OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

Working Paper 278

Privacy by design for public digital money

Martin Summer

OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

OeNB Working Paper Series
59 pages

Privacy by design for public digital money

As central banks develop digital currencies for public use, a critical challenge is protecting the privacy of granular data trails that digital payments leave behind. This paper argues that privacy should be a built-in feature of digital money, not a trade-off with crime prevention. Drawing on advances in privacy-enhancing technologies and strategic game-theoretic analysis, it shows that strong privacy and verifiable compliance can coexist. Three design principles are proposed for privacy-protective CBDCs, along with a PET dashboard mapping technologies to system layers.

Authors

Martin Summer
Oesterreichische Nationalbank,
Economic Research Division,
martin.summer@oenb.at

JEL classification
E42, G28, D82, O33

Keywords
CBDC, privacy-enhancing technologies (PETs), economics of privacy

Privacy and compliance are not opposing forces

Integrating a technical design framework with a strategic game-theoretic model, the paper shows that privacy and auditability can be engineered as two separate design dimensions. Modern privacy-enhancing technologies enable cryptographic shielding of routine payments while automatically triggering disclosure above policy-defined thresholds.

Architecture hard-codes incentives

Architectural choices made today will shape institutional credibility for decades. Payment data reveal not only consumption but also vulnerability. Privacy-enhancing technologies have advanced faster than most CBDC design processes could absorb — opening a window to strengthen protections before architectures are locked in

Core PET building blocks are production-ready

Authenticated encryption, zero-knowledge proofs, threshold cryptography, and network-layer protections are used in live systems today. The remaining gap is not scientific uncertainty but engineering maturity, regulatory clarity, and institutional commitment to making privacy a genuine design priority.

Privacy by design for public digital money

Martin Summer*

Abstract

As central banks explore issuing digital currencies for public use, a critical design challenge is how to protect the privacy of the granular data trails digital payments leave behind. While privacy is widely recognised as a goal, policy debates often frame it as a trade-off with crime prevention—limiting ambition and reinforcing legacy design choices that assume privacy and enforcement are fundamentally incompatible. This risks replicating the data practices of commercial platforms in public infrastructure. This paper charts an alternative approach. Recent advances in privacy-enhancing technologies (PETs) now enable both strong privacy protections and verifiable compliance through programmable, rule-based auditability. By embedding such capabilities directly into system architecture, central banks can make privacy a built-in feature of digital money—strengthening institutional trust. Building on recent advances in cryptography and strategic analysis, we offer a conceptual framework that treats privacy and auditability as distinct design dimensions, and distil three design principles for privacy-protective CBDCs that remain compatible with enforcement needs. We also introduce a "PET dashboard" that maps specific technologies to CBDC system layers, highlighting where collaboration across central banks, academia, and industry is most needed.

Keywords: CBDC, privacy-enhancing technologies (PETs), economics of privacy

JEL classification: E42, G28, D82, O33

*Economic Research Section, Oesterreichische Nationalbank. I thank Mike Alsonso, Urs Birchler, Rainer Böhme, Miguel Diaz, Helmut Elsinger, Jon Frost, Geoffrey Goodell, Christian Grothoff, Michel Habib, Markus Knell, Michael Lee, Peter Lindner, Fernando Perez Cruz, Fabian Schär, Hyun Song Shin, Helmut Stix, Nikola Tarashev and Harald Uhlig as well as an anonymous referee and seminar participants at the BIS and at UCL Future of Money Initiative for helpful comments and enlightening discussions. The support of the BIS who hosted the author as a Central Bank Research Fellow from March–May 2025 is gratefully acknowledged. The views expressed here are those of the author and not necessarily the views of the Bank for International Settlements (BIS) or Oesterreichische Nationalbank (OeNB).

Contents

Non-technical summary

Context. Central banks are developing the infrastructure for public digital money at a decisive juncture: privacy-enhancing technologies (PETs) are advancing fast and have already reshaped the research landscape—but have yet to fully inform the policy conversation. Choices taken now will determine whether tomorrow's payments ecosystem is anchored in public trust or in the data-extraction logic that dominates much of today's private sector.

Problem. In several current blueprints, privacy features appear secondary: everyday payments may be visible by default and supervisors may encounter broad data access—conditions that could affect public credibility.

Contribution. Combining recent conceptual and strategic models reveals that privacy and auditability can be treated as two levers—letting central banks

- cryptographically shield low-value retail flows while automatically disclosing transactions that breach a policy threshold;
- sharpen financial-crime detection by filtering out irrelevant data;
- anchor public trust beyond what purely legal safeguards can deliver.

Key insights.

1. Privacy has strategic implications: weak safeguards may reinforce data-driven market power.
2. Core PET components have matured in research and pilots, and some appear close to deployment in constrained settings.
3. Architecture hard-codes incentives: architectural choices made today shape institutional credibility over the long run.

Road-map. The paper

- charts the design space on two axes—how much routine transaction data remains private, and how readily rule-triggered events can be audited;
- overlays that map with a strategic model to identify which combinations remain stable under adversarial adaptation;
- distils three design principles (two levers, calibrated cap, tamper-proof commitment);
- offers a PET dashboard that matches technologies to each CBDC layer and flags readiness gaps.

Implications.

- Project teams can move privacy from patch to core without sacrificing compliance.
- Legislation can avoid locking in early technical assumptions and instead set principles and governance, leaving room for future PET advances.

Bottom line. A privacy-by-design approach looks increasingly practicable, subject to governance choices and engineering trade-offs.

1 Introduction

Like browsing the internet or sending emails, digital payments generate data about users that are partially protected by law but largely remain outside their control. These data are stored and processed by commercial banks, payment service providers, and merchants. What began as a by-product of digital transactions has since become a valuable strategic asset—supporting monetisation, competitive positioning, and granular monitoring and supervision of transactions. Governments see value in such data for national security, law enforcement, and financial supervision. Companies leverage payment data to personalise services, optimise operations, and generate revenue through aggregation, advertising, or resale. Privacy—the ability of individuals and entities to decide how their data are collected, used, and shared—safeguards them from unwanted exposure and exploitation. Yet this fundamental protection remains conspicuously absent from most current digital payment infrastructures.

A new generation of public digital payment systems—especially retail central bank digital currencies (CBDCs)—now offers a chance to change course. Unlike entrenched commercial platforms, CBDCs are still in the design phase and can embed privacy and auditability into their architecture from the start. The challenge is not just technical but conceptual: how to reframe privacy—not as a concession to be granted, but as a structural foundation of institutional trust and system legitimacy.

Already 40 years ago, the computer scientist David Chaum anticipated the need for protections that go beyond legal guarantees. He proposed a cryptographic protocol that could safeguard payment and identity data through technical means—empowering individuals to retain sovereignty over their information and enabling institutions to reduce their exposure to data-related risks (Chaum, 1985).[1] These pioneering ideas have since matured into a rich family of tools—collectively known as privacy-enhancing tech-

[1] Chaum wrote: "*Computerisation is robbing individuals of the ability to monitor and control the ways information about them is used... The foundation is being laid for a dossier society... Uncertainty about whether data will remain secure... can have a 'chilling effect'... The obvious solution for organisations is to devise more pervasive, efficient... record-keeping systems... However, this would exacerbate the problem... and would likely be unacceptable to many.*" (quotation abridged by the author; ellipses indicate omitted text).

nologies (PETs)—that can preserve privacy in payments without sacrificing oversight. Today, these technologies offer a chance to reclaim ground lost in the evolution of digital payments—and to restore privacy as a core design goal.

Privacy in digital payments is often dismissed as a niche concern, yet this fundamentally underestimates what is at stake.[2] Weak privacy protections threaten not only personal autonomy but also economic efficiency and institutional legitimacy. When payment data are easily accessed, exploited, or misused, the consequences extend beyond individual harm: trust erodes, markets distort through data-driven incumbency advantages, and institutional legitimacy suffers. Inadequate safeguards thus pose not only ethical but also strategic challenges for central banks and other public institutions.[3]

Privacy, in this light, is not just a matter of individual rights but a structural element of institutional legitimacy. Policymakers who view privacy as a mere trade-off should reconsider: the reputational costs of a CBDC associated with intrusive surveillance or commercial exploitation of user data could be lasting and severe. Such risks may undermine not only public acceptance of digital money, but also the credibility and effectiveness of central banks in their broader mandates—including monetary policy transmission, financial stability oversight, and crisis response capacity.

This paper argues that privacy should be treated not as a conditional add-on to a retail CBDC, but as a default design state—enforced architecturally and relaxed only under rule-based, auditable conditions. We make the case that privacy and compliance are not mutually exclusive; designs exist that achieve both. It focuses on privacy and auditability as primary design criteria.[4]

While the underlying issues—data visibility, institutional control, and privacy trade-offs—apply broadly across digital payment infrastructures, this paper focuses on retail

[2]See Uhlig et al. (2023) for a recent discussion.

[3]Early economic research already framed privacy as a fundamental feature of money; see Kahn et al. (2005).

[4]We discuss other system criteria—scalability, deployability, and operational resilience—in subsection 6.2 but reference them here only to note feasibility bounds. We use scalability to mean operational performance at target scale (e.g., transaction throughput and user-perceived latency) and deployability to include interoperability and operational resilience. We do not analyse performance engineering in depth; we reference it only to note feasibility bounds for the architectures discussed here.

central bank digital currencies for two reasons. First, CBDC design is still open: unlike entrenched commercial systems, it offers a rare opportunity to embed privacy and auditability directly into the architecture. Second, central banks operate under a public mandate and can credibly commit to safeguarding structural privacy in ways that private actors often cannot. The design space explored here is therefore specific in scope but general in relevance: the underlying tensions are widely shared, but CBDC is the context where they can be addressed by construction.

There is a further reason why CBDC design matters beyond the CBDC itself. Public payment infrastructure sets de facto standards for the broader digital economy: if a non-profit, publicly mandated system tolerates routine data observability, it becomes difficult to demand stronger protections from private actors whose business models depend on data extraction. Conversely, a privacy-by-design CBDC can serve as a benchmark—demonstrating that strong privacy and effective compliance are jointly achievable and raising expectations across the payments ecosystem. The private sector has so far rarely delivered privacy-by-design payment systems at scale, owing to surveillance-based business models, limited interoperability incentives, regulatory uncertainty, and weak user trust in private governance. Cryptocurrency projects have demonstrated technical feasibility in niche settings, but they do not yet constitute evidence of mass-market institutional viability. A credible public alternative could shift these dynamics.

Section 2 outlines what payment systems reveal about users, explains how current systems rely on institutional oversight rather than built-in constraints, and introduces a two-axis framework—privacy and auditability as separate dimensions—based on Auer et al. (2025). This framework offers a more nuanced, technology-informed view than the common framing of privacy versus law enforcement. Section 3 examines the societal and economic risks of weak privacy. Section 4 explains why these risks must be addressed in CBDC design. Section 5 applies these insights to the CBDC debate, showing why many prototypes fall short. Finally, Section 6 presents a strategic design framework, combining the Auer et al. (2025) design space with the Capponi et al. (2025) model. It proposes three design principles and a PET dashboard to guide CBDC development.

2 What payment data reveal, who can access it—and why protection remains fragile

Every time we tap a card, scan a phone, or click "pay now," we generate a data trail. Unlike web browsing or social media posts, digital payment data reflect actual economic decisions: they reveal not only consumption but also preference, priority, and vulnerability. Payments data encompass what we buy, where we are, and when we act—revealing how we live. Beyond the explicit content of transactions, metadata—frequency, timing, and patterns—independently reveal routines and relationships. This combination—actual economic decisions plus behavioural metadata—makes payment data uniquely sensitive and uniquely valuable. And yet, in today's digital economy, this information is routinely collected, aggregated, and monetised, often under frameworks that emphasise compliance and business efficiency rather than user agency. Central banks designing public digital payment systems must therefore understand not only the technical landscape, but the structural risks inherent in routine data collection.

2.1 What data are revealed by digital payments?

Despite jurisdictional and technological differences, the types of data disclosed by digital payments are remarkably consistent across systems. They can be grouped into three broad categories:[5]

- **Personal identifiers:** These include names, user or account IDs, phone numbers, email addresses, and, in some cases, physical addresses. They link transactions to identifiable individuals.
- **Payment method details:** This refers to information required to process the payment—such as bank account numbers, card credentials, digital wallet addresses, or payment app identifiers.

[5]For this categorisation, I draw on European Data Protection Board (2020); PCI Security Standards Council (2022); Vives et al. (2024); Jones (2024). A structured overview of data elements in typical payment scenarios is provided in Appendix A.

- **Transaction data:** This includes the amount, date and time, merchant category, location of the transaction or device. In some contexts, merchant-side commercial data—such as itemised purchase details—may be linked to the payment record, although such data are not part of the payment message itself.

These core data elements form the routine substrate of digital payments—collected, transmitted, and stored with each transaction.

Beyond these core categories, digital payments also generate rich metadata: contextual information about how, when, and with whom payments occur. Patterns in metadata can reveal social connections, economic dependencies, or personal habits. For example, repeated transactions with the same counterparty can suggest close personal or professional ties; night-time purchases at pharmacies may indicate health concerns; recurring transfers can reflect subscription services, rent, or loan repayments.

Some data are revealed not only about the person initiating a transaction but also about third parties—such as recipients—who may have no direct relationship with the payment service provider. These are sometimes referred to as silent party data. For instance, if Alice sends a payment to Bob via a third-party platform, Bob's account information may be processed even though he never interacted with the service himself. In such cases, the data subject may be unaware of the processing and unable to exercise any control over it.[6]

Taken together, digital payments generate a rich and often involuntary data trail. Even innocuous individual data elements become powerful through aggregation: revealing sensitive personal characteristics, forecasting behaviour, and enabling automated decisions about creditworthiness, insurance eligibility, or access to services. Understanding what data are generated—and who is affected—is essential before we can assess what protections are needed and what technologies can deliver them.

[6]See PCI Security Standards Council (2022), sections 4.1–4.3.

2.2 How are payments data protected today?

Digital payments are governed by a patchwork of legal and technical protections that vary widely in scope, enforcement, and effectiveness. While many frameworks aim to safeguard sensitive financial information, they differ across jurisdictions and are often limited in their ability to guarantee meaningful user control.

In the European Union, two complementary instruments define the regulatory baseline: the General Data Protection Regulation (GDPR) and the Revised Payment Services Directive (PSD2). The GDPR lays out general principles for personal data protection—such as lawfulness, transparency, data minimisation, and accountability—and grants individuals a set of enforceable rights, including access, rectification, erasure, and objection (European Parliament and Council, 2016). PSD2 supplements this framework with rules specific to payment services, including requirements for strong customer authentication and limited data access for third-party providers based on user consent and purpose limitation (European Parliament and Council, 2015).[7]

Beyond Europe, the regulatory landscape is more fragmented. In the United States, privacy protections are sector-specific and vary by state, leading to gaps and inconsistencies (White & Case LLP, 2024). China's Personal Information Protection Law mirrors aspects of the GDPR but grants broad access to state authorities (Cremeers, 2021). India's 2023 Digital Personal Data Protection Act provides individual rights on paper. However, it includes expansive exemptions for government use. (PRS Legislative Research, 2023; Burman, 2023) Emerging data regimes in Latin America and Africa face challenges not only in scope but also in enforcement capacity (Antequera, 2023; Data Protection Africa, 2023).

Even where strong laws exist, protections often stop at the border. In today's interconnected payments infrastructure, data routinely cross jurisdictions—passing through processing centres, cloud services, and intermediary banks that may operate under weaker legal or institutional safeguards. Moreover, legal regimes rarely address the full scope of

[7]For interpretative guidance on PSD2 data minimisation, see EDPB Guidelines 06/2020, Sections 4 and 6.

metadata—such as transaction frequency, timing, or merchant category—which can be highly revealing and are central to profiling and behavioural inference.

In what follows we focus on privacy-rights instruments (e.g., GDPR) and architectural privacy. Card-scheme security frameworks address credential compromise and fraud control and are outside scope here, even though authorisation data are personal data in the legal sense.[8]

2.3 Why current payment data protection is fragile

Legal and procedural data protection arrangements rely heavily on institutional trust and discretionary access. Most systems implement what Auer and Böhme (2025) call "soft privacy": legal and procedural controls that aim to limit access but allow for override under specific conditions. Similarly, enforcement relies on "soft auditability": access to stored transaction data following ex post legal process. But this model is under strain.

Most retail payment systems verify a customer's identity at onboarding (KYC) and then transact under persistent identifiers (e.g., IBAN, wallet ID). From that point forward, the system tracks all activity under that identifier.

Under the risk-based approach, ongoing monitoring looks for patterns and risk indicators over time. The transaction format itself does not encode whether a payment is licit or illicit.[9] Consequently, pre-emptive detection would require broad surveillance of many users who present low risk—raising necessity and proportionality concerns that regulators increasingly flag.[10]

In practice, enforcement relies on transaction metadata—such as value, frequency, counterparties, and merchant category—but these signals do not reveal intent. Illicit actors can easily mimic legitimate behaviour, opening many wallets and fragmenting large

[8]Authorisation tokens (PANs, expiry, service codes, tokenised credentials) can be linked to individuals; their protection is primarily an operational security matter, not a privacy-rights guarantee.

[9]For example, two monthly SEPA credit transfers of €1,000 have identical fields (amount, date, debtor/creditor IBAN, remittance text). One is rent to a landlord; the other is a transfer to a money-mule account. Ex ante, both records look the same; only *pattern-level* signals (counterparty history, networks, velocity, jurisdiction) or specific flags elevate the latter for review under the risk-based approach.

[10]See European Data Protection Board letters on AML/CFT data sharing, noting large-scale processing risks and the need for strict necessity and proportionality assessments.

transfers into many small payments (a technique known as smurfing). This dynamic overwhelms enforcement capacity, which does not scale as fast as evasion tactics. Investigations then retreat into selective, ex post sampling—dependent on subpoenas and rule-triggered alerts to re-identify the user behind a suspicious alias. As Capponi et al. (2025) show, such regimes enable bad actors to scale faster than supervisory systems can adapt.

Current arrangements also carry a distinct cybersecurity liability. Systems that store transaction data in plaintext—or encrypted with keys routinely available to operators—create centralised honeypots: high-value targets whose compromise exposes the payment histories of entire populations. This is precisely the "dossier society" risk that Chaum identified four decades ago—not only as a threat to personal autonomy, but as a structural security vulnerability inherent in centralised plaintext record-keeping. Hard privacy architectures reduce this attack surface by ensuring that no single point of compromise yields mass plaintext access. They also limit damage from insider abuse, since routine operators never hold data in readable form. This does not eliminate risk: implementation bugs, compromised wallets, threshold-governance failures, traffic analysis, and flawed hardware assumptions introduce new attack surfaces. The correct comparison is therefore not between architectural privacy and perfect security, but between two risk distributions—one in which breaches silently expose everything, and one in which exceptional access is constrained, multi-party, and auditable.

The mismatch between enforcement architecture and adversarial tactics, compounded by these cybersecurity vulnerabilities, does not imply that privacy and oversight are incompatible. Rather, it reveals that the prevailing trade-off framing conflates two distinct dimensions that can, in principle, be engineered separately. Auer et al. (2025) address this confusion with a two-axis framework in which privacy and auditability are treated as two different design levers—each implementable through institutional mechanisms (soft) or cryptographic enforcement (hard). Before applying that framework, it is useful to clarify the key terms that structure the remainder of this paper. Privacy is the broader objective: limiting who can observe, infer, and use payment-related information. Anonymity

is only one possible means of achieving privacy—it refers to the inability to link a transaction to an identifiable person. Pseudonymity offers weaker protection: identifiers are masked in ordinary operation, but transactions may remain linkable and, under certain conditions, re-identifiable. Confidentiality, by contrast, refers to shielding the content of transactions—amounts, counterparties, metadata—from unauthorized observation, regardless of whether identities are known. Hard privacy, as used in this paper, therefore need not mean full anonymity in every context; it means that routine data access is constrained by architecture rather than left to institutional discretion.

2.4 A two-axis map of privacy and auditability

Figure 1, adapted from Auer et al. (2025), maps payment architectures along two axes to move beyond the simplistic "privacy-versus-security" framing and reveal a wider design space:

1. **Privacy (vertical).** We ask whether personal data are operationally readable to routine actors—either in plaintext or encrypted but with keys held centrally—or whether they are cryptographically protected with access limited to specific, authorised circumstances.[11]

2. **Auditability (horizontal).** We ask how personal information becomes accessible to authorised parties *when conditions are met*. Auditability is soft when privileged staff can disclose data at their discretion (subject to policy and oversight). It is hard when disclosure is governed by pre-defined, cryptographically enforced rules—such as threshold release requiring multiple independent parties, with all access attempts logged in tamper-evident audit trails.

Rather than a single trade-off, the matrix yields four qualitatively distinct regimes—each with its own implications for user behaviour, enforcement capacity, and compliance

[11] *plaintext* here is shorthand for "effectively readable to those operating the system." Data encrypted with keys that are routinely available to operators offers little additional privacy relative to plaintext; what matters is key governance—who holds keys, under what process they can be combined or used, and what evidence is produced.

costs.

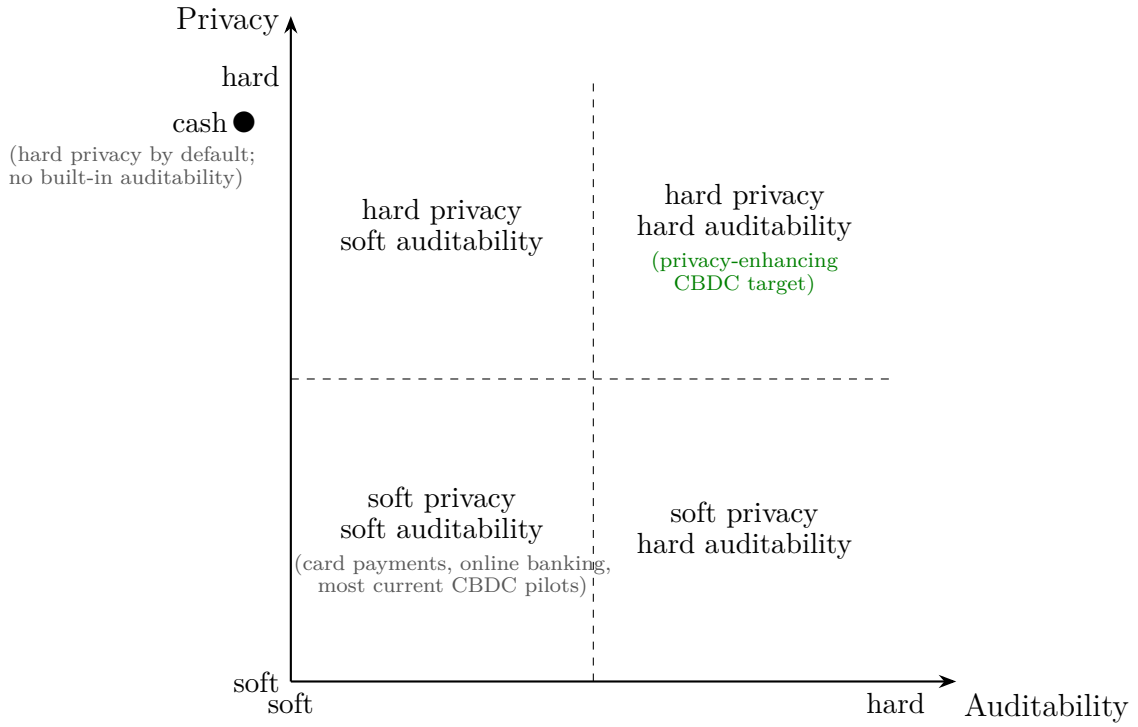

Figure 1: Design space for payment systems (adapted from Auer et al., 2025). Designs navigate a plane, not a line. Examples are placed indicatively; real implementations may straddle quadrant boundaries. Cash sits outside the grid: it provides hard privacy by default but no built-in auditability. *Note.* Axes are orthogonal for exposition; they denote two requirement dimensions. Real implementations will exhibit couplings and cross-effects, which PET compositions and governance mechanisms must manage.

Let us now place some familiar and hypothetical examples in this grid. Cash sits outside the grid: it offers hard privacy by default, but virtually no built-in auditability. Any forensic traceability must come from external means—such as banknote tracking or surveillance tools—not from the payment system itself.

Card payments and online banking fall into the bottom-left quadrant: both privacy and auditability are soft. Transaction data are typically encrypted at rest and in transit, but the institutions that store them also control the decryption keys. This means data can be accessed and turned into clear text internally, subject to institutional policies.

Auditability depends on ex post access by authorised personnel—typically via legal request or regulatory examination. Privacy, therefore, relies on governance rather than

architectural constraint.

A hypothetical CBDC design based on a two-tier architecture—where users access wallets through private intermediaries, while the central bank sees only pseudonyms—also lands in this quadrant. By "pseudonymous," we mean account-level identifiers not bound to a natural person in the ledger view. Importantly, even without KYC the resulting *transaction graph* (who pays whom, when, how often, and in what approximate patterns) is typically observable to some parties and can be highly identifying over many payments. Graph-based clustering and auxiliary-data joins have re-identified users at scale in multiple domains; such methods also introduce the risk of *false associations* when heuristics misclassify flows. Intermediaries handle onboarding and transaction processing, and retain full visibility over user flows.

Although the central bank receives only pseudonymous identifiers, the identity mapping is stored and can be accessed under legal procedure. Privacy in this setup is institutionally granted, not cryptographically protected. Auditability remains soft: authorities must request access on a case-by-case basis rather than relying on system-enforced disclosure rules.

This mapping reveals a salient pattern: many existing systems cluster in the bottom-left quadrant, where privacy and auditability rely mainly on institutional discretion rather than architectural guarantees. While this arrangement may seem pragmatic, it exposes users and institutions to risks—explored in Section 3 that extend far beyond individual privacy preferences. Understanding these systemic risks is essential to appreciating why CBDC design represents both an opportunity and an institutional imperative.

The remainder of the paper applies this two-axis lens to privacy-centred CBDC architecture. Other system criteria (e.g., performance at scale) are referenced only insofar as they bound feasibility.

3 Why weak privacy protection in payments is a problem

The systemic risks of weak payment privacy are often underestimated. Digital payments generate rich data trails that are often presented as sources of convenience and innovation. In jurisdictions such as the European Union, these data are subject to comprehensive protection laws that are increasingly adopted elsewhere. Yet even under strong legal regimes, serious privacy risks persist.

To be clear, transparency in payment data is not without value. Access to transaction records supports fraud detection, law enforcement, supervisory learning, and some forms of economic research. It can also enable convenience-enhancing services that many users appreciate. These benefits are real, but they do not by themselves justify routine mass observability—because the associated harms are structural, externalised, and difficult for individuals to discipline through private choice.

This section explains why legal safeguards and user benefits from data sharing do not by themselves resolve concerns about payment data privacy. Legal protections and data-driven conveniences must be weighed against three structural risks: (i) surveillance and overreach by public and private actors; (ii) negative externalities that affect even non-consenting parties; and (iii) market-power dynamics that entrench incumbent advantages.

These deeper structural problems are often overlooked—but they can outweigh the actual or perceived benefits of data-driven services, making stronger architectural safeguards not merely desirable but strategically necessary.

Public-sector surveillance. Government access to payments data often begins with narrow objectives but can escalate rapidly. The U.S.-led Terrorist Finance Tracking Program (TFTP) provides a cautionary example. Following the 9/11 attacks, it granted sweeping access to SWIFT transaction data[12] — the messaging system that underpins

[12]SWIFT (Society for Worldwide Interbank Financial Telecommunication) is a member-owned cooperative that provides secure messaging infrastructure for over 11,000 financial institutions worldwide. It facilitates communication of payment instructions but does not process or hold funds itself.

international bank transfers.

Though intended for counter-terrorism, the programme sparked sustained criticism in Europe over its expansive reach and potential for mission creep—ultimately requiring a negotiated agreement to limit scope and establish oversight mechanisms.[13] More recent examples illustrate how digital payment infrastructure can be repurposed in moments of political tension. During Canada's 2022 truckers' protests, the government invoked emergency powers to freeze protestors' accounts without court orders. Similar tactics were reported in the 2019–2020 Hong Kong protests, where accounts linked to pro-democracy activists were reportedly suspended or frozen. Without robust guardrails, even democratic states can use payment infrastructure to exert political pressure or suppress dissent—raising questions about perceived institutional neutrality and legitimacy.

Private-sector control. Commercial platforms pose two distinct risks to data privacy and financial access: vulnerability from centralised data storage, and opacity in access control.

The 2017 Equifax breach exposed the personal data of over 140 million individuals, highlighting the systemic insecurity of large-scale data repositories.[14] Even when privacy protections exist on paper, they can collapse through security failures—especially when data are held in a single place and protected only by institutional procedures, creating precisely the kind of centralised vulnerability discussed in Section 2.

Separately, companies like PayPal and Venmo have frozen user accounts with little explanation, often relying on opaque policies or automated filters. These suspensions have disproportionately affected politically sensitive users or those caught in poorly explained compliance procedures.[15] The structural issue here is governance: access to financial services is increasingly mediated by private actors, whose discretionary rules are governed by contract rather than enforceable rights.

These examples may fall outside the narrow legal definition of "payment service

[13]See Wikipedia Contributors (2024) for an overview.
[14]U.S. House Oversight Committee (2018).
[15]See Electronic Frontier Foundation, "22 Rights Groups Tell PayPal and Venmo to Shape Up," 2022.

providers," but they reveal a broader vulnerability. When financial infrastructure—including account access, transaction metadata, and identity management—is governed by internal policy rather than architectural constraint, the risk of exclusion or misuse rises. In such environments, privacy is not just about hiding flows from third parties. It is about preventing the arbitrary withholding, monetisation, or repurposing of financial access. Embedding privacy at the infrastructure level—through cryptographic guarantees and rule-bound disclosure—offers a more robust safeguard than reliance on institutional goodwill, whether the provider is a fintech platform, a bank, or a vertically integrated tech firm.

Externalities and market power. Individual decisions to share payment data may appear benign, but the social consequences are not. Research shows that data disclosure generates externalities: what one user reveals can expose others who never consented to be profiled (Bergemann et al., 2022; Choi et al., 2019; Acemoglu et al., 2022). For example, if many users in a demographic group share transaction patterns, algorithms can infer the likely behaviour of holdouts—effectively nullifying their privacy choice. This dynamic leads to systematic over-collection and concentration of payment data. Dominant platforms exploit data flows to reinforce market power—through personalized pricing, algorithmic exclusion, and targeted offers—while entrants struggle to compete without access to comparable datasets (Acquisti et al., 2024; Croxson et al., 2023).[16] The result is a self-reinforcing cycle: incumbents accumulate data advantages that raise barriers to entry and reduce competitive pressure.

Recognising these dynamics, the BIS has argued that well-designed CBDCs could serve as a public-sector anchor in digital payments, restoring competitive neutrality in platform-dominated markets (Bank for International Settlements, 2022). But this promise hinges on privacy. Without robust protections, public money risks reproducing the same extractive dynamics it seeks to displace. Dominant players will retain their advantage, while privacy-respecting alternatives remain marginalised. Conversely, strong

[16]See also BIS (2019) and Zuboff (2020) for broader treatments of how behavioural data, including payment activity, are monetised across sectors.

privacy-by-design could shift the competitive landscape—reducing entry barriers, discouraging surveillance-based monetisation, and restoring trust in digital infrastructure.

4 Why all this matters for the design of retail CBDC

The question of privacy in central bank digital currency (CBDC) design is not merely a matter of user preference or technical convenience; it is a strategic design decision with far-reaching implications for public trust, adoption, and long-term credibility. If users suspect that their payment data could be exposed, monetised, or otherwise misused, they may turn to alternatives—such as cash (limiting CBDC adoption), foreign CBDCs (undermining monetary sovereignty), or private digital currencies perceived to offer stronger privacy (shifting activity beyond regulatory oversight). Given the established business models of many digital platforms, these concerns are not unfounded. In practice, uptake often hinges on the assessments of opinion leaders in civil society—consumer advocates, privacy NGOs, investigative journalists, technologists, and academic experts—who scrutinize design choices and shape public understanding. A negative assessment from such actors can prove difficult to reverse, regardless of subsequent technical improvements. Even where privacy provisions have been carefully negotiated—as in the case of the digital euro or digital pound—there remains a risk that these communities will judge them insufficient, thereby influencing public perception and adoption.

Empirical studies consistently find that individuals often trade privacy against convenience, cost, or usability—a pattern sometimes described as the "privacy paradox."[17] Survey-based evidence, including our own empirical research on CBDC adoption potential,[18] shows that privacy valuations are heterogeneous and, on average, modest in size. But this apparent indifference need not imply that privacy protections are unimportant. In environments where data sharing creates externalities, individual behaviour may rationally underweight long-term or system-level risks. Moreover, when privacy erosion is

[17] See e.g. Acquisti et al. (2024) for a comprehensive review.
[18] See Abramova et al. (2023) and Elsinger et al. (2025).

gradual and distributed, individuals may fail to perceive the cumulative threat until it is difficult to reverse. Privacy in this context becomes a public good. Its value lies not only in personal concealment, but in constraining power asymmetries, limiting data-driven market dominance, and preserving institutional trust. Individual users cannot secure these systemic benefits through their own choices—they require architectural guarantees.

Crucially, weak privacy measures raise not only fears of state surveillance; they also heighten concerns that sensitive transaction data could be exploited by the private intermediaries that distribute and manage a CBDC. In two-tier architectures, these intermediaries may have operational access to transaction flows, creating opportunities for commercial profiling, targeted marketing, or data resale — potentially replicating the very dynamics that erode trust in private payment platforms.

The legitimacy of central banks rests on more than monetary and financial stability; it also depends on upholding public trust in the integrity of money. If users believe that a CBDC exposes their payment activity to surveillance or misuse— whether by government institutions or profit-driven intermediaries—trust may erode rapidly, drawing criticism from privacy advocates and civil-society organizations, and ultimately eroding public confidence. Such erosion would jeopardize domestic take-up, in particular among demographically privacy-sensitive groups. Armantier et al. (2024) find in a representative U.S. survey that women and older citizens demand substantially higher compensation for sharing transaction data, suggesting that privacy-light designs could deter these groups from adoption. At scale, this could undermine financial sovereignty by pushing transactions toward less accountable payment rails. And once a precedent is set for outside access to granular CBDC data, the same leverage could be used to pressure the central bank in other policy areas—ranging from credit allocation to asset purchases—chipping away at its operational independence. Data access, in other words, creates institutional vulnerability that extends beyond the payments domain.

Architectural privacy protections are necessary but not sufficient for public trust. Even cryptographically enforced guarantees may fail to convince privacy-conscious users if the underlying system remains opaque. Credibility requires more than sound engi-

neering: it demands public documentation of what data are hidden, what is logged, and under what precise conditions disclosure can occur; independent security audits with published findings; open-source verification tools where feasible; and sustained engagement with external experts, including the civil-society actors whose assessments shape public perception. Some users will remain sceptical regardless—but a verifiable, transparently governed architecture narrows the gap between technical reality and public understanding far more effectively than assurances of institutional good faith alone.

Yet the same risks reveal an opportunity. By establishing credible privacy guarantees from the outset, central banks can set a new benchmark for trust, transparency and institutional resilience in public digital infrastructure. Strengthening privacy standards is therefore not merely ethical—it is strategically essential. The question, then, is why current CBDC architectures fall short of this goal—and what design alternatives exist. We turn to these questions in Section 5.

5 Why privacy remains under-delivered in current CBDC architectures

Despite widespread claims that privacy is a core design objective, many current CBDC designs fall short of what appears technically feasible today, given stated objectives. Central banks often assert that protecting user privacy is essential for public trust and adoption. Yet a closer look at their design choices reveals a recurring pattern: Privacy is often constrained by prevailing compliance and oversight considerations. These projects frequently frame privacy and auditability as mutually exclusive, implying that any enhancement of one must come at the expense of the other.

This framing obscures the real possibility of designing payment systems that support both strong privacy and targeted transparency. Technological advances—especially in privacy-enhancing technologies (PETs) — have made such integration increasingly feasible (Auer et al., 2025; Asrov and Samonas, 2021; Arora et al., 2025; Vives et al., 2024).[19]

[19]These references refer to papers authored by researchers at major central banks and international

Clinging to the assumption of a necessary trade-off reflects an outdated view—one that risks becoming embedded in the very architecture of future money.

5.1 Why current CBDC designs fall short on privacy

A review of actual design proposals—listed and briefly described in Table B1 in the appendix — reinforces this diagnosis. Most projects cluster in the soft privacy, soft auditability quadrant of the design space introduced in Section 2. Most CBDC design proposals adopt some form of custodial model where private-sector intermediaries play a key role in hosting digital wallets or CBDC accounts as well as for distribution and transaction processing. In these frameworks, privacy considerations are typically limited to the relationship between end users and the central bank. Some designs—such as the Bahamas' Sand Dollar—adopt a model where the central bank keeps only pseudonymous wallet IDs while the full name-lookup table stays with private payment service providers. Law enforcement can access this mapping only under court warrant.

However, transaction data often remain fully visible to distributing intermediaries or payment service providers, with few provisions for protecting user privacy at that level. Even pseudonymised or token-based systems—such as Sweden's e-krona or Uruguay's e-Peso—maintain full traceability and grant authorities the ability to re-identify users under broadly defined conditions. The architectural capability for mass surveillance remains intact; only policy constraints limit its use.

What these designs share is not just a technical pattern, but a prevailing mindset: privacy is treated as a variable to be weighed, not a principle to be embedded. This contrasts with hard privacy approaches, where architectural constraints—not institutional discretion—enforce data minimisation and limit surveillance by design.

The persistent reliance on soft privacy models means that many current CBDC systems fall short—not for lack of stated intent, but relative to what is now technologically achievable.

institutions. For further proposals on how new cryptographic protocols can enable both privacy and regulatory auditability, see also Buterin et al. (2024).

Public documents on recent pilots emphasise institutional controls; few disclose hard privacy properties beyond policy and process, and several designs leave transaction metadata broadly visible. By contrast, Appendix C lists PET compositions—combinations of cryptographic techniques—that can deliver privacy for routine payments with rule-based auditability. For example: credential-based selective disclosure (prove attributes without revealing identity) + bounded zero-knowledge proofs (prove compliance without revealing amounts) + threshold release (require multiple parties to authorize disclosure) + network-layer protections (hide transaction graphs from bulk observers). What remains to reach production is interoperability testing, post-quantum readiness, and verifiable governance. We treat these as implementation tasks rather than showstoppers and point to the specific gaps in Appendix C.

5.2 Why partial workarounds are not enough

While many central banks publicly stress their commitment to privacy, official documents and pilot reports often reveal a more ambivalent picture. Some institutions appear to acknowledge—albeit implicitly—that their proposed systems fail to deliver robust privacy. Instead of addressing this shortfall at the structural level, they introduce narrow exceptions or features aimed at limiting visibility in specific contexts. Examples include low-value offline payments with enhanced anonymity, tiered account structures with minimal identification for basic access, or privacy zones defined by transaction size or purpose.

These efforts can mitigate specific concerns but often retain wide default visibility. In most cases, they rely on traceable identifiers, centralised visibility, and regulatory override. These measures provide limited improvements, but may fall short of addressing the structural visibility embedded in current designs. The deeper issue remains unaddressed: data visibility is assumed by default, and privacy must be justified, restricted, and explicitly carved out.[20]

[20]The Big Brother Watch report (Big Brother Watch, 2023) provides extensive documentation of this pattern. Whether in the form of "private shekels" in Israel, offline functionality for smaller transactions the digital euro, or reduced KYC thresholds in the Nigerian eNaira, the shared feature is clear: these are concessions to privacy, not structural commitments. As computer security experts have emphasised,

A notable example is the European Central Bank's plan to support offline functionality in the digital euro. This would allow low-value transactions to be conducted using secure hardware, aiming to replicate some of the privacy benefits of cash (European Central Bank, 2022, 2023). This proposal signals a sincere effort to address public concerns and acknowledge the importance of privacy and user autonomy.

Yet offline modes face an inherent constraint: they can provide bounded local privacy and resilience, but they do not remove the need for eventual synchronisation with the central ledger. Once a device reconnects, transactions must be reconciled—reintroducing the same data-visibility questions that apply to online payments. The privacy gain is therefore temporary and local unless the reconciliation path itself is designed with strong protections. Researchers have flagged further limitations. As Grothoff and Dold (2021) and Chaum et al. (2021) argue, restricting privacy to the offline tier risks framing it as a constrained exception rather than a general design norm. Maintaining large-scale offline capabilities also raises difficult technical challenges—ranging from device security to resilience, auditability, and fraud prevention.[21]

Beyond the technical challenges, tying strong privacy to physical co-presence creates an equity problem.[22] In dense urban areas, reaching a point of sale or ATM may take minutes; in rural settings it can mean kilometre-scale travel, limited opening hours, and queueing. People with reduced mobility, care obligations, shift work, or low access to public transport face higher effective costs. In practice, this makes privacy *location-dependent*: those farther from retail centres or cash infrastructure pay more—in time, money, or both—for the same level of privacy.

Moreover, privacy restricted to offline use excludes remote commerce (e.g., online

meaningful protection cannot be achieved through isolated features appended to inherently transparent systems. Privacy must be addressed as a system-wide design problem—an idea developed further in the next section.

[21]During network partitions, a system must trade off consistency and availability. Retail payment systems typically prioritise consistency (no double-spend, atomic settlement) and accept bounded unavailability. Offline modes carve out a separate regime with different trust assumptions and higher fraud tolerance; see Grothoff and Dold (2021) for a discussion with reference to the CAP theorem.

[22]I have to thank Rainer Böhme for pointing out this often overlooked limiting aspect of delivering privacy through offline digital instruments.

purchases, bill payments at a distance), which are precisely the contexts where many users need privacy most. If privacy is available only where physical co-presence is feasible, it risks becoming a convenience for some users and a costly privilege for others.

A design objective is *functional equivalence*: routine low-value payments should afford comparable privacy *regardless of geolocation or mobility*. Offline modes can complement this (e.g., resilience, power-outage tolerance), but relying on offline as the primary privacy channel creates systematic access asymmetries that policy makers typically aim to avoid.

These critiques do not challenge the motivations of providing privacy through an offline functionality but suggest that meaningful privacy requires more than localised features. It calls for a broader architectural shift—one that treats privacy not as an optional layer, but as a baseline condition across the system.

Despite ongoing refinements, most CBDC designs continue to rely on fragmented models of privacy. Selective bolt-on privacy features to otherwise transparent systems may help defuse political pressure, but they fall short of confronting the structural visibility embedded in current proposals. The default assumption remains: the system is transparent, and privacy must be negotiated.

This architectural inertia reflects not only technical conservatism, but also widely shared central-bank priorities around operational resilience and continuity with existing two-tier systems.[23] A shift in perspective is needed. Rather than selectively granting privacy under special conditions, future CBDC systems must treat privacy as the default state and visibility as the exception.

Achieving this transition likely requires more than updated policy language or incremental technical adjustments. As Arora et al. (2025) argue, effective privacy must be embedded across the full architecture—from onboarding and identity management to transaction processing, compliance, wallet design, and the network/transport layer. Without network-layer protections, transport metadata (e.g., IP/routing information, timing, packet sizes) enable bulk observers to reconstruct large parts of the transaction graph—

[23]See Auer et al. (2022) for a policy-oriented synthesis of CBDC research, which highlights the "triple imperative" of competition, data privacy, and system integrity as key goals that current design efforts attempt to balance.

even if the transaction content itself is encrypted. In other words, hiding who pays whom requires protections beyond encrypting amounts. It seems clear that central banks cannot do this alone. Delivering privacy at scale is primarily an integration and governance challenge spanning cryptography, systems, networks, economics, and oversight.

Meeting this challenge will require central banks to engage more deeply with the research community and expand their technical partnerships. In many cases, it will mean investing in new forms of collaboration beyond traditional IT domains. Building truly privacy-enhancing public payment systems is not only a question of political will—it is also a matter of institutional learning and long-term capacity.

Yet the gap between what is already feasible in the lab and what most pilot projects have attempted in practice is widening. A growing body of academic prototypes now demonstrates retail-scale payments that combine strong cryptographic privacy with rule-based disclosure and audit—even under throughput and latency constraints comparable to existing card networks. While none of these systems is "drop-in" ready for a national rollout, they show that the technical ingredients for hard privacy plus hard auditability already exist. The challenge for central banks is therefore less about inventing new cryptography and more about integrating proven building blocks into an accountable, production-grade architecture.

6 From diagnosis to design: Principles and practical options for privacy-based CBDC architecture

This integration challenge requires a strategic framework that combines technical feasibility with institutional viability. Two influential frameworks illuminate different aspects of the challenge, yet they operate in distinct conceptual spaces:

The Auer et al. (2025) framework positions payment architectures in a technical design space with two axes: privacy and auditability. For each dimension, enforcement can be rule-based and cryptographic ('hard') or institutional and discretionary ('soft').

The Capponi et al. (2025) three-player game operates in a strategic space with different coordinates: identity privacy (real names versus pseudonyms at onboarding) and transaction privacy (clear flows versus cryptographic shielding). Here, auditability is not an independent strategic choice but emerges implicitly from how those privacy levers are set.

To harness both insights, we must translate between these two spaces. The strategic game identifies which privacy configurations remain stable when users, criminals, and supervisors optimise their behaviour. It also assesses the welfare properties of different equilibria. But its coordinates don't directly correspond to the technical implementation axes. By mapping the Capponi et al. (2025) strategic choices onto the Auer et al. (2025) grid, we can identify not merely what central banks can build technically, but what they should build to achieve stable, welfare-maximizing outcomes.

The translation requires careful attention to how each strategic lever affects the technical dimensions. The Capponi et al. (2025) model has two binary choices—identity privacy (on/off) and transaction privacy (on/off)—that must be mapped onto the Auer et al. (2025) continuous axes of privacy and auditability hardness.

The mapping rules are:

- **Transaction privacy → Privacy axis**: Switching transaction privacy on (encrypted flows) or off (clear-text flows) directly corresponds to moving up or down the Auer et al. (2025) privacy axis.
- **Identity privacy → Auditability axis**: Switching identity privacy affects auditability, but the mapping is conditional. Real-name onboarding enables hard auditability only if the identity registry itself is protected by the same rule-based cryptography that governs transaction disclosure—for example, requiring threshold authorization to access the name-lookup table. If real names can be accessed through institutional discretion rather than cryptographic rules, auditability remains soft regardless of the onboarding model.

This mapping generates four possible combinations, summarized in Table 1. Each cell represents a different architectural approach, combining the Capponi et al. (2025)

strategic levers with their corresponding positions on the Auer et al. (2025) grid. The color coding anticipates the welfare implications: only one combination—hard privacy with hard auditability—delivers both strong user protection and enforceable compliance, while the others suffer from either privacy erosion or enforcement failure. This welfare assessment, however, rests on specific modeling assumptions that we unpack below.

Table 1: Mapping Capponi et al. (2025) levers onto the Auer et al. (2025) design grid

Transaction privacy (privacy axis)	**Identity privacy (auditability axis)**	
	Pseudonymous onboarding *(soft auditability)*	Real-name onboarding *(hard auditability)*
Encrypted flows *(hard privacy)*	Hard P + Soft A	Hard P + Hard A
Clear-text flows *(soft privacy)*	Soft P + Soft A	Soft P + Hard A

Notes: Colours are consistent with Figure 2: = hard P + hard A (target region), = hard P + soft A (dominated in the stylised model), = soft P + hard A, = soft P + soft A. The two-axis map is one coherent interpretation; welfare claims are under the stylised assumptions stated in the text.

Strategic outcomes The welfare interpretation —which favors hard privacy with hard auditability—relies on three modeling ingredients from Capponi et al. (2025). These assumptions are important because different assumptions could yield different rankings.

First, *unique participant identification* (UPI) consolidates accounts at the *user* level. This prevents scaling illicit activity by splitting volume across many accounts (smurfing), which is feasible under identity-privacy but not under UPI.[24] Second, enforcement is modelled via a detection function where the probability of catching illicit activity increases with per-user volume. This assumption makes structuring/smurfing attractive when UPI is absent (criminals can split flows across many identities) and makes volume thresholds meaningful when UPI is present (one identity, all flows consolidated). Third, the model

[24]We use UPI in a functional sense: a one-to-many mapping from real-world identity to accounts is precluded by system design. This is distinct from "KYC" as a compliance process. The mechanism relevant for the welfare result is the *consolidation* implied by UPI, not the breadth of personal data collected.

considers how illicit actors adapt their strategies. If enforcement technology is static, optimal privacy might involve complex rules with multiple transaction-value 'windows' where privacy applies. But if illicit actors can innovate—developing new evasion techniques incrementally or through technological breakthroughs—they will exploit any such complexity. In that case, optimal privacy collapses to a simpler rule: a single threshold below which privacy applies, with automatic disclosure above it. See the baseline model (static technology frontier) yielding potentially disjoint transaction privacy intervals and the innovation extensions that induce threshold rules. These ingredients together explain why real-name onboarding combined with transaction privacy up to a calibrated threshold can dominate in the stylised environment. Outside these assumptions, rankings can differ.

Under the stylised environment of Capponi et al. (2025) — with user-level consolidation (UPI), detection that rises in per-user volume, and bad-user technologies as specified—the welfare-optimal region lies in hard privacy with hard auditability. This is implemented as real-name onboarding plus transaction privacy up to a calibrated threshold. Game-theoretic analysis yields several stable equilibria, but a benevolent designer with commitment power would select the welfare-maximizing configuration: real-name onboarding combined with selective transaction privacy. This means encrypted flows below a carefully calibrated threshold—set between current cash anonymity limits and typical retail spending—with automatic disclosure above it.

When projected onto the Auer et al. (2025) canvas, this optimal equilibrium falls within the upper-right quadrant (hard privacy, hard auditability). Crucially, the welfare-maximizing point should be interior to this quadrant, not at its corner. Pure hard enforcement without any discretionary fallback requires the disclosure rules to be specified with surgical precision—a technical challenge that risks paralyzing investigations when edge cases arise.

The regulatory backlash against Tornado Cash illustrates this risk. The cryptocurrency mixer provided cryptographic anonymity with no mechanism for authorized disclosure. When it was used to launder funds for North Korea's Lazarus Group, OFAC

sanctioned the protocols see (US. Department of Treasury, 2022) and arrested two core developers—though parts of these sanctions were later overturned (Hussein, 2023). The episode underscores that hard privacy without an audit fallback invites severe legal push-back.

The strategic analysis also explains why the other combinations fail:

- **Hard privacy with soft auditability** (red) lets criminals open unlimited pseudonymous wallets and fragment flows; supervision is overwhelmed and welfare collapses.
- **Soft privacy with soft auditability** (brown) exposes amounts but not identities, so audits still fail to scale.
- **Soft privacy with hard auditability** (blue) is enforceable but offers honest users no meaningful privacy—everything is visible by default.

Figure 2 visualizes this strategic overlay, combining the Auer et al. (2025) technical canvas with the Capponi et al. (2025) welfare analysis. The red-hatched quadrant shows the welfare-dominated region where hard privacy meets soft auditability—a combination that enables criminal evasion while overwhelming enforcement capacity. The green quadrant represents the stable privacy window where both cryptographic protection and rule-based disclosure operate effectively.

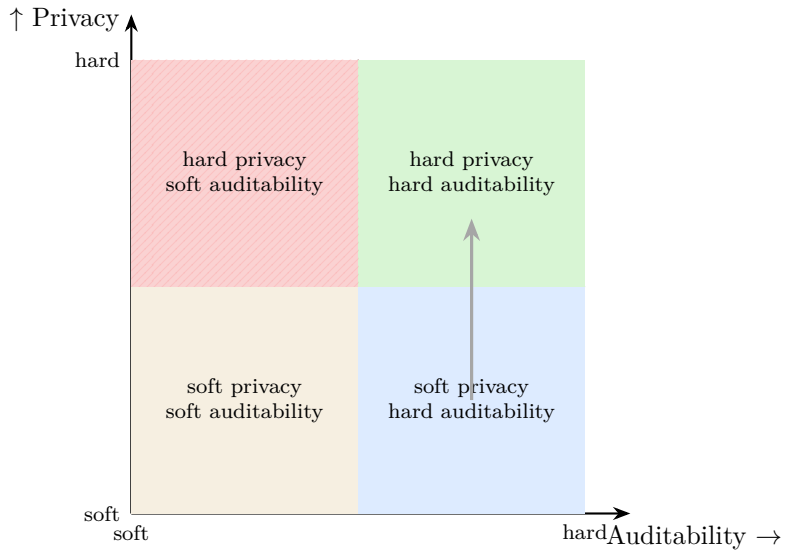

Figure 2: Auer et al. (2025) canvas with Capponi et al. (2025) stability overlay. Colours aligned with Table 1. The hatched quadrant (hard P + soft A) is dominated in the stylised model.

Under the model's assumptions, one design target is the strategically stable region in the upper-right quadrant. But this conclusion rests on specific modelling choices, and different assumptions—about enforcement technology, criminal adaptation, or governance structures—could yield different optimal configurations. The key implementation challenge lies in calibrating the disclosure threshold that separates cryptographically shielded transactions from those subject to automatic transparency. This threshold must balance competing demands: it should be high enough to protect routine consumer spending (coffee, groceries, local transport) while low enough to capture transactions that warrant regulatory attention. In practice, this means setting the boundary between current cash anonymity limits and typical retail spending ceilings, with flexibility to adjust as enforcement capabilities and threat patterns evolve.

Our mapping and welfare conclusions are model-dependent. They rely on (i) user-level consolidation (UPI), which blocks scalable multi-account evasion; (ii) an enforcement technology whose detection probability increases with per-user volume; and (iii) assumptions on the bad-user technology frontier and its capacity to innovate. Alternative hard/hard instantiations that relax these ingredients (e.g., pseudonymous credentials with weaker consolidation, different enforcement functions, or heterogeneous governance) could yield

different rankings.

6.1 Principles for assessing privacy-enhancing technologies

The strategic framework yields three concrete design principles for privacy-enhancing CBDC architecture:

Principle 1 — Co-design privacy and auditability as dual requirement dimensions: specify separate guarantees, implement jointly, and manage their interactions so gains in one don't erode the other.

The conventional "privacy versus compliance" framing assumes a single trade-off. The integrated framework shows this is false: privacy (cryptographic shielding) and auditability (rule-based disclosure) can be engineered separately. Modern privacy-enhancing technologies enable both simultaneously. The key insight: hard privacy protects honest users while hard auditability constrains criminal behaviour—but both elements appear important for favourable outcomes in the stylised setting.[25]

Principle 2 — Calibrate the privacy threshold to social value, not technical convenience

---

The disclosure threshold determines which transactions remain private and which trigger automatic visibility. The threshold can be calibrated to contexts where privacy has high social value (e.g., routine retail spending), with disclosure for larger transfers. The strategic model suggests setting the threshold between current cash anonymity

[25]Note: We treat privacy and auditability as distinct requirement dimensions with separately measurable guarantees. They are co-designed (not technically independent in implementation): proofs, logs, and key-escrow mechanisms couple the two. But the design goal is architectural: improving one dimension should not automatically degrade the other outside of explicitly calibrated policy thresholds.

limits (typically €1,000-3,000 in European jurisdictions) and median monthly consumer expenditure—with room to adjust as enforcement capacity improves.

An alternative version places strong privacy on the payer side while enforcing auditability primarily at receipt. In this model, merchants undergo know-your-business (KYB) checks, provide signed receipts, and deposit funds through accountable channels. Thresholds then apply to withdrawal amounts or aggregate merchant receipts rather than per-payment identity release. This occupies the same "hard privacy + hard auditability" region via a different access path.[26]

Principle 3—Governance-enforcing architecture: non-bypassable and verifiable overrides. The system should *constrain the operator by design*. Disclosure or privileged actions must be possible only via rule-based, *non-bypassable* paths and produce *verifiable evidence*. Concretely: use threshold release or decryption, separation of duties, append-only audit logs with public commitments, and automatic proofs that any override satisfied the stated policy. No single administrator should be able to read, alter, or disclose data unilaterally.[27]

An important qualification applies: threshold systems do not remove the possibility of exceptional disclosure; they make it conditional on multi-party cooperation under predefined rules. If a sufficient number of key holders collude, protected data can still be decrypted. The policy advantage lies not in making disclosure impossible but in requiring coordination across multiple institutions, enforcing separation of duties, and producing tamper-evident evidence of every access event. Central-bank independence may strengthen this arrangement, since it raises the institutional difficulty of compelling coordinated disclosure across several autonomous actors—though independence is one governance advantage among others, not an absolute guarantee.

These principles collectively point toward a specific technological approach: crypto-

[26]See, for example, the model of GNU Taler: https://www.taler.net/en/index.html.

[27]This principle is agnostic to trusted hardware. TEEs can help, but the core property is *operator-constraining, verifiable control*, which can also be realised with threshold cryptography, MPC, and transparency/audit commitments.

graphic systems that shield routine transactions while automatically triggering disclosure above policy-defined thresholds, anchored by real-name identity management with rule-based access controls. Implementing this vision requires navigating a complex landscape of privacy-enhancing technologies, which the next subsection maps systematically.

6.2 From principles to practice: a plain-English map of PETs

Our aim in this subsection is to explain, in policy terms, what privacy-enhancing technologies (PETs) can do for a retail CBDC today. This section is designed for decision-makers who need to understand the practical capabilities and limitations of current technology without becoming cryptography experts. We deliberately avoid mathematical formalism and focus instead on what each technology does, where it fits in a CBDC architecture, and what questions to ask before adoption.

A plain-English PET framing in four verbs Rather than organizing PETs by cryptographic primitive or mathematical property, we group them by what they do—their functional purpose in a payment system. Four capabilities matter most:[28]

1. **Hiding**: The first capability is to *hide* transaction data. This means keeping information confidential in transit, at rest, and—for limited tasks—even while being processed, so that intermediaries, network observers, and even system operators cannot read them without authorization.[29] The typical tools for hiding data include authenticated encryption, which uses public-key methods for key exchange and signatures to protect data in transit and at rest. Another important tool is the cryptographic commitment—think of it as a sealed digital envelope that hides a value now but lets you prove later that you did not change it. Zero-knowledge proofs take this further: they allow you to prove a fact, such as "this amount is under the threshold," without revealing the underlying data. For specialized tasks,

[28] More formal definitions and variants are in Auer et al. (2025); Arora et al. (2025); Vives et al. (2024).

[29] See surveys oriented to central-bank use cases: Auer et al., 2025; Arora et al., 2025; Vives et al., 2024.

fully homomorphic encryption can enable limited computations on encrypted data, though it remains too slow for full retail processing.[30] Some researchers also include trusted execution environments—secure hardware chips—in the hiding toolkit, though these come with important caveats around vendor trust, vulnerability patching, and governance of what code can run where.[31]

2. **Splitting**: The second capability is to *split* control across multiple parties so that no single institution can see or change sensitive data alone.[32] This enables what we might call "trust through division" rather than relying on one trusted operator. Secret sharing is a foundational technique here: it splits a secret—such as an encryption key—into pieces, where any threshold number of pieces (say, three out of five) can reconstruct it, but fewer pieces reveal nothing. Secure multiparty computation, often abbreviated as MPC, allows multiple parties to jointly compute a result without any party seeing others' private inputs. Distributed key generation creates cryptographic keys jointly, with no single party ever holding the complete key. These tools support threshold signatures and threshold decryption, where multiple independent parties—for example, three of five authorized officers—must jointly authorize an action. No single party can act alone. This capability is crucial for governance: it ensures that exceptional actions, such as disclosure or system recovery, require coordination across institutional boundaries, with each participant's role cryptographically verified. Splitting does not make disclosure impossible—a sufficient quorum can still decrypt—but it ensures that access is coordinated, rule-bound, and leaves an auditable trace.

[30]FHE lets a server compute directly on encrypted numbers without first decrypting them—useful for simple sums or rule checks on protected data. Today it is slower and more resource-hungry than alternatives, so it fits limited tasks rather than full, low-latency retail processing. Performance is improving, but widespread live use remains an engineering goal.

[31]TEEs can help, but come with caveats: (i) you must trust the chip vendor's design and updates; (ii) flaws can leak data, requiring rapid patching; (iii) remote attestation needs clear governance—who can run what code; (iv) treat TEEs as one defense layer, not the only one, and avoid single-vendor lock-in.

[32]See Auer et al., 2025; Arora et al., 2025; Vives et al., 2024.

3. **Proving**: The third capability is to *prove* compliance or policy adherence without revealing raw data.[33] Instead of revealing "Alice paid Bob €47.32," a privacy-preserving system can prove "this payment is under the threshold" or "the payer's balance remains non-negative" or "no double-spend occurred"—all without disclosing amounts, parties, or other details. The key tools here combine commitments and zero-knowledge proofs to create binding, verifiable statements about hidden data. For instance, a user might commit that a payment amount is some value x, and then prove that x is below the threshold without revealing x itself. Anonymous credentials extend this capability: they enable unlinkable proofs of attributes or limits. A user can prove "I am over 18" or "I have not exceeded my monthly privacy cap" multiple times without these proofs being linkable to each other or to the user's identity. Cryptographers call this selective disclosure: reveal exactly what policy requires, nothing more.

4. **Learning safely**: The fourth capability is to *publish or learn safely*—that is, to enable statistical reporting, aggregate analysis, or machine learning on payment data without exposing individual transactions or enabling re-identification.[34] Differential privacy adds carefully calibrated noise to aggregate statistics so that individual records cannot be reverse-engineered. Synthetic ledgers generate artificial datasets that preserve the statistical properties of real data without containing real transactions. Federated analytics computes aggregates across distributed data without centralizing raw records. These tools are valuable for reporting, research, and regulatory disclosure. But it is important to understand that they are not the primary shield for live payment privacy—that role belongs to the first three capabilities.

Where the four verbs fit in a retail-CBDC stack These four capabilities map to different parts of a CBDC architecture, and understanding this mapping helps clarify what each tool is meant to achieve. *Hide* and *Prove* work together in the transaction

[33]See Auer et al., 2025; Arora et al., 2025; Vives et al., 2024.
[34]See Auer et al., 2025; Arora et al., 2025; Vives et al., 2024.

layer. Payments are encrypted so that their content is hidden, and compliance rules are checked via cryptographic proofs without ever decrypting the underlying data. This pairing enables the core privacy-with-compliance objective: routine flows remain opaque to observers, while policy adherence is continuously and verifiably enforced.

Split underpins the governance-critical functions that sit above the transaction layer. Key custody is distributed so that no single party controls the master keys. Identity escrow—the mechanism that maps pseudonyms back to real names—requires coordination among multiple authorities. Exceptional disclosure paths, such as those triggered when a transaction breaches a policy threshold, follow multi-party protocols with tamper-evident logs. These mechanisms ensure that no individual administrator or institution can unilaterally access protected data.

Publish or learn safely operates in the reporting and analytics layer, separate from the live transaction layer. Central banks need to produce statistics, regulators need aggregate reports, and researchers need access to anonymized data. The tools in this family enable such outputs without compromising the privacy of individual users. Importantly, failures or compromises in this layer should not directly threaten the confidentiality of live payments, because the data flows are architecturally separated.

A crucial point bears emphasis: no single PET solves all problems by itself. Effective privacy requires composition—combining multiple techniques in ways that reinforce rather than undermine each other—and robust governance to manage the interfaces and ensure the whole system works as intended. The challenge is not merely selecting technologies from a menu, but integrating them into a coherent architecture with clear trust assumptions, well-defined failure modes, and accountable oversight.

What works now versus what needs engineering Decision-makers need to know not just what is possible in theory, but what is ready for deployment today. We can group PET capabilities into three readiness categories.

The first category comprises production-proven building blocks that are ready to use today with appropriate integration. Authenticated encryption and public-key infrastructure—

the foundations of secure communication—are mature, standardized, and widely deployed. Digital signatures and cryptographic commitments are well understood, efficient, and battle-tested. Zero-knowledge proof systems for specific statements, such as proving that a value lies within a range or that a balance is non-negative, already work at scale in production systems. Examples include certain cryptocurrency platforms and supply-chain tracking applications. Threshold cryptography for distributed key management and tamper-evident logs is used in cloud services and certificate authorities. Basic network mixing and batching techniques can hide transaction graphs from network observers. Trusted execution environments, when deployed in controlled settings with clear attestation policies and supply-chain vetting, provide additional defence layers. These building blocks are not experimental. They are used in production systems today, though not yet in the specific configuration required for a retail CBDC.

The second category includes capabilities that are technically feasible but require dedicated integration work and testing at scale. End-to-end system compositions that combine zero-knowledge proofs, threshold governance, and identity escrow to implement "privacy below threshold, automatic disclosure above" are within reach. The individual pieces exist and function, but they need careful integration, performance tuning, and stress testing under realistic retail conditions. Wallet-grade performance and user experience are also in this category. Cryptographic proofs must complete in milliseconds on consumer devices, with battery-efficient implementations and graceful error handling when networks are slow or devices are under load. Cross-vendor and cross-jurisdiction interoperability is another integration challenge: different implementations must verify each other's proofs, honor each other's policies, and operate under potentially different legal frameworks. These are solvable engineering challenges, not fundamental research problems. The main barriers are coordination, standardization, and sustained funding rather than scientific uncertainty. A common concern is whether encrypted transactions must be revealed in raw form to enable settlement across platforms. In principle, the answer is no: cryptographic proofs can attest that a payment satisfies the receiving system's policy requirements without disclosing amounts, identities, or other protected fields.

What interoperability does require is agreement on shared proof standards, settlement logic, threshold-governance rules, and incident-handling procedures across vendors and jurisdictions. The practical difficulty, in other words, lies in the governance and standardisation of proofs rather than in any inherent need to strip privacy at system boundaries.

The third category comprises capabilities that remain in active research and development. General compute-over-encrypted-data using fully homomorphic encryption works in principle, but performance remains orders of magnitude slower than needed for retail payments. Active research is narrowing this gap, but widespread deployment is not imminent. Robust network-graph unlinkability under adversarial traffic analysis—hiding not just transaction content but also the pattern of who pays whom—is harder than it sounds, especially against sophisticated observers who can analyse timing, volume, and routing metadata across multiple network vantage points. Post-quantum cryptography migration is another area of active work. Standards are emerging; the U.S. National Institute of Standards and Technology has selected candidate algorithms. But transitioning existing systems to quantum-resistant methods while maintaining acceptable performance and proof sizes will take years. Central banks should track progress in these areas and participate in standardization efforts, but they should not yet depend on these capabilities for production deployment.

A dashboard for decision-makers Rather than assigning simplistic red, yellow, or green maturity ratings, we provide a decision-support table that answers the questions decision-makers actually need to ask. For each PET family, the table explains what the technology does, where it fits in a CBDC architecture, and what questions to ask your technical team before adopting it. For deeper technical details, consult the comprehensive surveys by Auer et al. (2025), Arora et al. (2025), and Vives et al. (2024).

How this supports our policy requirements The framework we have developed rests on treating privacy and auditability as two different design dimensions, each enforceable through architecture rather than policy alone. The four PET capabilities map

Family	What it does	CBDC fit (examples)	What to verify before adoption
Hide	Keeps amounts/links confidential; protects data at rest/in transit/in use	Confidential amounts; shield routine payer identity; encrypted storage; TEEs for limited computations	Throughput/latency under peak retail load; failure modes; TEE attestation model and supply-chain risk; key lifecycle
Prove	Verifies rules without revealing raw data	Prove "under threshold", "no double-spend", "balance ≥ 0"; attribute proofs via anonymous credentials	Soundness and policy coverage of statements; proof size; audit verifiability by public authorities
Split	Distributes control and removes single points of failure	Threshold unmasking (governance committee); shared custody of master keys; MPC for compliance oracles	t-of-n governance defined in law; liveness under outages; incident playbooks; cross-jurisdiction operations
Publish and learn	Enables safe statistics and research	Releases of aggregates; synthetic datasets; federated models	Privacy budget management; re-identification risk; separation from the live transaction path

directly onto this framework. Routine payments satisfy strong confidentiality through *Hide*: transaction content is encrypted and remains opaque to all parties except those explicitly authorized under the system's rules. Compliance is continuously verified through *Prove*: cryptographic proofs establish that each payment satisfies policy requirements—such as falling below the disclosure threshold or maintaining a non-negative balance—without revealing the underlying amounts or identities. When exceptional disclosure is required, it follows a non-bypassable, logged path implemented through *Split*: multiple independent parties must jointly authorize access, and every such action produces tamper-evident evidence. Finally, *Publish and learn safely* addresses a separate concern—statistical reporting and research—without compromising the confidentiality of the live transaction rail.

This composition delivers the dual requirements articulated in Section 6.1: privacy for routine flows, enforced by cryptography; and auditability for exceptional cases, enforced by rule-based, multi-party protocols. Neither dimension is sacrificed for the other. Instead, they are implemented as complementary architectural features, each with its own technical realization and governance structure.

What remains hard—and tractable The open work is chiefly engineering rather than fundamental research. The challenge is not inventing new cryptographic primitives but making proven techniques work together at retail scale, under realistic operational constraints, and across institutional and jurisdictional boundaries. Each individual PET is usable today in controlled environments. The risk lies at the interfaces: how components interact under load, how keys are managed across their lifecycle, how governance rules are enforced during both routine operations and incidents, and how the system responds when something fails.

This situation has important implications for institutional capacity. Institutions need enough in-house PET competence to assess compositions, avoid lock-in, and verify claims; we return to the institutional implications in the conclusions. Cryptographic systems must be independently verifiable, not black boxes. This means demanding test vectors that allow independent verification of correctness, open verifiers that third parties can run, and reproducible audits that can be re-executed as systems evolve. Central banks do not need to become cryptography research labs, but they do need the capacity to evaluate technical claims, engage knowledgeably with vendors and the research community, and make informed trade-offs between different architectural options.

In short, institutions need in-house PET competence, independent verification, and to prioritise architectural soundness over convenience. The tools exist. The question is whether institutions are prepared to use them.

7 Conclusions

This paper has argued that privacy in retail central bank digital currencies (CBDCs) should not be treated as a marginal feature or a post hoc adjustment—nor as a perpetual trade-off against compliance. Privacy should be treated as a foundational design principle, essential for institutional legitimacy and public trust. Weak privacy protections risk undermining confidence in public digital money, with far-reaching consequences for adoption, competitive neutrality, and the credibility of central banks in their broader

mandates.

Recent advances in privacy-enhancing technologies (PETs), combined with strategic insights from game-theoretic modelling, now offer a viable path toward privacy-enhancing CBDC architectures. Our analysis—integrating technical capabilities with strategic stability—identifies a design space where strong cryptographic privacy and rule-based auditability can coexist. Under specific modelling assumptions, the welfare-optimal configuration lies in this region: encrypted flows below a calibrated threshold, automatic disclosure above it, with both dimensions enforced architecturally rather than institutionally. This technical insight transforms the policy debate: privacy and compliance are not opposing forces but complementary engineering challenges that modern cryptography can address simultaneously.

Section 6.2 established that the core building blocks are production-ready. Authenticated encryption, zero-knowledge proofs for specific statements, threshold cryptography, and network-layer protections are used in live systems today. Several research prototypes demonstrate retail-scale feasibility in controlled settings—proving that the technical ingredients exist and can work together. The remaining gap is not scientific uncertainty but engineering maturity, regulatory clarity, and above all, institutional commitment to making privacy a genuine priority rather than a rhetorical flourish.

CBDC architects face a critical choice. They can design payment infrastructure around current institutional practices and technological comfort zones, accepting their limitations as permanent constraints. Or they can establish architectural requirements according to democratic principles and long-term policy needs, then invest in maturing the necessary capabilities. This is not a choice between proven and experimental technology—the core tools are mature—but a choice about what to optimise for. Infrastructure designed around today's institutional constraints will entrench those constraints for decades. Infrastructure designed around policy objectives—privacy, auditability, competitive neutrality—creates the conditions for those objectives to be met.

The risk of technological conservatism is not merely that it delays progress. Infrastructure choices, once made, are extraordinarily difficult to reverse. Payment systems

are not software that can be patched iteratively; they are platforms on which entire economies depend. As argued in the introduction, public payment infrastructure sets de facto benchmarks for the broader digital economy. A CBDC designed with weak privacy will entrench that weakness as a precedent—making it far harder to demand stronger protections from private actors whose business models depend on data extraction.

Legislators have a crucial role in creating the conditions for privacy-by-design to succeed. Legal frameworks should support architectural flexibility rather than hard-coding premature technical assumptions into law. A principle-based legislative approach—one that defines guardrails, accountability mechanisms, and governance structures without prescribing implementation details—can create the regulatory space public digital infrastructure needs to evolve. Legislation should establish clear requirements: that privacy protections be verifiable rather than merely promised; that exceptional disclosure follow rule-based, logged procedures; that governance ensure separation of duties and prevent unilateral access. But it should avoid locking in specific technologies or threshold values that may need adjustment as enforcement capabilities and threat patterns evolve.

Central banks, for their part, must recognise that delivering meaningful privacy at retail scale is not simply a matter of vendor selection or system configuration. It requires sustained investment in multidisciplinary expertise—spanning cryptography, system design, network security, economics, and governance. It requires the institutional capacity to assess proposed compositions, avoid single-vendor lock-in, and validate technical claims through independent verification. And it requires a willingness to prioritise architectural soundness over short-term convenience, even when this means investing in capabilities that extend beyond traditional central banking domains.

The technical capabilities exist. Privacy-enhancing technologies have matured from academic curiosities to production-ready tools. The strategic framework developed in this paper shows how privacy and auditability can be treated as independent, mutually reinforcing design dimensions rather than opposing objectives. The question now is whether institutions are prepared to act on these insights—whether they will design for the payment systems democratic societies need, rather than accepting the limitations of systems

designed for a pre-digital era.

The decisions made today about CBDC architecture will shape payment infrastructure for decades. By treating privacy as a foundational design principle now, rather than a feature to be added later, central banks can build systems that strengthen rather than undermine democratic institutions. They can demonstrate that public infrastructure can deploy advanced technology while preserving citizen rights. And they can set standards that influence far beyond the payments domain—establishing that surveillance and data extraction are design choices, not technological necessities. With strategic commitment, international collaboration, and sustained investment in research and engineering, central banks have the opportunity to pioneer a new model of digital public infrastructure. Whether they seize this opportunity will determine not only the future of public money, but the broader relationship between citizens, technology, and the state in the digital age.

References

Abramova, S., Böhme, R., Elsinger, H., Stix, H., and Summer, M. (2023). What can central bank digital currency designers learn from asking potential users? In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 151–170, Anaheim CA. Available at: https://www.usenix.org/conference/soups2023/presentation/abramovaUSENIX Association.

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2022). Too Much Data: Prices and Inefficiencies in Data Markets. *American Economic Journal: Microeconomics*, Available at: https://pubs.aeaweb.org/doi/10.1257/mic.2020020014(4):218–256.

Acquisti, A., Bonatti, A., Goldfarb, A., Johnson, G. A., Miller, A. R., and Tucker, C., editors (2024). *The Economics of Privacy*. National Bureau of Economic Research Conference Report. London, Chicago.

Antequera, C. (2023). Available at: https://www.clarkemodet.com/en/articles/data-protection-in-latin-american-countries/Data protection in Latin American countries.

Armantier, O., Doerr, S., Frost, J., Fuster, A., and Shue, K. (2024). Nothing to hide? Gender and age differences in willingness to share data. BIS Working Paper, Available at: https://www.bis.org/publ/work1187.pdfBank for International Settlements.

Arora, R., Du, H., Kazmi, R. A., and Le, D.-P. (2025). Privacy-Enhancing Technologies for CBDC Solutions. Technical report, Available at: https://www.bankofcanada.ca/2025/01/staff-discussion-paper-2025-1/Bank of Canada.

Asrov, K. and Samonas, S. (2021). Privacy Enhancing Technologies: Categories, Use Cases and Considerations. Working Paper, Federal Reserve Bank of San Francisco, Available at: https://www.frbsf.org/wp-content/uploads/Privacy-Enhancing-Technologies_FINAL_V2_TOC-Update.pdfSan Francisco.

Auer, R., Böhme, R., Clark, J., and Demirag, D. (2025). Privacy-enhancing technologies for digital payments: Mapping the landscape. Technical report, Bank for International Settlements.

Auer, R., Frost, J., Gambacorta, L., Monnet, C., Rice, T., and Shin, H. S. (2022). Central Bank Digital Currencies: Motives, Economic Implications, and the Research Frontier. *Annual Reviews of Economics*, 14:697–721.

Bank for International Settlements (2022). CBDCs: An opportunity for the monetary system. Annual Economic Report, Available at: https://www.bis.org/publ/arpdf/ar2021e3.pdfBank for International Settlements.

Bank of England (2023). *The Digital Pound: A New Form of Money for Households and Businesses? Consultation Paper*. Dandy Booksellers Ltd, London.

Bergemann, D., Bonatti, A., and Gan, T. (2022). Available at: http://arxiv.org/abs/2004.03107The Economics of Social Data.

Big Brother Watch (2023). CBDC a privacy eroding pound? Technical report, Available at: https://bigbrotherwatch.org.uk/wp-content/uploads/2023/11/CBDC-a-privacy-eroding-pound-FINAL-1.pdfBig Brother Watch.

BIS (2019). Big tech in finance: Opportunities and risks. BIS Annual Report, Bank for International Settlements, Available at: https://www.bis.org/publ/arpdf/ar2019e3.pdfchapter 3.

Burman, A. (2023). Understanding India's New Data Protection Law. Technical report, Available at: https://carnegieendowment.org/research/2023/10/understanding-indias-new-data-protection-law?lang=enCarnegie Endowment for International Peace.

Buterin, V., Illum, J., Nadler, M., Schär, F., and Soleimani, A. (2024). Blockchain privacy and regulatory compliance: Towards a practical equilibrium. *Blockchain: Research and Applications*, Available at: https://linkinghub.elsevier.com/retrieve/pii/S20967209230005195(1):100176.

Capponi, A., Lee, M. J., and Zhu, B. (2025). Privacy-Enhanced Payment Systems. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5136803Technical report.

Chaum, D. (1985). Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, Available at: https://dl.acm.org/doi/10.1145/4372.437328(10):1030–1044.

Chaum, D., Grothoff, C., and Moser, T. (2021). How to issue a central bank digital currency. Technical report, Available at: https://www.ssrn.com/abstract=3965032Swiss National Bank.

Choi, J. P., Jeon, D.-S., and Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, Available at: https://linkinghub.elsevier.com/retrieve/pii/S0047272719300131173:113–124.

Cremeers, R. (2021). China's Emerging Data Protection Framework. Technical report, Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3964684Leiden University - Leiden Institute for Area Studies.

Croxson, K., Frost, J., Gambacorta, L., and Valletti, T. (2023). Platform based business models and financial inclusion: Policy trade offs and approaches. *Journal of Competition Law & Economics*, 19:75–102.

Data Protection Africa (2023). Available at: https://dataprotection.africa/data-protection-in-africa-progress/Mapping the progress (and delays) for data protection in Africa.

Elsinger, H., Stix, H., and Summer, M. (2025). Consumer Preferences for a Digital Euro: Insights from a Discrete Choice Experiment in Austria. BIS Working Papers 1302, Bank for International Settlements. Available at: https://www.bis.org/publ/work1302.htmAlso available as OeNB Working Paper, https://www.oenb.at/dam/jcr:f7d4024e-43e9-44b5-b208-c6adf3215e19/2025-07-10_Elsinger_neu.pdf.

European Central Bank (2022). The case for a digital euro: Key objectives and design considerations. Technical report, Available at: https://www.ecb.europa.eu/pub/pdf/other/key_objectives_digital_euro~f11592d6fb.en.pdfECB.

European Central Bank (2023). Progress on the investigation phase of a digital. Technical report, Available at: https://www.ecb.europa.eu/paym/digital_euro/investigation/governance/shared/files/ecb.degov230713-fourth-progress-report-digital-euro-investigation-phase.en.pdfECB.

European Data Protection Board (2020). Guidelines 06/2020 on the interplay of the Second Payment Services Directive (PSD2) and the GDPR. Technical report, Available at: https://edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202006_psd2_afterpublicconsultation_en.pdfEuropean Data Protection Board.

European Parliament and Council (2015). Directive (EU) 2015/2366: Revised Payment Services Directive (PSD2). Technical report, Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32015L2366European Parliament and Council.

European Parliament and Council (2016). Regulation (EU) 2016/679: General Data Protection Regulation (GDPR). Technical report, Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679European Parliament and Council.

Goodell, G. (2020). Available at: http://arxiv.org/abs/2006.05892Privacy by Design in Value-Exchange Systems.

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: https://www.ssrn.com/abstract=3891121*SSRN Electronic Journal*.

Grothoff, C. and Dold, F. (2021). Available at: https://www.taler.net/papers/euro-bearer-online-2021.pdfWhy a Digital Euro should be Online-firstt andBearer-based.

Hussein, F. (2023). Founders of crypto mixer arrested, sanctioned after US cracks down on Tornado Cash. Available at: https://apnews.com/article/cryptocurrency-treasury-crypto-sanctions-russia-north-korea-88115029d0a033b7b8b3e3a34dccf00c?utm*AP-news*.

Jones, N. (2024). *Understanding Payments: A Whistle-Stop Tour into What You Thought You Knew*. Routledge, Abingdon, Oxon New York, NY.

Kahn, C. M., McAndrews, J., and Roberds, W. (2005). Money is Privacy. *International Economic Review*, Available at: https://onlinelibrary.wiley.com/doi/10.1111/j.1468-2354.2005.00323.x46(2):377–399.

PCI Security Standards Council (2022). Payment Card Industry Data Security STandard (PCI DSS), Version 4.0. Available at: https://listings.pcisecuritystandards.org/documents/PCI-DSS-v4_0.pdfTechnical report.

People's Bank of China (2021). China_e_yuan_2021071614584691871. Technical report, Available at: http://www.pbc.gov.cn/en/3688110/3688172/4157443/4293696/2021071614584691871.pdfPeoples Bank of China.

PRS Legislative Research (2023). Available at: https://prsindia.org/billtrack/digital-personal-data-protection-bill-2023The Digital Personal Data Protection Bill., 2023.

Sarmiento, A. (2022). Available at: https://linkinghub.elsevier.com/retrieve/pii/S2666143822000175Seven lessons from the e-Peso pilot plan: The possibility of a Central Bank Digital Currency.

Schumacher, L. V. (2024a). *Decoding Central Bank Digital Currencies: Are CBDCs Friend or Foe?* DigitalEkho Edition, Zürich.

Schumacher, L. V. (2024b). *Decoding Digital Assets: Distinguishing the Dream from the Dystopia in Stablecoins, Tokenized Deposits, and Central Bank Digital Currencies*. Palgrave Macmillan, Cham.

Schumacher, L. V. (2024c). *Decoding the Digital Euro: Friend or Foe?: Make up Your Own Mind*. DigitalEkho Edition, Zürich.

Schumacher, L. V. (2024d). *Decoding the Digital Pound: Are CBDCs Friend or Foe?: Access for Yourself*. DigitalEkho Edition, Zürich.

Stuewe, S., Virza, M., Maurer, M., Lovejoy, J., Böhme, R., and Narula, N. (2024). Beware the Weak Sentinel: Making OpenCBDC Auditable without Compromising Privacy.

Sveriges Riksbank (2022). Sveriges_Riksbank_e-krona-pilot-phase-2. Technical report, Available at: https://www.riksbank.se/globalassets/media/rapporter/e-krona/2022/e-krona-pilot-phase-2.pdfSveriges Riksbank.

Tomescu, A., Bhat, A., Applebaum, B., Abraham, I., Gueta, G., Pinkas, B., and Yanai, A. (2022). Available at: https://eprint.iacr.org/2022/452.pdfUTT: Decentralized Ecash with Accountable Privacy.

Uhlig, H., Alonso, M., and Frost, J. (2023). Privacy in Digital Payments—Escaping the Panopticon. *Georgetown Journal of International Affairs*, Available at: https://muse.jhu.edu/article/91364324(2):174–180.

US. Department of Treasury (2022). U.S. Treasury Sanctions Notorious Virtual Currency Mixer Tornado Cash. Technical report, US. Department of Treasury.

Vives, G. T., Madars, V., Youngblom, R., Calabia, C., Vaughan, N., Said, Z., Mundakkal, C., and Mulings, S. (2024). Available at: https://static1.squarespace.com/static/59aae5e9a803bb10bedeb03e/t/675322efbb994b0f8891fe4c/1733501682224/BoE+MIT+DCI+Paper+Enhancing+the+Privacy+of+a+Digital+Pound.pdfEnhancing the Privacy of the Digital Pound.

White & Case LLP (2024). US Privacy Data Privacy Law Compliance Checklist. Technical report, Available at: https://www.whitecase.com/sites/default/files/2024-08/us-data-privacy-laws-compliance-checklist-2024.pdfWhite & Case.

Wikipedia Contributors (2024). Available at: https://en.wikipedia.org/wiki/Terrorist_Finance_Tracking_ProgramTerrorist Finance Tracking Program.

Zuboff, S. (2020). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, New York, NY, first trade paperback edition edition.

A Overview of typical data collected in digital payments

The following tables summarise the typical payment data collected in peer-to-peer (P2P) and retail transactions in digital payment ecosystems. While specific requirements may vary by jurisdiction, this overview captures the broadly applicable data elements typically involved. We build on Vives et al. (2024) and Jones (2024) for this overview.

Table A1: Typical data in peer-to-peer (P2P) transactions

Data Element	**Description**
Payer Identifier	Name, user ID, email, or mobile phone number
Payee Identifier	Name, user ID, email, or mobile phone number
Transaction Amount	Exact payment value
Payment Method	Bank account number, wallet address, payment app identifier
Transaction Timestamp	Precise time and date of transaction initiation
Transaction Reference	Unique transaction identifier, description provided by payer

Table A2: Typical data in retail transactions

Data Element	Description
Customer Identifier	Name, user ID, email, phone number, loyalty card or similar identifier
Merchant Identifier	Name, merchant ID, category, physical location
Transaction Amount	Exact payment value, including details of taxes or fees
Payment Method	Card number (partial or tokenised), bank account number, wallet address
Transaction Timestamp	Precise time and date of transaction initiation
Merchant-side purchase data	Category or itemised details of goods or services, where collected by the merchant and linked to the payment record
Transaction Reference	Unique identifier provided by the merchant or payment system

Note: Core payment data (identifiers, amounts, timestamps, references) are routinely collected and stored by payment intermediaries and banks. Merchant-side data—such as itemised purchase details—are collected separately by merchants and may be linked to payment records through loyalty programmes, point-of-sale systems, or data-sharing agreements. Both categories may be shared further with third-party processors, data brokers, analytics providers, and occasionally regulatory or governmental authorities, depending on jurisdiction-specific rules and compliance requirements.

B An overview of actual and discussed CBDC privacy solutions

Table B1: Overview of CBDC privacy features in selected projects. Evidence from public pilot descriptions suggests that most tested approaches prioritise institutional controls; only limited information is available on architectural (hard) privacy, and in several cases transaction graphs appear observable by default.

CBDC	Country	Stage	Privacy Model	Architecture	Data Access	Concerns
Digital Euro	Eurozone	Research	Pseudonymised online, limited offline anonymity; no full anonymity	Two-tier with intermediaries; offline component allows more privacy	ECB sees encrypted data; intermediaries have more access; AML/CTF obligations	Offline privacy unclear; privacy level subject to political discretion; no true anonymity
Digital Pound	United Kingdom	Consultation	Pseudonymised transactions; reversible with legal order; private use of data possible	Two-tier; intermediaries provide wallets and handle KYC	Private providers may use pseudonymised data for commercial purposes	Pseudonymity, not anonymity; surveillance and profiling risk
E-krona	Sweden	Pilot	Traceable transactions; banking privacy laws apply; no anonymity	Token-based on permissioned DLT; alias lookup system	Central alias mapping; ledger is traceable; AML regulations apply	"All electronic payments leave traces"; offline privacy limited
e-CNY	China	Pilot	"Managed anonymity": low-value pseudonymous, high-value traceable	Account-based; central ledger	Central bank and intermediaries have data access; biometric ID required	Extensive surveillance possibilities; privacy subordinated to state interests
Jam-Dex	Jamaica	Full rollout	Linked to photo ID and tax number; weak pseudonymity	Account-based; centralised architecture	Wallet providers can identify users; access granted to authorities on request	Plans to link to national ID system; lack of transparency
eNaira	Nigeria	Full rollout	Identity and biometrics required; no anonymous usage tiers	Account-based on private blockchain; CBN ledger control	CBN and intermediaries access identity and transaction data	Biometric surveillance; limited demonstrated enforcement
e-Peso	Uruguay	Pilot (ended)	Pseudonymised encrypted vaults; traceable token IDs	Token-based with centralised infrastructure	Traceable by central bank with court order; no unlinkable payments[35]	Weak privacy; low adoption; project discontinued
Digital Shekel	Israel	Proof of concept	Limited privacy; "conditional anonymity" for small payments only	Two-tier, potentially DLT-based; intermediaries run wallets	Designed to facilitate AML, tax collection; privacy viewed as secondary	Privacy conditional on policy; programmable privacy only

Table B1 draws on the overview by Big Brother Watch (2023) and on the following official sources: European Central Bank (2022), European Central Bank (2023), Bank

[35]In privacy-enhancing systems, a payment is said to be *unlinkable* if it cannot be traced back to the payer or correlated with other transactions by the same user, even by an adversary with access to the system's internal records. This contrasts with pseudonymous systems, where transactions can be linked to one another and potentially re-identified.

of England (2023), Sveriges Riksbank (2022), People's Bank of China (2021), https://boj.org.jm/a-primer-on-bojs-central-bank-digital-currency/, Sarmiento (2022), https://www.boi.org.il/en/economic-roles/payment-systems/future-payment-methods/digital-shekel-cbdc. Other comparisons between the privacy solutions between different CBDC models can be found in Schumacher (2024b) or Schumacher (2024c), Schumacher (2024a) and Schumacher (2024d).

C PETs used in research about PET and CBDC

Table C1: Overview of selected privacy-focused CBDC designs in the research literature.

Paper	PETs Used	Components Targeted	Key Features and Innovations
Chaum et al. (2021)	Blind signatures (anonymity-enhanced signatures)	Transaction processing, compliance, wallet design	Token-based CBDC using blind signatures for strong anonymity. Enables auditability at withdrawal and deposit; software-only and quantum-resistant.
Gross et al. (2021)	zk-SNARKs, commitments, nullifiers	Transaction processing, wallets, onboarding	Supports a privacy pool architecture with private, semi-private, and transparent transfers. Compliance via ZK turnover and balance caps.
Goodell (2020)	Blind signatures, zero-knowledge proofs, identity-based encryption	Wallets, onboarding, settlement infrastructure	Emphasises privacy by design and owner-custodianship; enables non-custodial wallets and separates issuance from infrastructure.
Tomescu et al. (2022)	Zero-knowledge proofs of knowledge, rerandomisable signatures, commitments, pseudorandom functions	Transaction processing, compliance, auditability	Scalable decentralised e-cash architecture (UTT) with anonymity budgets for privacy accounting; supports high-throughput, accountable transactions, threshold cryptography and sharding.
Stuewe et al. (2024)	Pedersen commitments, zero-knowledge range proofs	Transaction validation, auditability infrastructure	Extends OpenCBDC with auditability via confidential transactions; enables privacy-enhancing audits and scalable validation through commitments and ZK proofs.

D Strategic options for privacy-centred CBDC design

This appendix distils the paper's key insights into actionable principles and strategic options for central banks and public authorities engaged in the design and implementation of retail CBDCs. It is intended as a high-level guide to support institutional reflection and interdisciplinary collaboration.

1. Set Strategic Objectives Beyond Compliance

- Move from framing privacy as a trade-off to defining it as a design goal on par with compliance, resilience, and interoperability.
- Clearly articulate the desired privacy properties of a CBDC across transaction types, user tiers, and value thresholds.
- Anchor privacy requirements early in system architecture and design briefs, rather than addressing them post hoc through bolt-on features.

2. Adopt a Phased Architecture Strategy

- Design modular systems that allow for future integration of advanced privacy-enhancing technologies (PETs).
- Avoid irreversible architecture choices that lock in soft privacy constraints.
- Create pilot environments or experimental layers within CBDC projects to test hard privacy models in parallel with baseline implementations.

3. Engage in Collaborative Capacity-Building

- Recognise that delivering privacy at scale exceeds the current scope of most central bank IT teams.

- Commission academic and industry research through challenge programmes or structured collaboration.
- Support the development and evaluation of open-source PET implementations suited for high-volume public payment systems.

4. Strengthen Interdisciplinary Dialogue

- Treat CBDC privacy as a socio-technical challenge involving legal, technical, behavioural, and institutional dimensions.
- Engage stakeholders from privacy advocacy groups, civil society, computer science, economics, and law early in the process.
- Convene interdisciplinary working groups to align system goals with democratic principles of data protection and informational self-determination.

5. Invest in Strategic Communication and Trust

- Frame privacy not as an obstacle to enforcement but as a foundation for public trust and democratic legitimacy.
- Communicate privacy-by-design objectives transparently to the public and to oversight institutions.
- Position strong data protection as a European value and competitive advantage in the global digital economy.

This appendix is intended to inform strategic discussions within and across central banking institutions. It does not prescribe a singular path, but offers building blocks for a coordinated and forward-compatible approach to privacy-centred digital currency design.

E Legislative considerations for the digital euro

The legislative process for the digital euro offers a critical opportunity to align the institutional framework with the principles of privacy-centred design. At present, there is a risk that overly prescriptive legal language—particularly when based on early architectural assumptions—may inadvertently lock in technical choices that constrain future innovation and reduce institutional flexibility.

To ensure that public digital money evolves in line with democratic expectations and technological progress, legislation should avoid embedding detailed operational parameters into primary law. Instead, it should focus on establishing high-level principles, governance structures, and accountability mechanisms that leave room for architectural refinement as privacy-enhancing technologies mature.

Such an approach would empower central banks to explore and adopt stronger privacy guarantees without requiring legislative reversal or reinterpretation. It would also send a political signal that legislative support for the digital euro is conditional on a genuine commitment to informational self-determination as a design goal—not merely as a communications strategy.

In this way, the legislator can play a proactive and constructive role: enabling innovation, safeguarding democratic values, and helping ensure that Europe's public digital infrastructure remains adaptable, inclusive, and trusted.

The Working Paper series of the Oesterreichische Nationalbank is designed to disseminate and to provide a platform for discussion of either work of the staff of the OeNB economists or outside contributors on topics which are of special interest to the OeNB. To ensure the high quality of their content, the contributions are subjected to an international refereeing process. The opinions are strictly those of the authors and do in no way commit the OeNB.

The Working Papers are also available on our website (http://www.oenb.at) and they are indexed in RePEc (http://repec.org/).

Publisher and editor Oesterreichische Nationalbank
Otto-Wagner-Platz 3, 1090 Vienna, Austria
PO Box 61, 1011 Vienna, Austria
www.oenb.at
oenb.info@oenb.at
Phone (+43-1) 40420-6666

Editor Martin Summer

Cover Design Information Management and Services Division

DVR 0031577

ISSN 2310-533X (Online)