376(6578):270–272.

[Agur et al., 2022] Agur, I., Ari, A., and Dell'Ariccia, G. (2022). Designing central bank digital currencies. *Journal of Monetary Economics*, 125:62–79.

[Auer and Böhme, 2021] Auer, R. and Böhme, R. (2021). Central bank digital currency: the quest for minimally invasive technology. BIS Working Papers 948.

[Auer et al., 2023] Auer, R., Böhme, R., Clark, J., and Demirag, D. (2023). Mapping the privacy landscape for central bank digital currencies. *Commun. ACM*, 66(3):46–53.

[Auer et al., 2020] Auer, R., Cornelli, G., and Frost, J. (2020). Covid-19, cash, and the future of payments. *BIS Bulletin*, (3). 9 pages.

[Auer et al., 2024] Auer, R., Dupont, A., Gambacorta, L., Park, J. S., Takahashi, K., and Valko, A. (2024). Quantum computing and the financial system: Opportunities and risks. Technical Report 149, Bank for International Settlements.

[Bank for International Settlements, 2021] Bank for International Settlements (2021). Central bank digital currencies: an opportunity for the monetary system. Annual Economic Report 2021, Chapter III.

[Baranski et al., 2023] Baranski, S., Dotan, M., Lotem, A., and Vald, M. (2023). Haze and Daze: compliant privacy mixers. Cryptology ePrint Archive, Paper 2023/1152.

[BIS Innovation Hub, 2023a] BIS Innovation Hub (2023a). Project Polaris: A handbook for offline payments with cbdc. Bank for International Settlements.

[BIS Innovation Hub, 2023b] BIS Innovation Hub (2023b). Project Tourbillon: Exploring privacy, security and scalability for cbdcs. Bank for International Settlements.

[Boissay et al., 2021] Boissay, F., Ehlers, T., Gambacorta, L., and Shin, H. S. (2021). Big techs in finance: on the new nexus between data privacy and competition. *BIS Working Papers*, (970).

[Buterin et al., 2023] Buterin, V., Illum, J., Nadler, M., Schär, F., and Soleimani, A. (2023). Blockchain privacy and regulatory compliance: Towards a practical equilibrium. Available at SSRN. Date Written: September 6, 2023.

[Chatzigiannis et al., 2021] Chatzigiannis, P., Baldimtsi, F., and Chalkias, K. (2021). Sok: Auditability and accountability in distributed payment systems. In *Applied Cryptography and Network Security: 19th International Conference, ACNS 2021, Kamakura, Japan, June 21–24, 2021, Proceedings, Part II*, page 311–337, Berlin, Heidelberg. Springer-Verlag.

[Chaum et al., 2021] Chaum, D., Grothoff, C., and Moser, T. (2021). How to issue a central bank digital currency. *CoRR*, abs/2103.00254.

[Chaum and Moser, 2022] Chaum, D. and Moser, T. (2022). ecash 2.0 inalienably private and quantum-resistant to counterfeiting. Technical report.

[Darbha and Arora, 2020] Darbha, S. and Arora, R. (2020). Privacy in cbdc technology. Bank of Canada Staff Analytical Note 2020-9.

COVID-19 pandemic. [Schlegel, 2022] discusses critical elements of cash distribution systems and how they can be supported.

[Domingo-Ferrer et al., 2021] Domingo-Ferrer, J., Sánchez, D., and Blanco-Justicia, A. (2021). The limits of differential privacy (and its misuse in data release and machine learning). *Commun. ACM*, 64(7):33–35.

[European Data Protection Supervisor, 2023] European Data Protection Supervisor (2023). Financial and payment services: Use of personal data should remain proportionate and fair. EDPS Opinion 38/2023.

[Frost et al., 2020] Frost, J., Gambacorta, L., Huang, Y., Shin, H. S., and Zbinden, P. (2020). BigTech and the changing structure of financial intermediation. *Economic Policy*, 34(100):761–799.

[Garratt and Lee, 2021] Garratt, R. and Lee, M. J. (2021). Monetizing privacy. Staff Report 958, The Federal Reserve Bank of NY, New York, NY.

[Gross et al., 2021] Gross, J., Sedlmeir, J., Babel, M., Bechtel, A., and Schellinger, B. (2021). Designing a central bank digital currency with support for cash-like privacy. *Available at SSRN*.

[Kahn et al., 2005] Kahn, C. M., McAndrews, J., and Roberds, W. (2005). Money is privacy. *International Economic Review*, 46(2):377–399.

[Kappos et al., 2018] Kappos, G., Yousaf, H., Maller, M., and Meiklejohn, S. (2018). An empirical analysis of anonymity in zcash. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 463–477.

[Keller et al., 2021] Keller, P., Florian, M., and Böhme, R. (2021). Collaborative deanonymization. In *Financial Cryptography and Data Security. FC 2021 International Workshops*, volume 12676 of *Lecture Notes in Computer Science*, pages 39–46. Springer.

[Li et al., 2023] Li, Y., Soska, K., Huang, Z., Bellemare, S., Quintyne-Collins, M., Wang, L., Liu, X., Song, D., and Miller, A. (2023). Ratel: Mpc-extensions for smart contracts. Cryptology ePrint Archive, Paper 2023/1909. `https://eprint.iacr.org/2023/1909`.

[Matte et al., 2020] Matte, C., Bielova, N., and Santos, C. (2020). Do cookie banners respect my choice? measuring legal compliance of banners from IAB europe's transparency and consent framework. In *IEEE Symposium on Security and Privacy*, pages 791–809. IEEE.

[Meiklejohn et al., 2013] Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2013). A fistful of bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 127–140.

[Möser et al., 2018] Möser, M., Soska, K., Heilman, E., Lee, K., Heffan, H., Srivastava, S., Hogan, K., Hennessey, J., Miller, A., Narayanan, A., and Christin, N. (2018). An empirical analysis of traceability in the monero blockchain.

[Nadler and Schär, 2023] Nadler, M. and Schär, F. (2023). Tornado cash and blockchain privacy: a primer for economists and policymakers. *Federal Reserve Bank of St. Louis Review*.

[Narula et al., 2018] Narula, N., Vasquez, W., and Virza, M. (2018). zkledger: Privacy-preserving auditing for distributed ledgers. In *15th {USENIX} symposium on networked systems design and implementation ({NSDI} 18)*, pages 65–80.

[Ruffing et al., 2018] Ruffing, T., Thyagarajan, S. A., Ronge, V., and Schroder, D. (2018). (short paper) burning zerocoins for fun and for profit-a cryptographic denial-of-spending attack on the zerocoin protocol. In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, pages 116–119. IEEE.

[Schlegel, 2022] Schlegel, M. (2022). Popular, but under pressure - cash in the digital age. In *Forum for Financial Market Stability, FMA Liechtenstein*, Vaduz, Liechtenstein.

[Tinn, 2024] Tinn, K. (2024). A theory model of digital currency with asymmetric privacy. Technical report, CEPR Discussion Papers.

[Todd, 2014] Todd, P. (2014). Stealth addresses. https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2014-January/004020.html.

[Tomescu et al., 2022] Tomescu, A., Bhat, A., Applebaum, B., Abraham, I., Gueta, G., Pinkas, B., and Yanai, A. (2022). Utt: Decentralized ecash with accountable privacy. *Cryptology ePrint Archive*.

[UK Payment Systems Regulator, 2018] UK Payment Systems Regulator (2018). Psr discussion paper: Data in the payments industry.

[van Schaik et al., 2022] van Schaik, S., Seto, A., Yurek, T., Batori, A., AlBassam, B., Garman, C., Genkin, D., Miller, A., Ronen, E., and Yarom, Y. (2022). SoK: SGX.Fail: How stuff get eXposed.

[Viand et al., 2021] Viand, A., Jattke, P., and Hithnawi, A. (2021). Sok: Fully homomorphic encryption compilers. *CoRR*, abs/2101.07078.

[Visa, 2024] Visa (2024). Visa reinvents the card, unveils new products for digital age. https://usa.visa.com/about-visa/newsroom/press-releases.releaseId.20686.html. Press Release.

[Wheatley et al., 2016] Wheatley, S., Maillart, T., and Sornette, D. (2016). The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B*, 89(1):7.

[Wüst et al., 2019] Wüst, K., Kostiainen, K., Capkun, V., and Capkun, S. (2019). Prcash: fast, private and regulated transactions for digital currencies. In Goldberg, I. and Moore, T., editors, *Financial Cryptography and Data Security*, volume 11598 of *Lecture Notes in Computer Science*, pages 158–178. Springer.

[Zafrani et al., 2022] Zafrani, E., Mizrahi, T., and Soffer, Y. (2022). Digital shekel - experiment on a distributed platform. Bank of Israel.

Previous volumes in this series

1241 January 2025	Estimating nonlinear heterogeneous agent models with neural networks	Hanno Kase, Leonardo Melosi and Matthias Rottner
1240 January 2025	The granular origins of inflation	Santiago Alvarez-Blaser, Raphael Auer, Sarah M. Lein and Andrei A. Levchenko
1239 January 2025	The use and disuse of FinTech credit: When buy-now-pay-later meets credit reporting	Yanfei Dong, Jiayin Hu, Yiping Huang, Han Qiu and Yingguang Zhang
1238 December 2024	Fiscal stimulus plans and households' expectations	Fiorella De Fiore, Marco Jacopo Lombardi and Albert Pierres Tejada
1237 December 2024	The macroeconomics of green transitions	Gregor Boehl, Flora Budianto and Előd Takáts
1236 December 2024	Savings-and-credit contracts	Bernardus van Doornik, Armando Gomes, David Schoenherr and Janis Skrastins
1235 December 2024	Aggregate debt servicing and the limit on private credit	Mathias Drehmann, Mikael Juselius and Sarah Quincy
1234 December 2024	Targeted Inflation Targeting: Some Evidence and Theory	Boris Hofmann, Cristina Manea and Benoît Mojon
1233 December 2024	Fire Sales of Safe Assets	Gabor Pinter, Emil Siriwardane and Danny Walker
1232 December 2024	Bond supply, yield drifts, and liquidity provision before macroeconomic announcements	Dong Lou, Gabor Pinter, Semih Uslu and Danny Walker
1231 December 2024	Central bank communication and media coverage	Fiorella De Fiore, Alexis Maurin, Andrej Mijakovic and Damiano Sandri
1230 December 2024	Whither inflation targeting as a global monetary standard?	Claudio Borio
1229 November 2024	Through stormy seas: how fragile is liquidity across asset classes and time?	Nihad Aliyev, Matteo Aquilina, Khaladdin Rzayev and Sonya Zhu
1228 November 2024	Retail fast payment systems as a catalyst for digital finance	Giulio Cornelli, Jon Frost, Jonathan Warren, Clair Yang and Carolina Velásquez

All volumes are available on our website www.bis.org.