BANCA D'ITALIA
EUROSISTEMA

UIF
Unità di Informazione Finanziaria per l'Italia

Quaderni dell'antiriciclaggio

Analisi e studi

Engaging with Privacy Stablecoins: A Framework for Scalable and KYC/AML-Compliant Adoption

Michele Manna

luglio 2026

numero 35

BANCA D'ITALIA
EUROSISTEMA

UIF
Unità di Informazione Finanziaria per l'Italia

Quaderni dell'antiriciclaggio

Analisi e studi

Engaging with Privacy Stablecoins: A Framework for Scalable and KYC/AML-Compliant Adoption

Michele Manna

n. 35 – luglio 2026

La collana Quaderni dell'antiriciclaggio *ha la finalità di presentare statistiche, studi e documentazione su aspetti rilevanti per i compiti istituzionali dell'Unità d'Informazione Finanziaria per l'Italia.*

La collana si articola in diversi filoni: il filone Statistiche *presenta, con periodicità semestrale, statistiche sulle segnalazioni ricevute e dati sulle attività dell'Unità; il filone* Rassegna normativa *illustra i principali aggiornamenti della normativa e della giurisprudenza in materia AML/CFT; il filone* Analisi e studi *comprende contributi sulle tematiche e sui metodi in materia di contrasto al riciclaggio e al finanziamento del terrorismo. I lavori pubblicati riflettono esclusivamente le opinioni degli autori, senza impegnare la responsabilità delle Istituzioni di appartenenza.*

Unità di Informazione Finanziaria per l'Italia
Per la pubblicazione cartacea: autorizzazione del Tribunale di Roma n. 1942013 del 30 luglio 2013
Per la pubblicazione telematica: autorizzazione del Tribunale di Roma n. 1932013 del 30 luglio 2013

Direttore responsabile
Enzo Serata

Indirizzo
Largo Bastia, 35 – 00181 Roma – Italia

Telefono
+39 0647921

Sito internet
https://uif.bancaditalia.it/

ISSN 2283-3498 (stampa)
ISSN 2283-6977 (online)

Stampato nel mese di luglio 2026

Grafica e stampa a cura della Divisione Editoria e stampa della Banca d'Italia

Engaging with Privacy Stablecoins:
A Framework for Scalable and KYC/AML-Compliant Adoption *

Michele Manna

Abstract

This paper examines whether privacy-preserving stablecoins ("privacy stablecoins") can achieve mainstream status as retail payment instruments. We first argue that stablecoins operating on public Layer-1 blockchains remain limited in everyday retail payments due to structural constraints in scalability, compliance, and user-level privacy. We then show that privacy stablecoins built on Layer-2 architectures can address these constraints by combining high-throughput with constrained privacy—user-level confidentiality supported by selective disclosure and due-process-based regulatory access. These features align with both the value of privacy in payments and the preconditions for widespread use. Although adoption remains nascent, advances in proof systems, data-availability infrastructure, and auditable privacy technologies may support broader acceptance. Ultimately, the regulatory stance will play a decisive role in shaping whether the industry pursues the innovation needed for privacy stablecoins to potentially mature into scalable, trusted, and compliant digital payments.

Sommario

Il presente lavoro esamina se le *stablecoin* che preservano la privacy ("*privacy stablecoins*") possono raggiungere un elevato livello di diffusione come strumenti di pagamento al dettaglio. In primo luogo, si sostiene che le *stablecoin* operate su blockchain pubbliche di Layer-1 rimangono limitate nell'uso per tali finalità a causa di vincoli strutturali in termini di scalabilità, conformità normativa e privacy a livello utente. Si mostra quindi che le *privacy stablecoins* basate su architetture di Layer-2 possono superare tali vincoli, combinando elevata capacità di elaborazione con una privacy delimitata ("*constrained privacy*"), ossia riservatezza a livello dell'utente sostenuta da meccanismi di disclosure selettiva e da forme di accesso regolatorio basate sul rispetto del giusto processo. Queste caratteristiche risultano coerenti sia con il valore della privacy nei pagamenti, sia con le condizioni necessarie per un utilizzo diffuso. Sebbene l'adozione sia ancora nelle fasi iniziali, i progressi nei sistemi di prova, nelle infrastrutture di disponibilità dei dati e nelle tecnologie di privacy verificabile potrebbero favorirne una più ampia accettazione. L'orientamento regolamentare giocherà un ruolo decisivo nel determinare se il settore intraprenderà il percorso di innovazione necessario affinché le *privacy stablecoins* possano evolvere in strumenti di pagamento digitali scalabili, affidabili e conformi.

JEL Classification: E42, G28, O33

Keywords: no-questions-asked principle, Layer-2, constrained privacy, regulatory compliance

---

* Financial Intelligence Unit for Italy (UIF), Bank of Italy. The views and opinions expressed herein are those of the author and do not necessarily reflect the views of the institution. The author wishes to thank all participants in the internal seminar held on 23 April 2026. Special thanks are due to Marco Militello.

Contents

1. Introduction and motivation

Stablecoins have grown rapidly yet remain far from functioning as mainstream monetary instruments. The combined market capitalization of Tether and USD Coin—the two dominant U.S. dollar-pegged stablecoins, together accounting for roughly 90 percent of the market—stands at about USD 260 billion. In contrast, U.S. M1 and M2 at end-2025 stood at USD 19.3 trillion and USD 22.6 trillion, respectively.[1] Size alone, however, does not confer monetary relevance. For an instrument to achieve mainstream status, it must be widely accepted for everyday payments, used as a unit of account in pricing and contracting, and integrated into household and corporate balance sheets. By these criteria, stablecoins have not yet attained meaningful monetary relevance.

This paper's first contribution is to show that the gap between current usage and broad monetary relevance in stablecoins reflects largely structural—rather than incremental—limitations.

On the supply side, payment providers require technological throughput sufficient not only to avoid operational failures but also to ensure that system efficiency does not deteriorate as user bases and transaction volumes grow. Legal scalability is equally critical: compliance mechanisms—most prominently KYC identification—must remain operationally viable as transaction volumes grow, without relying on case-by-case forensic interventions.

On the demand side, widespread adoption requires user protections analogous to those found in established payment systems: privacy safeguards, error-correction mechanisms, dispute resolution channels, and cognitive simplicity. Users should be able to interpret value effortlessly in the relevant unit of account, without incurring mental conversion costs. Public authorities, in turn, require effective Anti-Money Laundering and Countering the Financing of Terrorism (AML/CFT) compliance, tax observability, and—in some jurisdictions—support for capital-flow management.

Standard stablecoins operated on base Layer-1 (L1) blockchains do not meet these requirements. Under decentralized assumptions, public L1 blockchains remain far from matching established payment systems in throughput, typically measured in transactions per second. Equally important, regulatory compliance is generally implemented through off-chain overlays: when courts, supervisors, or auditors must rely on ad hoc forensic techniques to link pseudonymous identifiers to individuals, the system cannot achieve legal scalability.

As argued throughout the paper, ex-post compliance differs fundamentally from compliance embedded directly into protocol design (compliance-by-design).

Purely privacy-oriented cryptocurrencies do not remedy these frictions. Systems such as Monero offer strong obfuscation but lack fiat pegs, auditability, and transparent reserves. Zcash introduced zero-knowledge (ZK) proofs with selective disclosure, but it is not a stablecoin and does not address reserve management. For users, privacy, liquidity, and value stability must be jointly satisfied; privacy coins satisfy only the first of these requirements.

Privacy-preserving stablecoins ("privacy stablecoins") offer a different path. These instruments pair a stable-value asset with a privacy-preserving transaction layer that enables constrained privacy, i.e., default confidentiality combined with selective, rule-based disclosure. These features are typically realized in high-throughput Layer-2 or off-chain computation environments, anchored to a public

[1] Data sources: CoinMarketCap (stablecoin market capitalization); Board of Governors of the Federal Reserve System, *H.6 Money Stock Measures* (U.S. monetary aggregates, M1 and M2).

Layer-1 for settlement finality and data availability. This architecture allows privacy-by-design to coexist with auditability, regulatory assurance, and both technical and legal scalability.

The second aim of this paper is to offer a structured and accessible account of how privacy stablecoins operate and to assess their technological and commercial maturity. We do not claim that privacy stablecoins resolve all challenges: peg stability still depends on reserve quality and governance;[2] supervisory models must adapt to a more diversified payment ecosystem, in which technology firms and fintech providers increasingly operate alongside traditional banks; and the dominance of USD-denominated stablecoins may reinforce dollarization pressures in some jurisdictions. Importantly, privacy-preserving stablecoins do not eliminate AML risks; rather, they reconfigure them by enabling targeted supervisory access without undermining baseline confidentiality.

Our analysis draws on five strands of literature:

(i) payment economics, institutional scalability, and privately issued money (Adrian and Mancini Griffoli, 2021; Gorton, 2020; Gorton and Pennacchi, 1990; Gorton and Zhang, 2023; Kahn and Roberds, 2009; Mitchell *et al*., 2024; Rochet and Tirole, 2006);

(ii) privacy in money and its interaction with public policy objectives (Auer *et al*., 2025; Borgonovo *et al*., 2021; Giammusso and Nardelli, 2025; Kahn, 2018; Kahn, McAndrews and Roberds, 2005);

(iii) cryptographic foundations of confidential transactions and ZK proofs (Ben-Sasson *et al*., 2014; Ben-Sasson *et al*., 2018; Groth, 2016);

(iv) blockchain architectures and scalability (Mssassi and El Kalam, 2025; Nadler and Schär, 2023; Nardelli, De Sclavis and Iezzi, 2025; Reno and Roy, 2025);

(v) blockchain systems integrating privacy and compliance-by-design (Arner, Auer and Frost, 2020; Chaliasos, Firsov and Livshits, 2025; Chaudhary, 2023; Croman *et al*., 2016; Duffie, Olowookere and Veneris, 2025; Gross, Sedlmeir and Seiter, 2022; Habib, 2024; Mahrous, Caprolu and Di Pietro, 2025; Sarencheh, Kiayias and Kohlweiss, 2023; Sun, Zhang and Liu, 2022).[3]

The combination of these strands reflects our effort to translate an economic inquiry—why privacy in money is valuable and under what conditions a payment instrument achieves mass adoption—into a technologically grounded analysis that incorporates recent advances in blockchain systems.

Two foundational insights guide our contribution. First, following Kahn and co-authors, the demand for privacy in payments arises for legitimate purposes, such as commercial confidentiality, personal safety, and protection from profiling. Second, as Gorton and Zhang (2023) emphasize, while technologies and legal forms evolve, the underlying economic issues surrounding privately issued money persist.[4] We also draw on recent work—including Auer *et al*. (2025) and Giammusso and Nardelli (2025)—highlighting the inherent tension among privacy, auditability, and technical performance.

Relative to this literature, our contribution is to develop a unified framework for evaluating the criteria for mainstream monetary acceptance; diagnose why standard stablecoins fall short; and demonstrate

---

[2] Across all types of privately issued money, an alternative to maintaining adequate reserves is the creation of a public backstop, such as government-provided insurance.

[3] This taxonomy necessarily simplifies a literature that is both rich and overlapping, and some of the cited contributions span multiple strands. With respect to category (iv), we cite works most directly relevant to the questions addressed in this paper, while acknowledging the breadth of the broader blockchain literature.

[4] Given the general nature of this observation, similar conclusions have likely been reached by other authors writing on privately issued money and beyond.

how—and under what conditions—privacy stablecoins can satisfy these economic, legal, and technical requirements.

A challenge in structuring this paper was determining how much space to devote to technical elements. While brevity and accessibility are valuable, an overly minimal treatment would obscure how privacy-preserving stablecoins fulfill their multiple functions. We aim to strike a balanced compromise that offers a shared conceptual basis for policymakers, market participants, and innovators.[5]

The remainder of the paper is organized as follows. Section 2 presents key data on stablecoin usage. Section 3 examines the conditions under which a payment instrument becomes mainstream. Section 4 provides the necessary technical background, describing how transactions are executed on a blockchain's Layer-2 and subsequently settled on the base Layer 1. Building on the conceptual framework of Section 3 and the technological foundations of Section 4, Section 5 assesses whether privacy-preserving stablecoins could achieve mass adoption. Section 6 surveys existing commercial solutions. Section 7 concludes.

2. Data on stablecoin usage

As of end-2025, the aggregate market capitalization of stablecoins stood at approximately USD 300 billion. Around 94 percent corresponds to tokens pegged one-for-one to the U.S. dollar and backed by dollar-denominated deposits and U.S. Treasury securities.[6] A second segment, representing roughly 5 percent of the market, consists of crypto-collateralized stablecoins. The remaining share comprises algorithmic designs[7] and tokens pegged to non-USD fiat currencies such as the euro or the Singapore dollar.

The market is highly concentrated around Tether (USDT) and USD Coin (USDC)[8], as shown by the data reported in footnote 6; unless otherwise indicated, we treat these two instruments as representative of the sector. They are not, however, perfect substitutes. Tether exhibits deeper liquidity and broader global availability, representing roughly two-thirds of total stablecoin capitalization and serving as the primary instrument across trading platforms and crypto-financial markets. USD Coin, while somewhat less liquid, emphasizes regulatory compliance, transparency of reserves, and institutional oversight. It is compliant with the EU MiCA regime, whereas Tether is not (as of the time of writing).[9]

---

5 In the technical sections of the paper, the use of specialized terminology is at times unavoidable. To assist the reader, we provide explanations of key terms throughout.

6 Author's calculations based on data from CoinMarketCap. Applying a 94 percent share to an aggregate market size of USD 300 billion yields approximately USD 282 billion. Relative to the USD 260 billion combined capitalization of Tether and USD Coin cited in the Introduction, this implies that other U.S. dollar-pegged stablecoins account for roughly USD 20 billion.

7 Algorithmic stablecoins seek to maintain a target peg without explicit reserves, relying instead on supply-adjustment rules triggered by deviations of the market price from the reference value. Their classification as "stablecoins" remains debated, with some authors treating them as a subset of fiat-reference crypto-assets—alongside assets such as Bitcoin and Ethereum—rather than as reserve-backed instruments.

8 Crypto-assets are commonly referred to both by their full name (e.g., Bitcoin, Tether) and by the tickers used by data providers (e.g., BTC, USDT). Tickers are useful for avoiding ambiguity, as a given crypto-asset may circulate on multiple blockchains, whereas the ticker uniquely denotes the asset irrespective of the underlying network.

9 As clarified by ESMA (2025), CASPs operating platforms are expected to cease making all non-MiCA-compliant Asset-Referenced Tokens (ARTs) and Electronic Money Tokens (EMTs) available for trading.

These differences shape user preferences: USDT dominates trading, market-making, and cross-exchange liquidity provision, whereas USDC is more commonly adopted in contexts where regulatory alignment, reserve transparency, and institutional oversight are central considerations.

However, despite these differences, both coins have grown rapidly.[10] Their combined capitalization increased from USD 116 billion in 2023 to USD 263 billion by end-2025, following more than a fourfold expansion over the preceding three-year period (2020–23). Yet even this rapid growth leaves the sector small relative to conventional monetary aggregates: the total capitalization of all stablecoins amounts to only about 1 percent of the U.S. M2 money stock (Fig. 1).[11]

Fig. 1

Capitalization of Tether and USD Coin, combined

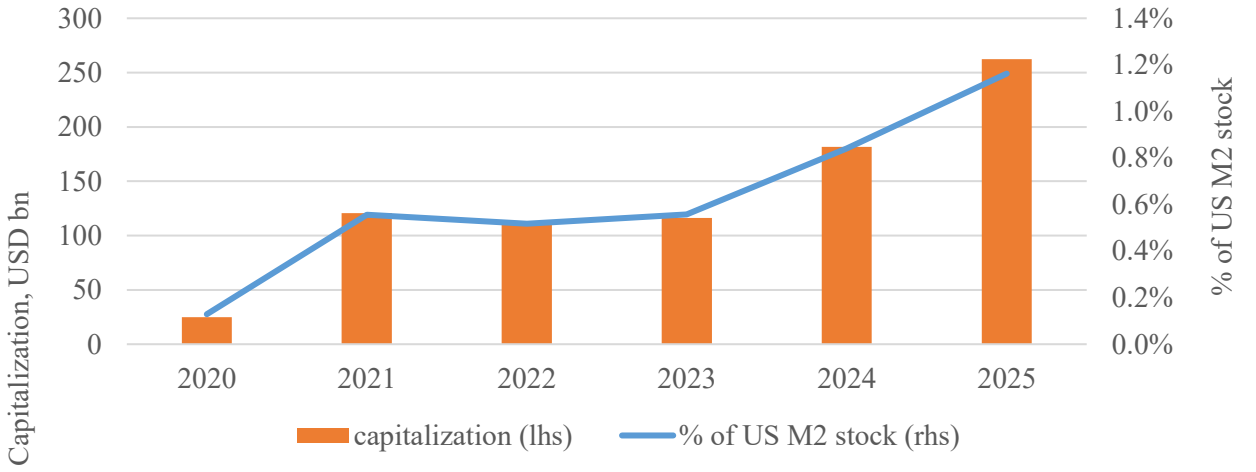

Source: Author's elaboration on data from CoinLore and the U.S. Federal Reserve.

Stock comparisons, however, do not capture the full picture. An alternative perspective is that the economic significance of a payment instrument may be better assessed through transaction flows rather than outstanding balances. In 2025, the estimated on-chain transaction volume of USDT and USDC reached USD 32 trillion. Benchmarking this figure against the ACH Network—the leading U.S. electronic payment network, which processed USD 93 trillion over the same period (Sundararajan, 2026; Nacha, 2026)[12]—yields a ratio of roughly one-third.

Superficially, such a comparison might suggest that stablecoins have already attained substantial economic relevance. The interpretation is, however, misleading.

Stablecoin transaction flows are overwhelmingly driven by trading activity, market-making, arbitrage, and exchange-related settlement, rather than by retail or commercial payments (Bolt, Lubbersen and Wierts, 2022; Lyons and Viswanath-Natraj, 2020). A simple calculation illustrates the point. Based on the figures above, a transactions-to-stock ratio (conceptually analogous to velocity in circulation) for stablecoins would exceed that of U.S. M2 by a factor of nearly 30—consistent with the literature showing that velocity becomes uninformative when an asset is used primarily for

[10] It is possible that the simultaneous growth of the two instruments reflects, in part, the complementary nature of their design features, which offer users scope for diversification across liquidity, transparency, and regulatory compliance profiles.

[11] We take a U.S. monetary aggregate as the reference point because both Tether and USD Coin are pegged to the U.S. dollar and function, in effect, as private substitutes for the dollar.

[12] The figures reported in Sundararajan (2026) draw on underlying data from Artemis Analytics.

portfolio rebalancing or short-term financial intermediation rather than for transactional purposes (Anderson and Rasche, 2001).[13] Even with conservative adjustments, the implied stablecoin flow/stock ratio remains implausibly high compared to M2, if interpreted as reflecting genuine retail payment activity (see Annex 1 for details).

It is therefore more plausible that only a modest fraction of the reported USD 32 trillion reflects genuine payment activity (IMF, 2025).[14] Since trading-related flows cannot currently be distinguished from non-trading transactions in public-blockchain data, stock measures provide a more reliable indicator of the economic scale of stablecoins and of their degree of penetration in agents' balance sheets.[15]

Despite their limited retail adoption to date, stablecoins offer a clear advantage: low transaction costs. Transferring USD 5,000 in USDT or USDC on low-fee networks typically costs between USD 0.01 and USD 0.05, with higher fees arising only on congested chains. These costs are far below the 1.5–3 percent merchant fees associated with card payments and the 6–7 percent average charges for cross-border retail payments (McKinsey, 2025b). The contrast is striking, particularly given the current G20 target of 3 percent for cross-border payments. Even domestic account-to-account transfers generally exceed a few basis points, depending on the jurisdiction. Cost competitiveness therefore does not appear to be the binding constraint on stablecoins' mass adoption.

Instead, as argued in the Introduction and developed in Section 3, stablecoins remain constrained by their current architectural and institutional design. When operated primarily on base-layer blockchains, they face a number of structural barriers: cognitive frictions for users, fragmented regulatory treatment, limited scope for compliance-by-design, and technical constraints related to throughput and settlement finality. These features—rather than any lack of user demand—explain why stablecoins have yet to achieve mainstream monetary relevance.

3. What makes a payment instrument mainstream

3.1 The features of a mainstream payment instrument

Payments arise to resolve two fundamental frictions: the mismatched timing of trading needs and the limited enforceability of intertemporal commitments (Kahn and Roberds, 2009). Although both money and credit mitigate timing mismatches, money has a comparative advantage when the enforcement of promises is weak or costly; it is accepted not for its consumption value but because it reliably circulates.

While many objects could in principle serve as money, only a limited set ultimately attains widespread acceptance. A useful benchmark in this respect is the no-questions-asked (NQA) principle (Gorton and Pennacchi, 1990; Gorton, 2020; Gorton and Zhang, 2023): an instrument becomes mainstream if it is accepted in exchange without material due diligence. Put differently, users are effectively "information-insensitive" with respect to instruments that satisfy NQA.

---

[13] When used in financial transactions, a single unit of money can circulate repeatedly as collateral or as source of liquidity across market venues.

[14] Stablecoins also circulate extensively off-chain, and such movements are not captured in public-blockchain data, further limiting the informational content of on-chain flow statistics.

[15] An attempt to filter out what Coinbase refers to as "inorganic activity" in stablecoin flows is presented in a report published on August 5, 2024 by Coinbase Institutional. Their analysis reaches conclusions consistent with ours regarding the limited role of stablecoins in retail payments.

For token-based instruments (e.g., banknotes), NQA concerns authenticity and stable purchasing power; for account-based instruments (e.g., deposits), it relates to issuer credibility and the certainty of settlement. By this criterion, instruments whose value fluctuates materially—such as unbacked crypto-assets like Bitcoin—require users to perform due diligence before accepting them in settlement, including checking exchange rates and assessing short-term volatility. Such requirements violate NQA and render these instruments unlikely candidates for mainstream payments.

We posit three additional properties that are jointly necessary for the mass-adoption test: dimensional scalability, legal (regulatory) scalability, and constrained privacy.[16] These requirements operationalize NQA at scale: they ensure that users, intermediaries, and authorities can accept the instrument without continual due diligence on performance, legality, or privacy boundaries.

Dimensional scalability denotes the capacity of an instrument and its supporting infrastructure to accommodate large increases in users and transaction volume, while preserving stable and acceptable throughput, latency[17], and fee levels. It also encompasses system state[18] growth—the on-chain storage requirements relevant for rollups and full nodes. Because payment platforms exhibit economies of scale and strong network externalities (Kahn and Roberds, 2009; Rochet and Tirole, 2006), performance must not deteriorate appreciably as adoption expands.[19] Otherwise NQA fails in practice when congestion forces users to conduct due diligence regarding fees, delays, or potential reversals.

Legal (regulatory) scalability refers to the ability of a payment instrument and its infrastructure to uphold compliance obligations—including KYC identification, AML/CFT requirements, Travel Rule data exchange, record-keeping, supervisory access, sanctions screening, consumer redress, and the recognition of settlement finality—in a manner that remains tractable as volumes grow and across heterogeneous technologies, jurisdictions, and regulatory frameworks. As argued in the Introduction, ex-post overlays, such as exchange-level KYC, are insufficient in a peer-to-peer environment: if courts or auditors must repeatedly engage in ad hoc forensic analyses to link pseudonymous identifiers to individuals or reconstruct transaction semantics, the system cannot achieve legal scalability.[20]

To further ground these requirements, we draw on Sarencheh, Kiayias and Kohlweiss (2023), who identify two core conditions: (i) the backing of a redeemable instrument must be auditable without compromising user privacy; and (ii) KYC and AML-CFT checks must be enforceable without disabling privacy-preserving operations. Together, these requirements define what the authors term an "auditability and regulation-friendliness" environment.

Constrained privacy is related to, but distinct from, legal scalability. Although anonymity is often associated with illicit use—hence the ongoing debate over cash (Camera, 2021)—privacy in payments has legitimate economic value when transacting parties cannot fully trust one another

---

[16] As a cautionary note, these terms are not standardized in the literature, even though the underlying concepts are well established. We introduce this terminology to clarify how technical constraints intersect with economic and legal requirements.

[17] In ICT systems, *latency* refers to the time elapsed between an input and the corresponding output. We return to this notion later when discussing its specific meaning in the context of blockchain architectures.

[18] In blockchain systems, the *system state* refers to the complete snapshot of all data on the network at a given point in time, including account balances and, where applicable, smart-contract code and associated variables. See, for example, the website *HashtagWeb3*; in the Reference section, we list websites separately, specifying their URLs and, where available, the date of publication of the cited page.

[19] Rochet and Tirole develop their analysis in the context of card payment systems. However, the underlying concepts they advance naturally extend to payment instruments more broadly.

[20] A further concern associated with ad hoc analyses is whether their effectiveness applies symmetrically across users, intermediaries, and jurisdictions.

(Kahn, McAndrews and Roberds, 2005). Empirical work further indicates that the attractiveness of a payment instrument increases with the privacy it offers (Borgonovo *et al*., 2021). However, privacy must be intentionally bounded ex ante so that lawful compliance and targeted supervision remain feasible; without such limits, privacy risks collapsing into opacity. Constrained privacy therefore goes beyond mere "compatibility" with regulation: users do not obtain absolute anonymity; instead, what remains private, for whom and under what conditions, is specified within the protocol itself (Duffie, Olowookere and Veneris, 2025).

To clarify the distinction: legal scalability—notably condition (ii) of Sarencheh, Kiayias and Kohlweiss (2023)—requires that privacy not impede the enforcement of KYC/AML rules. Constrained privacy, by contrast, involves embedding ex ante limits on privacy within the protocol so that it remains within acceptable regulatory boundaries.[21] In practice, this implies that certain attributes become visible only through lawful due-process procedures. Authorities must be able to obtain targeted access when legally justified, while the system's architecture must simultaneously preclude bulk visibility, dragnet searches[22], and passive monitoring of all transactions.

3.2 Stablecoins as mainstream payment instruments

Evaluated against the criteria set out above, stablecoins deployed primarily on public Layer-1 (L1) blockchains face several binding constraints.

First, public L1s must contend with the well-known blockchain trilemma, that is the scalability-decentralization-security trade-off (Buterin, 2023; Croman *et al*., 2016; Mssassi and El Kalam, 2025; Reno and Roy, 2025). Leading L1s process orders of magnitude fewer transactions per second (TPS) than global card networks and exhibit fees that rise substantially under periods of congestion. Bitcoin processes roughly 7 TPS under standard parameters (Li *et al*., 2018), while other L1s achieve higher—but still modest—rates, typically in the range of 20–60 TPS for systems such as Ethereum (pre-rollup) and Bitcoin Cash.

A substantial literature on consensus optimizations, parallel execution, signature aggregation, and networking improvements reports markedly higher experimental performance (surveys include Dabbagh *et al*., 2021; Esmaili and Christensen, 2025; Nasir *et al*. 2026). Two caveats apply. First, the proposed techniques are heterogeneous: some apply only to specific chains or depend on architectural assumptions incompatible with decentralization objectives. Divergent methodologies, hardware assumptions, and adversarial models make it difficult to establish a robust consensus on sustainable peak performance. Second, reported throughput figures frequently reflect laboratory conditions—such as clusters of homogeneous high-performance nodes, unconstrained bandwidth, or relaxed validation rules—that do not reliably translate to open, permissionless environments.

A balanced reading of the literature suggests that a sustainable upper bound for decentralized L1 throughput—under optimized consensus variants, parallel execution, and advanced networking—

[21] This also sheds light on why constrained privacy is a dimension of the NQA principle: users should not need to perform due diligence on privacy boundaries or compliance guarantees when deciding whether to accept the instrument.

[22] A *dragnet search* is a law-enforcement operation that, instead of targeting specific individuals, investigates a broad population indiscriminately. Physical dragnets traditionally involve cordoning off an area (e.g., a city block) and questioning everyone within it, regardless of suspicion. A digital dragnet extends this logic to data collection, such as accessing information from every mobile device in a given location or analyzing complete cell-tower logs. These practices contrast sharply with targeted, due-process-based investigative methods.

likely falls within 1,000–10,000 TPS.[23] Yet even this optimistic envelope remains well below the expanding capacity of established global payment networks, which operate at orders-of-magnitude higher throughput. Visa reportedly handles more than 20,000 TPS during peak periods and has stress-tested capacity above 65,000 TPS, with Mastercard at comparable levels. Large digital wallet systems substantially exceed these figures: Alipay reached 544,000 TPS during the 2019 Singles' Day event, while WeChat Pay recorded 409,000 TPS over the same period.[24]

A similar picture emerges for latency. The two "classical" blockchains exhibit multi-minute settlement finality, on the order of 60 minutes for Bitcoin and 12–15 minutes for Ethereum.[25] By contrast, in major credit card networks and large digital wallets, the time to transaction approval—the metric most relevant for user experience—is typically well under one second.[26] Some newer blockchains (such as Aptos and Solana) achieve comparably fast finality, but only by relying on significantly higher hardware requirements, which are not widely accessible to average users, and by operating with reduced decentralization. This again reflects the scalability-security-decentralization trade-off.

The evidence suggests that public blockchains remain far from matching established payment systems in throughput and latency at global scale. While continuous improvements to base-layer performance have expanded the technological frontier, they have not eliminated the underlying trade-off, since the trilemma reflects hard constraints, not transitory engineering limitations. High-throughput L1s typically achieve their performance by sacrificing robustness—e.g., by increasing susceptibility to liveness failures[27] or network outages—or by concentrating validation among a smaller set of participants. In short, public L1s are not optimized for mass-market payment scale, whereas permissioned systems can attain substantially higher throughput, albeit at the cost of openness and global reach.[28]

Second, pseudonymous addresses, low-cost address rotation, and the ease of cross-chain transfers (chain-hopping) render ex-post forensics probabilistic and resource-intensive. Regulatory frameworks such as MiCA in the European Union govern issuers and intermediaries but generally do not embed identity within on-chain transfer primitives, leaving peer-to-peer flows largely opaque from a supervisory standpoint. Furthermore, many public blockchains provide only probabilistic finality. By contrast, deterministic finality—under which a transaction becomes irrevocable once included—requires governance structures and coordination guarantees comparable to those found in

---

[23] In this context, *sustainable* refers to throughput achievable under real-world traffic conditions using contemporary hardware and proof systems, rather than peak capabilities observed only under idealized or laboratory conditions. *Parallel execution* is a scaling technique where multiple, unrelated transactions are processed simultaneously, to boost overall throughput (see the website *Forklog*).

[24] See the website *SDK Finance*.

[25] See the websites *Coincodex* and *Spark*.

[26] In contrast to *user-facing latency*, the *clearing and settlement latency* of credit and debit card payments is generally not real-time; in retail systems, these processes typically occur in end-of-day—if not, in some cases, monthly—batches.

[27] In the context of blockchains, *liveness failures* refer to situations where a decentralized network stalls, halts, or fails to make progress, resulting in valid transactions not being included in new blocks within a reasonable timeframe.

[28] Hyperledger Fabric is reported to achieve approximately 1,000 to 3,000 TPS; however, it is a permissioned blockchain network, in which participants are identified and mutually authenticated, and therefore does not exhibit the degree of openness of public blockchains.

recognized payment and clearing systems.[29] Most public L1s and associated stablecoin arrangements do not meet these conditions and are therefore not incorporated within such regimes.[30]

Third, public L1s often deliver misaligned privacy outcomes: they provide too much privacy at the edges for illicit flows—enabled by pseudonymous addressing and on-chain obfuscation tools—and too little for ordinary users, whose transfers and associated metadata remain publicly observable and can often be deanonymized (albeit typically requiring non-trivial analytical effort). Empirical work documents extensive address clustering and transaction-graph deanonymization (Heimbach *et al*., 2025). Even privacy-oriented systems such as Zcash, Dash, and Monero exhibit vulnerabilities, including significant information leakage under realistic assumptions (Biryukov and Tikhomirov, 2019).[31] By contrast, conventional account-based payment systems require identity verification at onboarding while protecting routine transaction confidentiality through established contractual and institutional safeguards.

These limitations are not intrinsic to stablecoins as monetary instruments. Rather, they stem from deploying them directly on public L1 blockchains without protocol-embedded mechanisms for scalability, compliance, and privacy (BIS, 2025). As we argue in the next two sections, Layer-2 (L2) architectures that separate high-throughput private execution from L1 settlement and data availability—such as ZK proof rollups with encrypted state and compliance-by-design features, including selective disclosure, authorization proofs, and audit-friendly reserve attestations—can, in principle, satisfy the NQA criterion while also meeting the institutional requirements associated with money-like instruments.

4. Execution of transactions on Layer-2 and interaction with Layer-1

Having concluded that current stablecoins do not structurally meet the conditions for mainstream status, in this section we discuss how these limitations may evolve when stablecoins are deployed on Layer-2 systems.

4.1 The respective roles of Layer-2 and Layer-1

At the foundational Layer-1 (L1), blockchains perform two core functions: (i) validating transactions and (ii) replicating computation across many independent nodes. Validation at L1 relies on Byzantine-fault-tolerant (BFT)-style consensus mechanisms—with Bitcoin as a notable exception— that allow the network to reach agreement on the order and correctness of transactions despite potentially malicious participants.[32] Because agreement must be achieved among numerous independent

[29] *Probabilistic finality* means that there remains a non-zero chance of a chain reorganization, although this probability decreases rapidly as additional blocks are added. *Deterministic finality*, by contrast, guarantees that once a block is finalized it cannot be reverted. Bitcoin and base-layer Ethereum are standard examples of systems with probabilistic finality, whereas modern BFT-style protocols (see below Section 4.1) implement deterministic finality. A downside of the latter approach is that such systems may be less robust under adversarial network conditions, more vulnerable to correlated validator outages, and prone to abrupt failure if quorum assumptions are violated.

[30] Instruments lacking robust legal finality cannot reliably serve as settlement media in institutional contexts, limiting their pathway to mainstream adoption.

[31] The work by Biryukov and Tikhomirov is widely cited and regarded as a standard source on blockchain deanonymization. Despite the time elapsed since its publication—during which underlying blockchain technologies have evolved—the core findings remain relevant and broadly accepted in the academic and technical literature.

[32] A *Byzantine fault* is one that presents different symptoms to different observers. The term takes its name from an allegory, the *Byzantine Generals Problem*, developed to describe a situation in which the system's actors must agree on a strategy, but some of these actors are unreliable in a way which causes other actors (who behave correctly) to

validators, the process is inherently time- and resource-intensive. In most prevailing L1 designs, full nodes re-execute all transactions, ensuring universal verifiability but limiting throughput and increasing latency.[33] This is sometimes described as the "cost of decentralization".

Layer-2 (L2) systems adopt a fundamentally different architecture. Rather than relying on fully decentralized consensus to order and execute transactions, an L2 typically delegates ordering to a *sequencer*—often a single operator or a small committee—whose role resembles that of a traffic controller. The sequencer rapidly collects and orders transactions into batches, which are then submitted to Layer-1 as rollup blocks. Execution is then verified by a separate component, the *prover*, which produces succinct validity proofs—often, though not exclusively, ZK proofs—certifying that the entire batch of transactions was processed correctly.[34] These two roles, sequencing and proving, are logically distinct, even if they may be implemented by the same entity in practice (see Annex 2).

At first glance, the sequencer-prover architecture may appear analogous to the functions performed by conventional financial intermediaries. This analogy is, however, only partial. L2 systems are logically distinct from—but not independent of—the underlying L1.[35] Crucially, an L2 must publish sufficient information to L1 to enable any party to verify or challenge invalid behavior. L1, in turn, does not re-execute the individual transactions contained in the rollup batch; it verifies only the accompanying succinct validity proof. If the proof verifies, the L2 state root is accepted; otherwise, the batch is rejected.

This architecture ensures that a sequencer cannot finalize an incorrect state. If it attempts to include invalid operations, the prover cannot generate a valid proof, and the L1 rollup contract will reject the batch (Chaliasos *et al*., 2024; Nainwal, Kamble and Awathar, 2025). Thus, the security of an L2 depends not on trusting the sequencer but on the cryptographic infeasibility of producing a validity proof for an incorrect state transition (Section 4.2 introduces the leading types of proof systems).

While we refer generically to "blockchains", the interaction described here is most concretely instantiated on Ethereum, which provides the smart-contract infrastructure needed to verify validity proofs and manage rollup state commitments. Other platforms—such as Solana and Avalanche—offer high-performance execution environments and support related scaling designs, but the rollup-centric architecture discussed in this section is most fully developed in the Ethereum ecosystem.[36]

In summary, L1 and L2 are not substitutes but complementary layers that perform specialized roles. L2 provides fast, low-cost, high-throughput execution—akin to a payment processor—while L1 functions as the globally shared, tamper-resistant ledger with final authority to prevent double

---

disagree on the strategy, and they may be unaware of the disagreement. In computing systems, *Byzantine fault tolerance* is the resilience of a fault-tolerant computer system to such conditions.

[33] Strictly speaking, participation in consensus is not open to all network users but only to those who satisfy eligibility requirements, which in most public blockchains remain relatively light. Under Proof-of-Stake (PoS) systems—such as the model now adopted on Ethereum—participants must post a stake to become validators, and block proposers and attestors are then selected through stake-weighted, pseudorandom mechanisms. This contrasts with the Proof-of-Work (PoW) approach originally used in Bitcoin, where any participant with sufficient computational resources may contribute to consensus.

[34] Validity proofs guarantee that a batch of transactions has been executed correctly, but they do not by themselves ensure that users reconstruct the underlying state unless all relevant transaction data is made available.

[35] We use the plural form ("systems") to reflect the fact that Ethereum currently hosts approximately 150 L2 networks (at the time of writing), each of which settles its state on Ethereum. For this reason, Ethereum is frequently described as a network of networks.

[36] Conceived in 2013 and launched in 2015, Ethereum is widely regarded as the first programmable blockchain to achieve large-scale adoption. Its smart-contract architecture enables the deployment of decentralized applications across a broad range of use cases, including games, data-management tools, and decentralized finance (DeFi) protocols. Ethereum completed its transition from a Proof-of-work to a Proof-of-Stake consensus mechanism in 2022.

spending. This division of labor preserves the security guarantees of public blockchains while enabling performance characteristics compatible with mass-market payment systems.

4.2 Rollup types and an introduction to zero-knowledge proofs

Two principal rollup families are widely deployed in practice: optimistic rollups and zero-knowledge (ZK) rollups (Song, Qu and Wei, 2024).[37]

In optimistic rollups, the term "optimistic" reflects the assumption that each batch posted to L1 is valid unless challenged within a pre-specified dispute window. The sequencer orders transactions and submits the resulting state root, together with the necessary data, to the L1 contract. In principle, any participant may file a fraud proof upon detecting an invalid state transition. If a challenge succeeds, the batch is rolled back, and the dishonest party is penalized. While this design simplifies the sequencer's function and avoids the computational burden of generating proofs for every batch, it introduces latency and uncertainty: finality—particularly for withdrawals—is deferred until the challenge period has elapsed.

ZK rollups operate differently. Each batch is accompanied by a validity proof demonstrating that all transactions were executed correctly according to the L2's rules. The sequencer rapidly orders transactions, while the prover constructs a succinct cryptographic proof for the entire batch. The L1 verification contract checks only this proof rather than re-executing the underlying transactions. This design provides deterministic correctness once the proof has been verified: if any transaction is invalid—such as an attempted double-spend—it becomes computationally infeasible to produce a valid proof for the resulting state.[38] The principal trade-off lies in the computational intensity of generating proofs, although these costs are amortized across the batch and continue to fall due to engineering improvements.

At the core of ZK rollups is the zero-knowledge proof mechanism, which enables a prover to convince a verifier that a given statement is true without revealing any information beyond the truth of the statement itself. Formally, ZK proofs satisfy three core security properties (Villar, 2025): completeness (or correctness), soundness, and zero-knowledge. These ensure, respectively, that true statements are accepted, false statements cannot be proven except with negligible probability, and no information other than the statement's truth is revealed.[39] Intuitively, ZK proofs enable public verification of private computation, precisely the mechanism that allows rollups to inherit L1 security without requiring L1 nodes to re-execute full computations (see Annex 3).

Two major proof systems implement these ideas in rollup settings: SNARKs and STARKs (Chainlink, 2024).[40]

---

[37] We refer readers to Song, Qu and Wei (2024) as an accessible introduction to rollup architectures and the transfer of information from L2 to L1. Although many valuable references exist on this subject, these authors present the key technical concepts with notable clarity and concision, offering a comprehensive overview within a compact six-page exposition.

[38] *L2 deterministic correctness* is distinct from *L1 deterministic finality*, which highlights why both layers are required.

[39] In completeness, both parties are assumed to be *honest*, meaning only that they follow the rules of the protocol (with no moral connotation). In soundness, the prover may be dishonest. In zero-knowledge, the verifier may be dishonest; the property guarantees that such a verifier still learns nothing beyond the validity of the statement.

[40] Other widely used proof systems include PLONK, a prominent SNARK variant with a universal and updatable structure (the initial generation of secret parameters needs to be done only once for all applications), and Bulletproofs, which are frequently employed for confidential-transaction schemes.

SNARKs—*Succinct Non-Interactive Arguments of Knowledge*—produce extremely small proofs, often only a few hundred bytes (succinctness), enabling fast on-chain verification and low L1 gas costs. They are non-interactive (requiring only a single prover-to-verifier message) and satisfy knowledge-soundness: any prover generating an accepting proof must possess a valid witness (typically an execution trace consistent with the batch).[41] Many SNARK systems rely on strong algebraic assumptions, typically implemented using elliptic-curve constructions (see Annex 4). While verification is highly efficient, proof generation can be computationally intensive, and the long-term post-quantum security of elliptic-curve-based systems remain uncertain (Chaliasos *et al*., 2024).

STARKs—*Scalable Transparent Arguments of Knowledge*—emphasize scalability and transparency. They rely solely on collision-resistant hash functions[42], avoiding elliptic-curve assumptions, and are widely believed to offer post-quantum security, based on current knowledge. STARKs exhibit particularly good performance on large computations, benefiting from substantial parallelisation during proof generation.[43]

In practice, many modern architectures combine the two proof systems. SNARKs are favored where minimizing on-chain verification costs is critical, owing to their succinct proofs, while STARKs are preferred in settings where transparency, the absence of trusted setup, and long-run scalability are prioritized.[44]

Two additional considerations apply across all rollups. First, to inherit L1 security, an L2 must publish sufficient state-transition data to Layer-1 (or to a data-availability, DA, layer offering equivalent guarantees[45]). Without robust data availability, users cannot reconstruct the L2 state or safely exit the system if the sequencer censors transactions or becomes unavailable. Second, although many L2s currently rely on centralized sequencers to achieve high performance, their fundamental security derives from validity proofs and data availability, not from the trustworthiness of the sequencer.

4.3 An illustrative example: Bob sends stablecoins to Alice on Layer-2

At the outset, a transaction on a Layer-2 system resembles any other blockchain transaction. Bob controls an account (an "address") on an Ethereum-based L2 and holds a balance of a given stablecoin—say, more than 100 USDC. Likewise, Alice has an address on the same L2. To transfer 100 USDC to Alice, Bob constructs a transaction on his device and signs it with his private key. The signed transaction is then submitted to the L2 sequencer, which receives the essential elements of the message: sender, receiver, amount, and signature.

---

[41] In the context of SNARKs, a *witness* is the private information (or "input") known only to the prover, which demonstrates that a given statement or computation is true (see the website Medium).

[42] A *hash function* is a mathematical algorithm that maps data of arbitrary size (input) to a fixed-length bit string (output or digest). It is designed to be deterministic, fast, and one-way (given the output, it is computationally very difficult to recover the input). A hash function is *collision-resistant* if it is computationally infeasible to find two distinct inputs that produce the same output.

[43] STARKs' trade-offs include significantly larger proof sizes, typically tens to hundreds of kilobytes, which increase on-chain data costs. In some configurations, verification can also be more resource-intensive, resulting in higher gas consumption relative to the highly succinct proofs produced by SNARKs.

[44] In the context of SNARKs, a *trusted setup* refers to the initial generation of secret parameters required by the proof system.

[45] DA is central to rollup security: it ensures that any user can reconstruct the L2 state independently, and failures in DA can render funds effectively frozen. Different rollup teams adopt different solutions—such as Ethereum calldata and Celestia—reflecting varying assumptions about cost, performance, and trust models. Importantly, data availability is conceptually distinct from proof validity: validity proofs guarantee correct computation, whereas DA guarantees that the underlying state-transition data is accessible.

A clarifying point is necessary. Bob can transfer funds to Alice only if her address is activated on the same L2. This is not guaranteed: multiple L2s operate in parallel (see fn. 35) so Alice's address may reside on a different L2, or she may hold an account only at the L1 level. Each L2 maintains its own isolated state, and Ethereum L1 does not track unified balances across different L2s. In this sense, L2 states exist "above" L1 and synchronize with it only when publishing state roots and accompanying validity proofs (in ZK rollups) or when batches are finalized after the challenge period (in optimistic rollups). The practical implications of cross-L2 transfers and bridging are discussed in the relevant footnote.[46]

Returning to the example: the sequencer processes Bob's transaction by deducting 100 USDC from Bob's balance and crediting the same amount to Alice. This update is not applied in isolation but as part of thousands of transactions aggregated into a single batch. The L2 prover then constructs a zero-knowledge proof certifying that all transactions in the batch complied with L2's rules: Bob's signature was valid, his balance was sufficient, and all associated state transitions followed the prescribed protocol.

A further clarification is required. Standard ZK rollups—such as zkSync or Scroll—provide cryptographic correctness guarantees but do not encrypt transaction data; the sequencer observes sender and receiver addresses, transferred amounts, and other metadata in plain text. Privacy-preserving rollups—for example, Aztec or Aleo—combine correctness with encryption, concealing transaction details while still enabling the prover to generate a valid ZK proof. Thus, privacy does not arise automatically from the use of ZK proofs. It requires an additional cryptographic layer integrated into the rollup design.

This distinction hopefully clarifies the term *privacy stablecoin* used in this paper. A privacy stablecoin is not a new instrument issued with monetary backing or contractual terms. Rather, it typically consists of an existing stablecoin (such as USDC or USDT) transacted with an L2 environment that employs privacy-preserving rollups, including variants outlined in Section 6. In this setting, privacy derives not from the stablecoin itself but from the infrastructure through which it is transferred.

5. Discussion: the potential of L2-based stablecoins to achieve mainstream status

5.1 Dimensional scalability

Section 3 showed that dimensional scalability remains a binding limitation for transactions executed directly on L1, especially when compared with the evolving capacity of established global payment networks. Recalling the key figures, even under optimistic and highly specific assumptions, the literature places sustainable throughput for decentralized L1s in the range of 1,000–10,000 TPS, with typical day-to-day performance well below 1,000 TPS. By contrast, the two leading credit-card networks reliably operate at more than 20,000 TPS, while major digital wallet systems (e.g., Alipay, WeChat) have recorded peak volumes on the order of half a million TPS.

Against this benchmark, L2 constructions routinely achieve execution rates in the range of 3,000–10,000 TPS. Several analyses anticipate order-of-magnitude gains as ZK rollups become more

[46] The canonical route is for Bob to withdraw his funds to L1 and then transfer them to Alice. An alternative is to bridge them from his L2 to the L2 on which Alice holds an account and execute the transfer there. This trust-minimized path is, however, slow and costly. As a result, participants often rely on cross-liquidity networks (such as Across, Hop, or Connext), which enable fast cross-L2 stablecoin transfers by using liquidity providers that advance the funds immediately and recover them once the underlying settlement completes.

efficient and as calldata costs decline.[47] Some projections estimate 50- to 100-fold improvements relative to Ethereum L1, potentially approaching 100,000 TPS for execution.[48] These figures, however, should be interpreted cautiously, as they are scenario-based estimates: achievable throughput on L2 ultimately depends on the data-availability bandwidth of the underlying L1—i.e., the rate at which the network can publish and verify transaction data—and on the cadence at which validity proofs are posted.

Latency exhibits a similar pattern. Ethereum L2 rollups offer near-instant confirmation, typically around 2–3 seconds.[49] This is slightly longer than established payment instruments, but the difference is negligible from a user perspective.

The conclusion is clear: while dimensional scalability remains a binding limitation for transactions executed directly on L1, L2 rollups substantially narrow the performance gap with mainstream payment infrastructures. They offer a credible path toward mass-market throughput and latency, contingent on continued advances in proof-generation speed and, for throughput especially, data-availability bandwidth.

5.2 Legal scalability

The first condition identified by Sarencheh, Kiayias and Kohlweiss (2023)—auditability of the reserves backing a redeemable instrument without compromising user privacy—is relatively straightforward. In this respect, privacy-preserving stablecoins do not differ in any fundamental way from ordinary L1-based stablecoins. Reserve transparency is determined by issuer-level reporting, custodial attestations, and applicable supervisory frameworks, not by user-level visibility into transaction data. As a result, privacy at the transaction layer does not impede reserve audits.

Condition (ii)—that KYC and AML-CFT checks be enforceable without disabling privacy-preserving operations—requires more nuance. As noted earlier, most public blockchains do not embed native identity primitives; compliance is typically imposed ex post at the level of intermediaries such as exchanges or custodians. Executing transactions on an L2 does not, by itself, resolve this limitation. However, privacy-preserving stablecoins can incorporate cryptographic mechanisms enabling protocol-native compliance. Achieving such outcomes depends on two important qualifications: first, they require deliberate architectural and cryptographic design choices and are not achieved by default; second, even with these mechanisms in place, certain limitations remain.

The relevant mechanisms include: (a) ZK proofs; (b) shielded pools or note-based models; (c) commitment schemes combined with nullifiers; (d) selective-disclosure keys; and (e) identity-bound proofs.[50] Below, we outline the non-technical intuition behind each component.[51]

---

[47] *Calldata costs* in blockchain technology, particularly within the Ethereum Virtual Machine (EVM) ecosystem, refer to the fees paid to include input data in a transaction (see the website Cyfrin).

[48] See the website Medium.

[49] See the website Cryptos.

[50] In this context, an *identity-bound proof* refers to a cryptographic mechanism—typically implemented using ZK techniques—that allows a user to demonstrate membership in an authorized or KYC-verified identity set without revealing which specific identity they hold. *Identity-bound nullifiers* extend this concept by enabling one-time or rate-limited actions (such as preventing double-voting or repeated use of a free service) while preserving user anonymity. Together, these tools support privacy-preserving compliance by ensuring that actions can be linked to a valid identity group without disclosing the underlying identity.

[51] For more technical background, we refer readers to the literature listed under item (v) in Section 1.

(a) A regulated intermediary—such as a bank, VASP, or licensed exchange—issues a credential to a user following successful KYC/AML checks. The user can subsequently provide a ZK proof demonstrating possession of this credential without revealing any underlying information. This enables regulators or authorized verifiers to confirm compliance attributes while preserving user anonymity, a paradigm sometimes described as anonymous yet permissioned access.[52]

(b) A shielded pool can be engineered as a permissioned domain accessible only to users who have passed KYC verification, with balances represented internally as encrypted notes. Privacy applies to transfers within the pool, while access is controlled through cryptographic admission checks.[53] Crucially, such behavior requires explicit protocol-level gating and does not arise automatically; it is achieved only when the appropriate technical components are deliberately integrated.

(c) Commitment schemes conceal transaction details such as amounts and recipients, while nullifiers prevent double-spending by ensuring that each committed note can be spent only once.[54] Together, these mechanisms uphold transaction integrity—including ensuring sufficient balance and preventing illicit value creation—without revealing user identities. From a regulatory perspective, they ensure that privacy does not compromise core soundness requirements.

(d) Selective-disclosure mechanisms allow users to reveal specific information—such as proofs of clean funds or transaction histories—to regulators or auditors upon request, while maintaining confidentiality by default. This supports compliance with statutory information-request processes without enabling indiscriminate surveillance.

(e) A trusted entity issues a digital KYC credential to a user, binding it cryptographically to the user's verified identity. The user can then generate a ZK proof demonstrating that a transaction originates from a KYC-verified individual without revealing which individual. This enables private yet permissioned transactions: only verified actors may interact with the system, while their identities remain concealed by default.

Taken together, these mechanisms enable a form of protocol-native, KYC-compatible privacy, consistent with condition (ii). Access to the shielded pool is restricted to verified users; transactions within the pools remain private while their integrity is enforced cryptographically; and regulators can obtain targeted information through selective-disclosure pathways. To some extent, this structure mirrors conventional payment systems, in which users access services only after meeting onboarding requirements, while thereafter benefiting from routine transaction confidentiality subject to lawful information requests.

Nevertheless, several grey areas remain. First, these mechanisms often rely on user-assisted disclosure: users must cooperate by providing selective-disclosure proofs when legally required, whereas some jurisdictions mandate investigative tools that do not rely on voluntary cooperation. Second, the resulting design creates a domain in which transactions are private by default—albeit

---

[52] Given the borderless nature of blockchains, KYC-compatible privacy ultimately depends on the legal interoperability of credential issuers across jurisdictions.

[53] A shielded pool can be visualized as a private room to which only authorized users have access. Inside the room, individuals can deposit funds, transfer them between sealed lockers, and withdraw them later; however, observers outside the room cannot determine which locker belongs to which user and how value moves between lockers. The system enforces the integrity of these movements while preserving confidentiality.

[54] Consider the following analogy. A user writes a note and places it inside an envelope, which is then sealed—this corresponds to the commitment, as the contents are hidden but the envelope is publicly observable. The user then tears a unique corner off the envelope to mark it as "used"; this is the nullifier, ensuring that the envelope cannot be reused. So long as the envelope is placed on a desk in an open corridor, anyone walking through can see that the envelope exists and has already been used, but no one can read the note inside.

within defined boundaries—which may still prompt concerns among authorities worried about excessive concealment. Third, users may transact within the gated environment in full compliance with regulatory requirements while simultaneously engaging in unrelated activity through fully anonymous or pseudonymous channels, complicating supervisory visibility over aggregate behavior.

5.3 Constrained privacy

In Section 3, we noted that legal scalability and constrained privacy are overlapping yet distinct concepts. The overlap is most evident in mechanisms such as ZK-KYC credentials, selective-disclosure keys, and shielded pools, all of which simultaneously support privacy and regulatory compliance.[55] In other respects, however, the two dimensions diverge. Many privacy-preserving systems rely on users voluntarily disclosing information, whereas regulators often seek independent, non-cooperative access pathways. Likewise, privacy-enhancing designs aim to minimize metadata leakage, while modern AML frameworks typically depend on risk scoring, transaction-pattern analysis, and sanctions screening—all of which presuppose some degree of metadata visibility.

This motivates a separate discussion of whether privacy-preserving stablecoins can satisfy the requirements of constrained privacy. The short answer is that they can, but doing so is not automatic. A privacy stablecoin achieves this standard only when the system guarantees both (a) privacy against the public and other unauthorized observers, and (b) bounded transparency for authorized actors under due-process procedures. These properties hold only if several conditions are jointly met: identity-gated access, ensuring that only KYC/AML-verified users can enter the private transaction environment; protocol-level privacy boundaries, so that confidentiality guarantees are enforced deterministically by code; and selective-disclosure mechanisms that allow regulators to obtain targeted information in legitimate cases. In brief, constrained privacy should be regarded as a design choice, not an emergency property, ensuring that privacy must remain bounded rather than absolute.

The table overleaf presents a simple two-dimensional classification of the concepts under discussion. We include it because "auditability and regulatory-friendliness" and "constrained privacy" are typically treated separately in the literature. For example, Sarencheh, Kiayias and Kohlweiss (2023) focus on auditability, whereas Kahn, McAndrews and Roberds (2005) analyze privacy in monetary systems.

A privacy stablecoin can occupy quadrant A, the ideal configuration, but reaching this position requires several mechanisms to be implemented jointly; identity-gated access to the privacy environment (such as ZK-KYC, or identity-bound proofs), encrypted transaction state, selective disclosure mechanisms, integrity proofs based on commitments and nullifiers, and governance arrangements that enable lawful access under defined procedures.

A privacy stablecoin may instead fall within quadrant B—high privacy but non-compliant—when privacy is enforced at the protocol level, but identity gating, selective-disclosure mechanisms, or lawful-access pathways are absent, whether through omission or deliberate design.[56]

Less intuitively, a privacy stablecoin may fall even within quadrant C—compliant but privacy-poor—if privacy mechanisms are incorrectly implemented or if privacy boundaries are incomplete. This

[55] One can thus argue that, at least partially, the difference between legal scalability and constrained privacy lies in the perspective taken, more than in the tools involved.

[56] Examples can include private L2s with ZK-shielding but no KYC binding, systems where only the user holds the viewing keys, privacy designs inspired by older privacy coins (Monero-like note models) but embedded into stablecoins, and ZK-rollups with full anonymity sets and no compliance layer.

possibility is often underappreciated. For example, generating correct ZK proofs is technically demanding, and implementation errors can inadvertently expose sensitive information. Similarly, a rollup may encrypt transaction amounts yet still leak metadata, thereby undermining effective privacy even though formal compliance requirements are met.

Finally, quadrant D—systems that provide neither privacy nor compliance—should, in principle, be impossible for any coherent privacy-preserving stablecoin architecture.

It is useful to compare this framework with ordinary L1 stablecoins, which operate on transparent ledgers with pseudonymous addresses and unencrypted state. Such instruments generally fall within quadrant C: they are compliant through issuer-level oversight but provide little privacy to users. Quadrant D is also possible in fully peer-to-peer transfers conducted outside regulated venues, where neither privacy nor compliance is reliably maintained. Conversely, it is conceptually impossible for ordinary L1 stablecoins to fall within quadrants A or B, since they lack both protocol-level privacy and built-in compliance mechanisms.

Table 1

A simple taxonomy of constrained privacy and auditability and regulatory friendliness

	System auditable/regulation friendly	System <u>not</u> auditable/regulation friendly
Constrained privacy achieved	**A.** Achievable with privacy stablecoins provided certain criteria are met: e.g., L2 with ZK-KYC, selective disclosure, identity-bound proofs	**B.** Privacy systems with strong cryptography but no regulated access pathways: possibly attractive to users, unacceptable to regulators
Constrained privacy <u>not</u> achieved	**C.** Transparent systems with full KYC/AML compliance: regulators are satisfied, users lack privacy (e.g., transparent L1 stablecoins)	**D.** Neither privacy nor compliance; system is anonymous and opaque (e.g., mixers, non-credentialed private pools)

5.4 The no-questions-asked (NQA) principle

We have deferred consideration of whether a privacy-preserving stablecoin can satisfy the NQA principle because, in our view, this is where the analysis shifts from predominantly objective criteria to more subjective and behavioral considerations—even though a meaningful objective component remains. Earlier, we offered a taxonomy based on discrete system properties; in contrast, the adoption of the NQA principle is best understood as occurring along a continuum rather than at a binary threshold.

The objective component concerns the denomination of the peg. A stablecoin pegged to the U.S. dollar can be regarded as a digital analog of the corresponding fiat currency. Consequently, any user operating in the United States—or in a jurisdiction that is formally or informally dollar-denominated—may reasonably accept such a stablecoin in exchange with minimal due diligence (Section 3). In this setting, the instrument can be said to satisfy the NQA principle, subject to the caveat that confidence in the underlying reserves must remain intact.

Conversely, a user located in a non-USD jurisdiction faces foreign-exchange risk and additional cognitive costs when converting between currencies. For such a user, a USD-pegged stablecoin may not satisfy the NQA principle. A plausible intermediate case is that of large corporations accustomed to hedging FX exposures, which may accept such a stablecoin as part of ordinary business operations.

The prevailing—and more subjective—component concerns the broader question of what leads individuals to accept a monetary instrument without material due diligence. Put differently: when does a particular brand, technology, or object become so entrenched that most people use it without reflection? Although objective characteristics—such as durability, functionality, or ease of use—certainly matter, the adoption of widely used products often hinges on social dynamics. We wear jeans not solely because of their physical qualities but because nearly everyone around us does as well.

This reasoning links directly to the literature on critical mass theory, which examines how behavior originating within a minority group, social movement, or technological niche can reach a tipping point and trigger self-sustaining, population-wide adoption (Granovetter, 1978; Macy, 1990; Marwell and Oliver, 1993).[57] A recurring theme in this literature is the existence of a threshold—commonly estimated at roughly 15–30 percent of a population—beyond which adoption accelerates rapidly. Progress toward this threshold may be gradual; once reached, diffusion to the majority can occur quickly.

It is also important to recognize that in systems governed by social diffusion, the product that becomes dominant need not be flawless. Rather, it must be affordable, sufficiently reliable, and compatible with existing behaviors and constraints. Switching to a technically superior alternative may be costly precisely because individuals are embedded in socioeconomic networks.

The building blocks of critical mass theory can be framed around three behavioral elements: (i) visibility, the sustainable choice must be easy for others to observe; (ii) legitimacy, the behavior must appear sensible and not excessively burdensome; (iii) reciprocity, individuals must perceive that their peers are making similar choices and that their own adoption will be socially reinforced.[58]

A full treatment of critical mass theory lies beyond the scope of this paper. Our purpose here is simply to emphasize that forecasts regarding the future use of privacy-preserving stablecoins cannot rest solely on legal, financial, or technological arguments. Once a visible and meaningful minority incorporates such instruments into routine economic activity—and this behavior becomes socially legitimate—it is entirely plausible that, within a relatively short period, a majority of adults will come to regard privacy-preserving stablecoins as a natural option for at least some everyday transactions alongside other payment instruments. Nor does it detract from this process that other solutions may be technically superior: what becomes mainstream does not need to be the best solution.

6. Developments in privacy-preserving stablecoin solutions [59]

A necessary premise is that this remains a nascent market segment. Whereas standard stablecoins are estimated to account for roughly 30 percent of all on-chain crypto-asset transaction volumes, the privacy stablecoin sub-segment appears to represent well below 1 percent of the market, possibly as low as 0.1 percent of the total.[60] These figures should be regarded as indicative, rather than precise:

---

[57] Notably, relevant contributions in the field long predate contemporary technological developments, including those in the ICT domain. This suggests that the psychological foundations underlying individuals' tendency not to question certain behaviors or purchasing habits are deeply entrenched.

[58] We follow here the classification presented on the Sustainability Directory website.

[59] In this section we outline ongoing commercial developments while deliberately refraining from naming specific brands, except for already established stablecoins. This paper is not intended to serve any marketing purpose, and neutrality is appropriate given the author's affiliation with a public institution. The qualitative considerations presented here—particularly those concerning the potential role of authorities in fostering such developments—draw on a broad reading of BIS (2022), BIS (2024), FATF (2024), McKinsey (2025a).

[60] See the TRM website; data based on transactions carried out from January to July 2025.

major blockchain analytics providers do not yet recognize a dedicated category for "privacy stablecoins" and no systematic or standardized data collection exists (to the best of our research). As an additional note of caution, given the rapid pace of innovation in the crypto-asset ecosystem, any dataset even one year old—such as the estimates above—may already have been overtaken by emerging design patterns or market developments.

Unsurprisingly for a niche and still-immature segment, the pace of innovation is high. New privacy-preserving concepts and prototypes appear regularly, while several projects are being refined or tested for real-world viability. Accordingly, the overview provided below should be understood as a work in progress, not a definitive taxonomy.

Current market solutions can be grouped into two broad categories:

(i) native privacy-preserving stablecoins, where confidentiality is embedded directly in the instrument, and
(ii) privacy-enhancing execution layers, which allow established stablecoins to circulate privately within a shielded environment.

Category (i) comprises stablecoins that are private by design. Here, privacy is an intrinsic property of the asset rather than an optional feature. Each token circulates exclusively within a privacy-preserving environment, and the issuer or protocol jointly engineers value stability and confidentiality as core, inseparable features. This integration can impose constraints: combining monetary stability and privacy within a single protocol may limit opportunities for specialization, potentially making the solution less robust or less operationally efficient than an architecture in which stability and privacy are provided by distinct layers.

Category (ii), by contrast, does not introduce new stablecoins. Instead, it enables standard fiat-backed stablecoins (e.g., USDT, USDC) to be transferred within a shielded environment, such as a ZK rollup or a privacy-preserving smart-contract pool. Privacy arises only once the asset enters this environment. The same stablecoins can thus circulate in two modes: transparently, on the public blockchain; or privately, within the privacy-enhancing layer.

From an AML perspective, this distinction is substantial. In privacy-enhancing execution layers, auditability and regulatory visibility remain feasible at the entry and exit points of the shielded environment. Users may opt for either private or public transfer paths depending on convenience, a flexibility that is operationally attractive but can create monitoring challenges. Native privacy stablecoins, by contrast, provide asset-level privacy across the entire lifecycle. There are no transparent perimeters or transition points: once issued, the token operates in a fully private state. This difference—privacy at the asset layer versus privacy at the execution layer—is critical for understanding where AML obligations can meaningfully apply.

In earlier sections, we argued that privacy-enhancing mechanisms could, under certain technical and regulatory conditions, satisfy the preferences of payment service providers, users, and public authorities, more effectively than ordinary stablecoins. One might therefore expect rapid uptake of these solutions. Yet, as noted at the outset of this section, their market relevance remains limited.

A tentative list of factors that constrain wider use includes both barriers to adoption and technological differentiation factors. In outlining these obstacles, we focus primarily on the second of the two categories listed above, (ii). This reflects a partly subjective but economically intuitive judgment: a new native privacy stablecoin faces the difficult challenge of succeeding against entrenched incumbents such as USDT and USDC, whose liquidity and network effects are hard to replicate. The

more promising path for privacy-preserving payments is therefore to allow established stablecoins to circulate within privacy layers.

A first barrier concerns the complexity of the stability mechanism at the execution-layer level. Even when monetary stability is provided externally by USDT or USDC, the privacy layer must ensure full one-to-one convertibility between the public (unshielded) version of the stablecoin and its shielded counterpart. This requires a reliable proof system and careful smart-contract engineering. These tasks are technically demanding and prone to subtle implementation errors.

A second barrier arises from the regulatory climate, especially following the 2022 action by the U.S. Treasury's Office of Foreign Assets Control (OFAC) involving the Tornado Cash smart contract. This episode is widely regarded as a watershed moment. Any system offering strong on-chain privacy may be perceived as bearing similarities to the Tornado Cash model, generating apprehension among developers, investors, and exchanges. This "regulatory chill" can limit the resources allocated to developing privacy-oriented Layer-2 circuits or smart contracts, independently of the intrinsic merits of these technologies.

Technological differentiation also plays a role. Execution-layer solutions differ considerably. A fundamental divide exists between architectures such as ZK rollups—where proof generation can be computationally intensive and integration with custodians is still evolving—and shielded smart-contract pools, where privacy relies more directly on the correctness of the on-chain logic. This fragmentation can further split liquidity and complicate user experience, reducing adoption.

At the technical level, the representation of the private state introduces further complexity. In many designs, user balances are not stored as explicit account balances; instead, they are encoded as cryptographic commitments, nullifiers, or other ZK-based constructs. While these mechanisms provide confidentiality, they also create operational vulnerabilities, especially if proof systems fail or if implementation errors occur.

A final barrier concerns the tension between privacy and DeFi composability. Established stablecoins such as USDT and USDC owe much of their success to their near-universal acceptance as collateral across decentralized finance and centralized exchanges. Private layers typically must either restrict such interactions or implement ZK circuits for each supported application. Both options impose substantial technical and economic costs. A privacy layer that cannot support private swaps, private lending, or private liquidity provision will struggle to match the utility offered by transparent versions of the same stablecoin.

Against this background, what emerges is that privacy layers built atop established fiat-backed stablecoins face primarily engineering and coordination challenges. Advances in ZK proofs, shared liquidity infrastructure, and privacy-preserving compliance frameworks could reasonably reduce these frictions over time. Ultimately, the decisive variable is likely to be regulatory acceptance, which would in turn catalyze investment and experimentation, accelerating the technological improvements needed for these systems to scale. If supervisors converge on models that permit privacy under verifiable auditability, the widespread use of established stablecoins within shielded execution environments becomes a plausible trajectory. Conversely, absent such convergence, privacy-enhanced stablecoin systems will likely remain niche.

7. Concluding remarks

In this paper, we have examined why stablecoins operating on public Layer-1 (L1) blockchains have so far achieved limited use in everyday retail payments, despite the attention they have attracted. We

argue that this outcome reflects structural constraints rather than insufficient demand or cost inefficiency. At the core of this diagnosis lies the no-questions-asked (NQA) principle, under which a mainstream payment instrument is accepted in exchange without material due diligence—a property that is not currently satisfied by L1 stablecoins.

Three conditions must be jointly satisfied to operationalize the NQA principle at scale: dimensional scalability, legal (regulatory) scalability, and constrained privacy. Current stablecoin arrangements fail along all three dimensions.

In terms of dimensional scalability, the throughput of L1 blockchains remains orders of magnitude below that of established payment systems, while latency is comparatively high. In terms of legal scalability, compliance is largely implemented through off-chain overlays, requiring resource-intensive forensic analysis to link pseudonymous addresses to real-world identities—an approach that does not scale. With respect to privacy, L1 systems produce misaligned outcomes: they offer insufficient confidentiality for ordinary users, whose transaction histories can be deanonymized (albeit often requiring non-trivial analytical effort), while enabling forms of obfuscation that are incompatible with the requirements of regulated intermediaries.

These limitations point to a broader conclusion: compliance mechanisms that rely on ex-post verification are inherently dominated by architectures in which compliance is embedded directly into protocol design. The challenge, therefore, is not to enhance existing overlays, but to develop payment systems that are compliant-by-design.

Against this background, we have argued that stablecoins deployed on Layer-2 (L2) infrastructures with privacy-preserving features ("privacy stablecoins") offer a credible alternative. These arrangements combine a stable-value asset with a transaction layer that provides encrypted execution, selective disclosure, and verifiable compliance. Implemented in high-throughput environments and anchored to L1 for settlement and data availability, they allow privacy-by-design to coexist with auditability and regulatory assurance.

Our analysis shows that, in principle, such systems can meet the operational requirements for achieving NQA at scale.

With respect to dimensional scalability, L2 rollups substantially narrow the performance gap with established payment infrastructures, offering a credible path toward mass-market throughput and low-latency execution, subject to continued advances in data-availability bandwidth and proof-generation technologies.

With respect to legal scalability, privacy-preserving architectures can incorporate cryptographic mechanisms that enable protocol-native compliance. These include identity-bound proofs, selective disclosure, and permissioned access to encrypted transaction environments. However, two important qualifications apply. First, such outcomes require deliberate design choices and do not arise automatically. Second, even well-designed systems leave residual challenges: many mechanisms rely on user-assisted disclosure; fully private environments may raise concerns about excessive concealment; and users may combine compliant activity within the system with unrelated behavior in fully anonymous settings, complicating supervisory oversight.

Privacy stablecoins can also satisfy the requirement of constrained privacy, combining user-level confidentiality with targeted, due-process-based regulatory access. This balance, however, depends critically on correct implementation. Errors in zero-knowledge proof systems or incomplete privacy boundaries may result in unintended information leakage—an often-underappreciated risk in technically demanding environments.

An additional condition for mainstream adoption concerns denomination. A stablecoin must be aligned with the relevant unit of account of its users. The current dominance of USD-denominated stablecoins implies exchange-rate risk outside the United States and other dollarized economies, limiting their potential to function as mainstream payment instruments in those jurisdictions. Privacy-preserving stablecoins are therefore more likely to gain traction globally if issued in multiple fiat denominations.

Taken together, these considerations suggest that privacy stablecoins offer a credible pathway toward scalable, KYC/AML-compliant systems. If the limited adoption of current stablecoins reflects structural constraints in meeting the NQA principle at scale—rather than insufficient demand—then privacy-preserving architectures have the potential to overcome these constraints.

This outcome is not, however, predetermined. The technological design space underpinning privacy-preserving payment systems is broad—reflecting the trade-offs identified above—and rapidly evolving. As such, market forces alone do not necessarily converge on architectures that optimally reconcile innovation with full regulatory compliance. This reflects the multiplicity of feasible technical pathways, only some of which embed compliance-by-design in a robust and scalable manner.

More concretely, left to themselves, market forces are likely to converge toward configurations that are sufficiently compliant to enable major participants to operate without significant reputational concerns, while offering a level of privacy attractive to users. In seeking this balance between compliance requirement and user demand for privacy, such configurations, however, need not be fully aligned with all regulatory objectives and may leave residual vulnerabilities, including in relation to anti-money-laundering (AML) requirements.

In this sense, the underlying issue is one of coordination: market forces are not guaranteed to converge toward architectures that optimally reconcile innovation with the full set of regulatory constraints. A purely passive, wait-and-see regulatory stance therefore carries the risk of entrenching suboptimal outcomes, thereby expanding the scope for money laundering. A more proactive approach is thus warranted to help coordinate the development of the ecosystem.

Public authorities should engage in structured dialogue with technology developers, financial intermediaries, and market participants to help steer innovation toward architectures that satisfy both privacy and regulatory requirements. Such engagement can reduce uncertainty and support capital allocation toward solutions that are more likely to be viable within regulatory frameworks.

From the perspective of private actors, regulatory clarity is equally critical. Investment in privacy-preserving infrastructures entails significant fixed costs and technological risks; credible signals that compliant architectures will be recognized within supervisory regimes are therefore essential to unlock sustained development and adoption.

References

Paper and reports

Adrian T. and T. Mancini Griffoli (2021). The rise of Digital Money. *Annual Review of Financial Economics*, Vol. 13, pp. 57–77.

Anderson R.G. and R.H. Rasche (2001). The Remarkable Stability of Monetary Base Velocity in the United States,1919-1999. *The Federal Reserve Bank of St. Louis, Working Paper*, 2001–008A.

Arner D., R. Auer and J. Frost (2020). Stablecoins: risks, potential and regulation. *BIS Working Papers*, no. 905.

Auer R., R. Böhme, J. Clark and D. Demirag (2025). Privacy-enhancing technologies for digital payments: mapping the landscape. *BIS Working Papers*, no. 1242.

Ben-Sasson E., I. Bentov, Y. Horesh and M. Riabzev (2018). Scalable, transparent, and post-quantum secure computational integrity. Available at: https://starkware.co/wp-content/uploads/2022/05/STARK-paper.pdf.

Ben-Sasson E., A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, M. Virza (2014). Zerocash: Decentralized Anonymous Payments from Bitcoin. 2014 *IEEE Symposium on Security and Privacy*.

Biryukov A. and S. Tikhomirov (2019). Deanonymization and linkability of cryptocurrency transactions based on network analysis. 2019 *IEEE European Symposium on Security and Privacy* (EuroS&P).

BIS (2022). *Project Genesis 2.0. Smart Contract-based Carbon Credits attached to Green Bonds.*

BIS (2024). *Project Mandala. Streamlining cross-border transaction compliance*.

BIS (2025). The next-generation monetary and financial system. *Annual Economic Report*, pp. 77-114.

Bolt W., V. Lubbersen and P. Wierts (2022). Getting the balance right: Crypto, stablecoin and CBDC. *De Nederlandsche Bank NV, Working Paper*, no. 736.

Borgonovo E., S. Caselli, A. Cillo, D. Masciandaro and G. Rabitti (2021). Money, privacy, anonymity: What do experiments tell us? *Journal of Financial Stability*, Vol. 54.

Buterin V. (2023). The Limits to Blockchain Scalability. Available at: https://vitalik.eth.limo/general/2021/05/23/scaling.html.

Camera G. (2001). Dirty Money. *Journal of Monetary Economics*, Vol. 47, pp. 377–415.

Chaliasos S., J. Ernstberger, D. Theodore, D. Wong, M. Jahanara and B. Livshits (2024). SoK: What Don't We Know? Understanding Security Vulnerabilities in SNARKs. Available at: https://arxiv.org/abs/2402.15293,

Chaliasos S., D. Firsov, B. Livshits (2025). Towards a Formal Foundation for Blockchain Rollups. Available at: https://arxiv.org/abs/2406.16219.

Chaudhary A. (2023). zkFi: Privacy-Preserving and Regulation Compliant Transactions using Zero Knowledge Proofs. Available at: https://arxiv.org/abs/2307.00521,

Croman K., C. Decker, I. Eyal, A. Efe Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. Gün Sirer, D. Song, and R. Wattenhofer (2016). On Scaling Decentralized Blockchains. Available on request at: https://www.researchgate.net/publication/292782219_On_Scaling_Decentralized_Blockchains_A_Position_Paper.

Dabbagh M., K.-K. R. Choo, A. Beheshti, M. Tahir, N.S. Safa (2021). A survey of empirical performance evaluation of permissioned blockchain platforms: Challenges and opportunities. *Computers and Security*, Vol. 100.

Duffie D., O. Olowookere, and A. Veneris (2025). A Note on Privacy and Compliance for Stablecoins. SSRN Working Paper. Available at: https://dx.doi.org/10.2139/ssrn.5242230.

ESMA (2025). *On the provision of certain crypto-asset services in relation to non-MiCA compliant ARTs and EMTs*.

Esmaili M. and K. Christensen (2025). Performance Modeling of Public Permissionless Blockchains: A Survey. *ACM Computing Surveys*, Vol. 57(7), pp. 1–35.

FATF (2024). *Targeted update on implementation of the FATF standards on virtual assets and virtual asset service providers*.

Giammusso S. and M. Nardelli (2025). A Vision for Evolving Secret Sharing in Regulation-aware Private Payments. Available at: https://www.researchgate.net/publication/399057174_A_Vision_for_Evolving_Secret_Sharing_in_Regulation-aware_Private_Payments

Gorton G.B. (2020). The Regulation of Private Money. *Journal of Money, Credit and Banking*, Supplement to Vol. 52 (S1), pp. 21–42.

Gorton G. and G. Pennacchi (1990). Financial Intermediaries and Liquidity Creation. *The Journal of Finance*, Vol. 45(1), pp. 49–71.

Gorton G.B. and J.Y. Zhang (2023). Taming Wildcat Stablecoins. *The University of Chicago Law Review*, Vol. 90(3), pp. 909–971.

Granovetter M. (1978). Threshold Models of Collective Behavior. *The American Journal of Sociology*, Vol. 83(6), pp. 1420–1443.

Gross J., J. Sedlmeir and S. Seiter (2022). How to design a compliant, privacy-preserving fiat stablecoin via zero-knowledge proofs. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4331465&__cf_chl_f_tk=K3zs3vtvttxHlaCP_7qGnxkSmIzkMNB6gsyMQjlTa4w-1782895874-1.0.1.1-phA.KxVpP0h.y6DOF3EUpBJS1cI99laNQ096_9p7ck4.

Groth J. (2016). On the Size of Pairing-based Non-interactive Arguments. In: *35th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Vienna, Austria, Proceedings, Part II.

Habib A. (2024). Analyzing Performance Bottlenecks in Zero-Knowledge Proof Based Rollups on Ethereum. Available at: https://arxiv.org/abs/2503.22709.

Heimbach L., Y. Vonlanthen, J. Villacis, l. Kiffer and R. Wattenhofer (2025). Deanonymizing Ethereum Validators: The P2P Network Has a Privacy Issue. In: *Proceedings of the 34th USENIX Security Symposium*, August 13–15, Seattle, WA, USA.

IMF (2025). *Understanding Stablecoins*.

Kahn C.M. (2018). Payment systems and privacy. *Federal Reserve Bank of Saint Louis Review*, Vol. 100, pp. 337–344.

Kahn C.M., J. McAndrews and W. Roberds (2005). Money is privacy. *International Economic Review*, Vol. 46(2), pp. 377–399.

Kahn C.M. and W. Roberds (2009). Why pay? An introduction to payments economics. *Journal of Financial Intermediation*, Vol. 18, pp. 1–23.

Li C., P. Li, D. Zhou, W. Xu, F. Long and A. Yao (2018). Scaling Nakamoto Consensus to Thousands of Transactions per Second. Available at: http://arxiv.org/abs/1805.03870.

Lyons R.K. and G. Viswanath-Natraj (2020). What keeps stablecoins stable? *NBER Working paper series*, no. 27136.

Macy, M. W. (1990). Learning Theory and the Logic of Critical Mass. *American Sociological Review*, Vol. 55(6), pp. 809–826.

Mahrous A., M. Caprolu and R. Di Pietro (2025). Stablecoins: Fundamentals, Emerging Issues, and Open Challenges. Available at: https://arxiv.org/abs/2507.13883.

Marwell G. and P.E. Oliver (1993). *The Critical Mass in Collective Action*. Cambridge University Press, Cambridge, UK.

McKinsey (2025a). *The stable door opens: How tokenized cash enables next-gen payments*.

McKinsey (2025b). *The 2025 McKinsey Global Payments Report - Competing systems, contested outcomes*.

Mitchell J., A.L. Thompson, M.D. Carter and R.S. Nguyen (2024). Legal Implications of Decentralized Payments. Mimeo.

Mssassi S. and A.A. El Kalam (2025). The Blockchain Trilemma: A Formal Proof of the Inherent Trade-Offs Among Decentralization, Security, and Scalability. *Applied Sciences*, Vol. 15(1).

Nadler, M. and F. Schär (2023). Tornado cash and blockchain privacy: a primer for economists and policymakers. *Federal Reserve Bank of St. Louis Review*, Second Quarter, pp. 122-136.

Nainwal A., A. Kamble and N. Awathar (2025). A Comparative Analysis of zk-SNARKs and zk-STARKs: Theory and Practice. Available at: https://arxiv.org/abs/2512.10020.

Nardelli M., F. De Sclavis and M. Iezzi (2025). A Hitchhiker's Guide to Privacy-Preserving Digital Payment Systems: A Survey on Anonymity, Confidentiality, and Auditability. Available at: https://arxiv.org/abs/2505.21008.

Nasir M.H., J. Arshad, M.M. Khan, M. Fatima, K. Salah, R. Jayaraman (2026). Scalable blockchains — A systematic review. *Future Generation Computer Systems*. Vol. 126, pp. 136–162.

Reno S. and K. Roy (2025). Navigating the Blockchain Trilemma: A Review of Recent Advances and Emerging Solutions in Decentralization, Security, and Scalability Optimization. *Computers, Materials and Continua*, Vol. 84(2), pp. 2061–2119.

Rochet J.-C. and J. Tirole (2006). Externalities and Regulation in Card Payment Systems. Review of Network Economics, Vol. 5(1), pp. 1–14.

Sarencheh A., A. Kiayias, and M. Kohlweiss (2023). PARScoin: A Privacy-preserving, Auditable, and Regulation-friendly Stablecoin. Available at: https://dl.acm.org/doi/10.1007/978-981-97-8013-6_13.

Song H., Z. Qu and Y. Wei (2024). Advancing Blockchain Scalability: An Introduction to Layer 1 and Layer 2 Solutions. Available at: https://arxiv.org/abs/2406.13855.

Sun N., Y. Zhang and Y. Liu (2022). A Privacy-Preserving KYC-Compliant Identity Scheme for Accounts on All Public Blockchains. *Sustainability*, Vol. 14.

Villar J.L. (2025). Zero-Knowledge Proofs Notes. Available at: https://web.mat.upc.edu/jorge.villar/doc/notes/DataProt/zk.pdf.

Websites

Chainlink, "What Is a Zero-Knowledge Proof?", 29 June 2024, https://chain.link/education/zero-knowledge-proof-zkp.

Cloudflare, "A (Relatively Easy To Understand) Primer on Elliptic Curve Cryptography", 24 October 2013, https://blog.cloudflare.com/a-relatively-easy-to-understand-primer-on-elliptic-curve-cryptography/.

Coincodex,"Layer-1 Performance: Comparing 6 Leading Blockchains", 8 April 2025, Layer-1 Performance: Comparing 6 Leading Blockchains | CoinCodex.

Cryptos, "Layer 2 Crypto Explained: Base vs Arbitrum vs Optimism 2025, 20 November 2025, Layer 2 Crypto Explained: Base vs Arbitrum vs Optimism 2025.

Cyfrin, "Calldata in Solidity", https://www.cyfrin.io/glossary/calldata-solidity#:~:text=calldata%20is%20a%20non%2Dmodifiable,function%20calls%20in%20smart%20contracts.

Forklog, “Parallelization: what it is and how it scales blockchains”, 4 March 2024, https://forklog.com/en/parallelization-what-it-is-and-how-it-scales-blockchains/.

Gate Learn, “The Current State and Future Outlook of Ethereum Layer2: A Sober Reflection Behind the Prosperity”, 29 September 2025, https://www.gate.com/learn/articles/the-current-state-and-future-outlook-of-ethereum-layer2-a-sober-reflection-behind-the-prosperity/12580.

HashtagWeb3, “Understanding Network State in Blockchains”, 13 March 2026, https://hashtagweb3.com/understanding-network-state.

Helius, “Zero-Knowledge Proofs: An Introduction to the Fundamentals”, 20 August 2024, https://www.helius.dev/blog/zero-knowledge-proofs-an-introduction-to-the-fundamentals.

Medium, “Data Availability Layers: A Comparison”, 12 December 2024, https://sunriselayer.medium.com/data-availability-layers-a-comparison-5188da1a97b8.

Nacha (2026). “Overall ACH Network Volume. Growing fast. Going Strong”, https://www.nacha.org/content/ach-network-volume-and-value-statistics.

Paar C., “Introduction to Cryptography”, https://www.youtube.com/watch?v=2aHkqB2-46k&list=PL3boZvi-wmN6r4HUGUpRSk5uhEcTNfjSS.

Polygon. “Finality”, https://docs.polygon.technology/pos/concepts/finality/finality/.

Rickard M. “Elliptic Curve Cryptography for Beginners”, 27 March 2022, https://mattrickard.com/elliptic-curve-cryptography.

Spark. “Blockchain Speed Comparison: TPS and Finality Across 15+ Chains”, Blockchain Speed Comparison: TPS and Finality Across 15+ Chains | Spark.

SDK Finance, “High-Volume Transactions: Lessons from the Largest Payment Systems”, 18 September 2025, https://sdk.finance/blog/high-volume-transactions-lessons-from-the-largest-payment-systems/.

Sundararajan S. (2026). Stablecoin Transactions Soared 72% in 2025, Hit $33T With USDC in Lead. yahoo!finance website, https://finance.yahoo.com/news/.

Sustainability Directory, “Critical Mass Theory”, 25 October 2025, https://lifestyle.sustainability-directory.com/term/critical-mass-theory/.

TRM, 2025 Crypto Adoption and Stablecoin Usage Report, 21 October 2025, 2025 Crypto Adoption and Stablecoin Usage Report | TRM Labs.

Annex

A.1 Velocity-of-circulation ratios

We defined four baseline quantities:

- (stock) (a) combined capitalization of Tether and USDC: USD 262 billion
- (flow) (b) estimated transaction volume in Tether and USDC: USD 31,600 billion
- (stock) (c) U.S. M2 monetary aggregate: USD 22,582 billion
- (flow) (d) transaction volume handled by the ACH network: USD 93,000 billion

Stocks refer to end-2025 levels, while flows are measured over the same year.

The ratios of flow to stock—conceptually analogous to measures of velocity—for Tether and USDC, on the one hand, and for M2, on the other, together yield:

[1] $$(b/a) \div (d/c) \cong 29$$

where *(b/a)* approximates the velocity of stablecoins and *(d/c)* approximates the velocity of M2.

To make the terms more comparable, data on banknotes should be treated symmetrically in the stock measure (c) and the flow measure (d). Since there are no reliable statistics on transactions in US dollar banknotes, as a first adjustment we deduct the outstanding stock of U.S. banknotes (USD 2,432 billion) from M2 and repeat the calculation of [1], to obtain a ratio of 26.

As a further adjustment, it should be borne in mind that the ACH network is not the only option for settling retail payments in the US. The figure of USD 93 trillion therefore captures only part of the total volume of non-cash retail transactions in instruments included in the U.S. M2 monetary aggregate. However, to the best of our knowledge, no reliable estimates of the market share of the ACH network are available.

That said, the ACH network is widely regarded as the dominant system for recurring payments, direct deposits and business-to-business transactions. To assess the robustness of the result, we assume that this network accounts for 80 percent of total non-cash transaction flows. Under this assumption, the ratio in [1], excluding banknotes, would decline to approximately 21 times. Even this figure remains implausibly high.

A.2 Tasks of the sequencer and prover on L2

This annex provides a simplified description of the roles of the sequencer and prover in ZK rollup architectures.

The sequencer receives users' transactions, orders them, and executes them locally to compute the updated L2 state. It then aggregates transactions into a batch and forwards this batch to the prover. However, the sequencer is not trusted to enforce correctness: it can propose any state transition, including an invalid one.

The prover is the component responsible for verifying correctness. It receives the sequencer's proposed state transition, encodes the rules of the L2 into a constraint system (typically a SNARK or STARK circuit), and attempts to generate a zero-knowledge proof attesting that all transactions in the

batch were executed correctly. When successful, it submits the new state root together with the succinct validity proof (SNARK or STARK) to Ethereum.

The generation of the zero-knowledge proof takes place off-chain rather than on Ethereum (L1). However, as noted in the main text, each ZK rollup relies on a verifier smart contract deployed on L1, which serves as the ultimate gatekeeper. This contract verifies the proof (rather than re-executing individual transactions), checks the consistency of public inputs (such as the batch hash), and updates the canonical rollup state root if the proof is valid. These steps are enforced deterministically by the L1 protocol.

Once the proof is successfully verified on L1, the updated L2 state root becomes final, withdrawals to L1 become claimable, and the batch is considered settled and immutable. Conversely, is the proof is invalid, the transaction is rejected at the smart contract level and the proposed state root is not accepted.

A.3 Zero-knowledge proofs in simple terms: the Ali Baba cave example[61]

In the Ali Baba cave illustration, Peggy (the prover) discovered the secret word that opens a magic door deep inside a cave. Victor (the verifier) wishes to verify whether Peggy indeed knows this secret word, without Peggy revealing it explicitly. Peggy, for her part, seeks to keep the secret confidential and avoid disclosing any additional information.

The cave has a circular shape with two paths, A and B, that meet at the magic door. Peggy enters the cave and chooses either path A or path B at random. Victor remains outside and cannot observe which path she has chosen. Once Peggy is inside, Victor calls out “A!” or “B!”—also at random—and asks Peggy to return via the specified path.

If Peggy truly knows the secret word, she can always comply with Victor’s request. If she initially chose the same path Victor names, she could simply return along the same route. If she chose the other path, she must open the magic door to switch sides before exiting.

Thus, an individual who knows the secret word can always emerge from the requested exit. By contrast, someone who does not know the secret word can only succeed only with probability 1/2 in each round: they can comply only when Victor happens to name the path initially selected.

If Victor and Peggy repeat this procedure a sufficiently large number of times, the probability that Peggy succeeds in every round purely by chance becomes negligible—approximately 0.000003% after 25 rounds. In other words, after enough repetitions, Victor can be almost certain that Peggy knows the secret word, even though he never learns the word itself, nor has Peggy revealed it to any third party.

A.4 Elliptic-curve hardness[62]

To understand the basics of Elliptic Curve Cryptography (ECC), a useful starting point is the concept of a trapdoor function. This is a mathematical function that is easy to compute in one direction (from

---

[61] The example is associated with Quisquater, Guillou and Berson (1990). We follow a standard presentation, as summarized in the Wikipedia entry.

[62] The aim of this annex is to offer a simple introduction to trapdoor functions, the RSA algorithm, and Elliptic Curve Cryptography (ECC). It draws on explanatory material from the websites of Cloudflare, Helius, and Rickard. A more technical yet still accessible discussion is provided in the cryptography lectures of Prof. Paar.

input to output) but extremely hard to reverse unless one possesses special, secret information—the "trapdoor".

A simple analogy illustrates the idea. Suppose Bob wishes to send a confidential message to Alice. Because the message is confidential, Bob wants to encrypt it so that—even if a malicious third party intercepts it—the original content remains hidden.

The procedure works as follows. Both Bob and Alice possess a private key. In addition, Alice has a public key, which she shares with Bob using any ordinary communication channel.[63] Using Alice's public key, Bob encrypts his message—for example, "the cat is on the table"—and obtains an unreadable output such as "t70g2g8tfrbnh8t". This encrypted output is what Bob sends over the network.

When Alice receives it, she uses her private key to decrypt the ciphertext and recover the original message. A third party intercepting the encrypted text would not be able to reverse it into readable form unless they had access to Alice's private key. Reversing the function without this secret information is computationally unfeasible, based on what current computing resources could achieve in any reasonable amount of time.

A traditional and still widely used form of public-key cryptography based on trapdoor functions is the RSA algorithm.[64] RSA relies on the fact that, given two large prime numbers p and q, it is easy to compute their product $n = p\,q$ but extremely hard to recover p and q from n alone.[65] This "one-way" property underpins RSA's security.

The main limitation of RSA is that, as computational power increases, the size of n must increase accordingly to maintain security. Modern secure RSA keys must be very large, and such key sizes are difficult to accommodate efficiently on resource-constrained devices such as smartphones, embedded computers, and cryptocurrency networks.[66]

ECC is based on a different trapdoor function but achieves comparable security with much shorter keys. In practical terms, ECC typically requires markedly less storage and bandwidth than RSA—roughly a factor of 10 to 15—to reach the same security level. This makes ECC especially attractive for applications where efficiency, speed, and reduced data footprint are essential.

An elliptic curve is the set of points that satisfy a mathematical equation of the type:

[1] $$y^2 = x^3 + ax + b \qquad \text{with } \Delta = 4a^3 + 27b^2 \neq 0$$

Two properties are useful to highlight. First, elliptic curves are symmetric about the horizontal axis: each point (x, y) of the curve has a corresponding point $(x, -y)$. Second, any non-vertical line intersects the curve in at most three places. Using these properties, a notion of "addition" of points on the curve can be defined: given two points P and Q, one can obtain a third point R such that

[2] $$P + Q + R = 0$$

Because of the symmetry property, this can be rewritten in the usual way as

---

[63] "Public" in *public key* does not mean that Alice must broadcast the key to everyone. It only means that the key does not need to be kept confidential. Alice may send it to Bob through an ordinary channel (for example, e-mail); even if a third party intercepts it, the security of the system is not affected. Only the private key must remain secret.

[64] The acronym is derived from the family names of Ron Rivest, Adi Shamir and Leonard Adleman, who publicly described the algorithm in 1977.

[65] The full RSA scheme involves additional mathematical steps, which are omitted here for simplicity.

[66] From a technical standpoint, the primary limitation stems from the computational cost of RSA signature generation and verification, rather than from key-storage requirements.

[2] $P + Q = -R$

where $-R$ denotes the reflection of R across the x-axis.

A further step is needed to obtain the version of elliptic curves used in cryptography. Instead of working over the real numbers, ECC works over a finite field, typically defined by taking all integer values between 0 and a large prime number p. This is where modular arithmetic—sometimes called "clock arithmetic"—comes into play.[67]

When Q approaches P along the curve, the line through them becomes the tangent at P. This allows the addition rule to extend naturally to the case where $P = Q$. Thus, the operation $P + P$ (the "addition" of the point with itself) becomes $2\,P$, which is the simplest instance of scalar multiplication. The core difficulty exploited by Elliptic Curve Cryptography lies in determining the integer used in this multiplication—that is, recovering how many times the base point has been added to itself.

In brief, ECC relies on three main elements: (i) a large prime number p defining the size of the finite field; (ii) a specific elliptic-curve equation over that field; and (iii) a public base point G of large prime order n.[68] A user's private key is a randomly chosen integer d with $1 \leq d \leq n$, and the corresponding public key is obtained by computing the scalar multiple $Q = dG$, meaning that it is added to itself d times according to the elliptic curve group law.

Recovering the private key d from the public key Q would require determining how many times the base point was added to itself. This is the elliptic-curve discrete logarithmic problem, which is believed to be computationally intractable with current technology. This hardness assumption underpins the security of ECC.

[67] For example, in modular arithmetic modulo 17, the sum of 12 and 15 is 10: $(12 + 15) mod\ 17 = 10$. In ordinary arithmetic, 12 + 15 is 27; dividing 27 by 17 leaves a remainder 10, which is the value of interest.

[68] The base point G can be thought of as a "starting point" on the curve from which all valid public keys are generated by repeatedly adding G to itself. Its order n is simply the number of times one can add G before the sequence loops back to the starting point. The fact that this number is a very large prime ensures a vast space of possible keys and prevents cycles or structural shortcuts that attackers could exploit.