BIS Working Papers

No 1359

The anatomy of stablecoin transactions

by Fabian Schär, Anneke Kosse, Tara Rice, Takeshi Shirakami and Jirapat Siridhasanakul

Monetary and Economic Department

June 2026

JEL classification: E42, O33, G28, C81, G23

Keywords: blockchain, payments, policy and regulation, stablecoins, transaction complexity

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

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

The Anatomy of Stablecoin Transactions

Fabian Schär[1,2], Anneke Kosse[3], Tara Rice[3], Takeshi Shirakami[3], and Jirapat Siridhasanakul[4]

[1]Faculty of Business and Economics, University of Basel
[2]Swiss Finance Institute
[3]Committee on Payments and Market Infrastructures, Bank for International Settlements
[4]Formerly Committee on Payments and Market Infrastructures, Bank for International Settlements

Working Paper Version: June 10, 2026

Abstract: Stablecoin transfers are often interpreted as payments. On programmable blockchains, however, they are frequently embedded in atomically executed transaction bundles that combine trading, lending, arbitrage, liquidity provision, and settlement. We show that ignoring this structure materially distorts the interpretation of stablecoin activity. Using 593 million event logs from 141 million Ethereum transactions involving three major U.S. dollar stablecoins, we develop a replicable framework to measure transaction complexity from archive node data, public contract labels, and event signatures. The analysis combines measures of token and contract co-usage, action type, computational complexity, urgency, and timing. Two results emerge. First, complexity is a first-order feature of stablecoin activity: nearly 60 percent of transfer events occur within complex transactions. Second, the three stablecoins are not used interchangeably: their use differs systematically across transaction structures, urgency, and timing, consistent with distinct institutional designs and economic functions. Analyses that treat transfers as standalone payments therefore risk misclassifying a large share of on-chain stablecoin use, with implications for empirical measurement, market monitoring, and policy.

JEL Codes: E42, O33, G28, C81, G23.
Keywords: Blockchain, Payments, Policy and Regulation, Stablecoins, Transaction Complexity.

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements or the Committee on Payments and Market Infrastructures. The authors are grateful to Remo Nyffenegger, Dario Thürkauf, and Dan Rees as well as the participants of research seminars at the Bank for International Settlements, the Bank of Canada and the Oesterreichische Nationalbank for helpful comments and suggestions. The authors used AI/ML-supported tools for proofreading, code testing, and debugging. All remaining errors are ours.

1 Introduction

Stablecoins have emerged as a core component of blockchain-based financial systems. Beyond facilitating peer-to-peer payments, they serve as settlement assets for often complex financial transactions, as well as sources of on-chain liquidity within smart contract-based financial markets. As their economic footprint expands, a precise understanding of how stablecoins are used in practice becomes central to assessing their implications for financial intermediation, market structure, and financial stability.

A growing literature examines stablecoins in the context of cross-border payments, monetary policy transmission, and systemic risk, often interpreting stablecoin transfers of value as economically analogous to simple payment transactions such as consumer payments, remittances, and simple settlement flows. While such uses are economically important, this perspective overlooks a defining feature of blockchain-based financial systems: the ability to compose multiple financial operations into a single, atomically executed transaction. Smart contracts allow payments, asset swaps, collateral adjustments, risk transfers, and other financial operations to be bundled into inseparable sequences that either execute jointly or fail together. In this setting, a stablecoin transfer often represents only one element of a broader economic interaction. Interpreting stablecoin transfers in isolation can therefore provide a misleading picture of their economic role, overstating payment activity and obscuring the function of stablecoins as inputs into more complex financial operations.

A central contribution of this paper is to formalize the distinction between stablecoin transfers and the transactions in which they are embedded. A transaction specifies the bundled sequence of operations to be executed; transfers, by contrast, are event logs emitted by the stablecoin contract itself, each recording an individual change in token ownership. Transactions thereby provide the economic context that links discrete on-chain steps into a coherent financial operation, while transfers capture only the granular asset flows that occur within it. This distinction is not merely taxonomic. A given transfer may reflect very different economic content: it can constitute a payment between independent parties, an intermediate leg of a multi-step settlement, or a purely technical movement executed for internal accounting, collateral rebalancing, or routing within a protocol. Treating transfers as standalone observations therefore conflates economically distinct events and can introduce material bias into measures of stablecoin usage, transaction volume, and on-chain velocity. We mitigate this bias by recovering the transactional context surrounding each transfer.

The paper's distinction between simple payments and more complex transactions is crucial for policy and standard setting work by central banks and securities regulators to enhance the safety and efficiency of financial market infrastructures (FMIs). In 2012, the Committee on Payments and Market Infrastructures (CPMI) and the International Organization of Securities Commissions (IOSCO) issued the Principles for Financial Market Infrastructures (PFMI), a set of international standards for FMIs (CPMI-IOSCO, 2012). Within this framework, the complex stablecoin transactions highlighted in this paper have significant policy relevance. In the FMI context, many of them would be considered analogous to "money settlement", that is, settlement conducted by an FMI with or between its participants to settle financial obligations arising from wholesale market transactions.

In contrast to retail payments (e.g., consumer payments, remittances or business to business payments), wholesale money settlements are systemic in nature. Reflecting this systemic importance, PFMI Principle 9 sets out risk management requirements, including those related to the quality of settlement assets used for money settlement. In July 2022, CPMI-IOSCO issued guidance on how to apply the PFMI to stablecoin arrangements, including guidance on Principle 9, setting out expectations for stablecoins used as settlement assets (CPMI-IOSCO, 2022). Understanding how, and to what extent, stablecoins are used in transactions analogous to money settlement is therefore important for policy considerations. This paper offers valuable insights in this regard, indicating that a significant share of stablecoin activity relates to such transactions.

Analyzing and understanding this broader transaction context is also where the economic stakes lie. Atomic execution and composability may substitute for intermediated settlement, mitigate counterparty risk, and enable forms of contingent contracting that are difficult to replicate on traditional financial platforms. Stablecoin usage therefore reflects not only demand for digital money, but also demand for programmable settlement and on-chain liquidity. Understanding the complexity of stablecoin transactions is thus a first-order question for financial economics research.

This paper examines the structural composition of stablecoin transactions and analyzes how stablecoin transfers are connected to other blockchain event logs. We construct a novel dataset comprising more than one half of a billion event logs and the associated transaction-level information for three U.S. dollar-denominated stablecoins. Rather than treating stablecoin transfers as standalone observations, we analyze their interaction with other tokenized assets and smart contract-based financial protocols, and we develop quantitative measures of transaction complexity that capture the depth, sequencing, and interdependence of on-chain operations.

Our analysis yields two main findings:

First, a substantial share of stablecoin activity extends beyond simple payments, and the magnitude of that share depends fundamentally on whether one observes transactions or transfers. At the transaction level, 31.6 percent of all stablecoin transactions generate multiple event logs and exhibit computational complexity well above that of a simple value transfer. The degree of complexity varies markedly across these transactions, ranging from straightforward bilateral atomic asset swaps to highly intricate operations involving more than 1,000 event logs and the coordinated transfer of dozens of assets across multiple counterparties and smart contract-based financial protocols. Because such complex transactions typically generate many stablecoin transfer events each, the picture at the transfer level differs sharply: 59.96 percent of all stablecoin transfers occur within complex transactions rather than as standalone value transfers. This wedge between roughly one third of transactions and nearly two thirds of transfers is itself a direct consequence of the distinction we formalize, and it carries first-order implications for empirical work. Because empirical analyses of stablecoins typically begin from transfer-level data, interpreting each transfer as a standalone payment misclassifies almost six in ten transfer events, overstates both stablecoin transfer counts and transferred volumes, and can yield misleading inferences about activity concentration and the economic role of stablecoins. A precise characterization of transaction structure, and of the function that stablecoins fulfill within those transactions, is therefore essential for empirical research on stablecoins.

Second, the three stablecoins in our sample are not used interchangeably. Each exhibits distinct patterns with respect to transaction complexity, co-usage, urgency, and timing, indicating meaningful differentiation within the on-chain financial ecosystem. The three stablecoins differ in how deeply they are embedded in smart contract-based financial infrastructure, in the set of tokens and protocols with which they typically co-occur, in the urgency with which their transactions are executed, and in the alignment of their activity with regional business hours. The differences are economically substantive, particularly for PYUSD: they reflect heterogeneity in institutional design, regulation, user base, and functional role rather than incidental variation.

Together, these findings recast stablecoins as constitutive elements of an emerging programmable financial platform rather than as digital analogues of traditional payment instruments. The reframing has substantive consequences for empirical work that links stablecoin activity to monetary aggregates, capital flows, and financial conditions, as well as for policy and regulatory frameworks that increasingly take stablecoins as objects of oversight. In each of these settings, inference is sensitive to assumptions about what a stablecoin transfer represents and to the comparability of different stablecoins. The analytical framework developed in this paper provides a basis for addressing both.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 provides the technical background on tokens and on the types of blockchain data used in the study. Section 4 describes the data collection process and the construction of the dataset. Section 5 presents the empirical analysis, beginning with patterns of token and smart contract co-usage, then turning to an action-set classification of transactions, measures of computational complexity, the urgency with which transactions are executed, and the business-hour alignment of stablecoin activity. The final section concludes.

2 Literature Review

The first stablecoin (BitUSD) was introduced in 2014, but, at the time, received little attention among the general public. In 2019, Libra was announced; it was a wake-up call to central banks (G7 Working Group on Stablecoins, 2019). The nascent research and policy work on Libra (and other stablecoins) lacked robust data and sufficient background. Since then, the FSB has consistently identified data gaps, including insufficient regulatory reporting, transparency, and cross-border coordination, as critical risks in the stablecoin and crypto-asset ecosystem. Their primary focus has been the need for robust data on reserves, user redemption rights, and interconnections with traditional finance (Financial Stability Board, 2019, 2023). To fill these gaps, researchers have employed a range of empirical methods.

Our paper sets out to tackle this shortcoming by measuring on-chain stablecoin activity without presuming what that "real" economic activity and usage would look like. We build on a rapidly evolving literature that (by necessity) has often relied on vendor-provided or model-based classifications. We instead read the signal embedded in the intrinsic complexity of transaction structures, and our classification draws only on inputs that can be independently checked: standardized event signatures from a public registry, which are cryptographic hashes that can be recomputed and verified, and the membership of exchange and lending protocols, obtained directly from their on-chain factory contracts. Token names are read from the standardized ERC-20 interface using our own archive node, and the only third-party labels we use serve to attach human-readable names to non-token contracts in the descriptive co-usage analysis. Critically, we do not rely on proprietary or otherwise non-verifiable attributions of end-user addresses. The resulting framework is transparent and reproducible across chains and time, and complements rather than replaces existing methods.

The closest strand of literature to our work is the analysis of on-chain data to estimate activity patterns of stablecoins, including cross-border investments and real-world payments.

Reuter (2025) provides a methodology for estimating international stablecoin flows using on-chain data. The analysis contains transaction-level data from six blockchains and uses a large language model (LLM) to analyze wallet domain names for cultural or linguistic clues, and identifies wallets that frequently interact with region-specific exchanges. The author then trains a machine-learning model (gradient boosting) using behavioral patterns, such as when users transact, which exchanges they use, and what tokens or contracts they interact with to predict the geographic region of wallets and map stablecoin flows globally.

Auer et al. (2025) examine the drivers behind cross-border flows of cryptoassets and stablecoins across 184 countries from 2017 to mid-2024. They identify distinct drivers between native cryptoassets and stablecoins, highlighting their dual role as speculative investments and transfer methods. The empirical analysis is based on data from third party providers (Chainalysis and Iknaio). These providers do not include transaction-level data but instead offer aggregated flows at the quarterly and entity-to-entity level, along with a breakdown of value bands by Iknaio. The authors focus on BTC and ETH, the two largest unbacked cryptoassets, and USDT and USDC, the two largest stablecoins with a peg to the U.S. dollar.

Visa, in collaboration with Allium Labs, developed the Visa Onchain Analytics Dashboard to shed light on how fiat-backed stablecoins move via public blockchains globally. Insights from this dashboard were first published in July 2025 (Sheffield, 2025). The author argues that raw on-chain stablecoin transaction data are difficult to interpret as these transactions include significant non-payment activity such as automated trading, arbitrage and liquidity operations that do not reflect genuine economic transactions. To improve interpretability, the author applies filtering techniques that remove transactions from bots, high-frequency trading and complex smart contract interactions. This is done by combining large sets of labeled addresses (e.g., exchanges, lending and minting addresses) with heuristic filters for unlabeled activity. These filters include counting only the largest transfers within a transaction to eliminate redundant internal flows and excluding wallets with very high activity (wallets with over 1,000 transactions or USD 10 million in 30 days as these are likely associated with automated or trading activity). The results suggest that while stablecoin adoption and usage are growing, stablecoin activity remains concentrated in large-value transfers and infrastructure-related use cases, with retail payments still representing a relatively small share.

Similar to Sheffield (2025), Higginson et al. (2026) estimate the volume of stablecoin payments from on-chain data. Key differences lie in the filtering criteria and interpretation of transaction activity. Higginson et al. (2026) propose filtering and inference techniques such as tagging known payment-related infrastructure (e.g., custody providers and card programs) and applying behavioral and size-based heuristic filters (i.e., B2B transaction ranges (4,500 - 600,000 USD)). The results suggest that the volume of actual stablecoin payments is much smaller than total stablecoin activity, though growing rapidly.

Batra et al. (2026) filter stablecoin transactions from on-chain data into "real" payments by applying a three-step, behavior-based approach. First, they remove non-economic activity such as bots, protocol mechanics (minting, bridging) and intermediary routing (e.g., DEX hops). Second, they analyze wallet behavior patterns (e.g., transaction sizes, frequency, directionality (one-way vs cyclical), counterparty diversity and timing) to infer whether a transfer resembles a payment or trading activity. Third, they apply a classification rule, where only transactions that are directional, non-reversible, and consistent with recurring or purposeful transfers between distinct parties are counted as payments.

The research has branched out into many other related topics. Recent research on cross-border crypto flows finds large dollar-backed volumes with stablecoins accounting for a substantial share, distinct drivers for native assets versus stablecoins, and weakened geographic frictions alongside limited effectiveness of traditional capital flow management measures (Auer et al., 2025). Consulting and industry studies argue that gross stablecoin volumes overstate real-economy usage, with the majority of activity tied to exchange microstructure and arbitrage, yet they also document fast growth in business-to-business and remittance use cases (Batra et al., 2026; Higginson et al., 2026; Sheffield, 2025). The macro-financial research indicates that issuer portfolio rebalancing associated with stablecoin demand shocks can transmit into money and FX markets, including short-end U.S. Treasury yields and the broad dollar (Cerutti et al., 2024; Aldasoro et al., 2026). Methodologically, advances in AI-driven flow attribution and comprehensive node-level data collection (through Chainalysis and Iknaio) have improved geographic measurement and broadened chain coverage but do not include transaction-level breakdown of the data (Reuter, 2025).

Our paper differs in two ways. First, the source and structure of the data. Many, especially the earlier empirical papers, are based on exchange-level data collected from commercial data vendors (e.g., Coingecko, CCData), which means that they cover activities in the secondary market rather than on-chain activity. Examples include Ahmed and Aldasoro (2025), Aldasoro et al. (2024), Aldasoro et al. (2026), Cerutti et al. (2024), Kosse et al. (2023).

Those papers that do use on-chain data, such as Auer et al. (2025), Cerutti et al. (2024), Batra et al. (2026), Sheffield (2025), and Higginson et al. (2026), primarily use data provided by commercial data vendors (e.g., Chainalysis, Iknaio, Allium Labs, Artemis Analytics).

Second, and related, while these approaches provide valuable insights, they largely depend on assumptions about how different types of economic activity manifest on-chain. Such assumptions, whether embedded in machine learning models, rule-based filters, or tagging frameworks, are inherently difficult to validate due to the absence of ground-truth data. Hence, results may be sensitive to model specification, filtering thresholds, and classification criteria.

By measuring stablecoin usage through the intrinsic complexity of transaction structures and drawing only on publicly verifiable inputs rather than behavioral assumptions about what individual transfers represent, our approach provides a replicable and methodologically robust framework for evaluating and interpreting on-chain stablecoin activity.

3 Technical Background

This section provides the technical background necessary to interpret blockchain stablecoin data. Interpreting stablecoin transfers requires an understanding of how stablecoins are implemented at the protocol level and of what occurs when a transaction is executed on-chain. While several tokenization approaches exist, the mechanism described below is the dominant one at the time of writing and is used across most major blockchains and smart contract platforms.

Stablecoins are typically implemented as smart contracts. A smart contract is deployed at a distinct address on the blockchain and contains state variables and executable code. The code is exposed through a set of standardized functions that allow users and other smart contracts to interact with the token. These functional interfaces are standardized to ensure interoperability across tokens and across smart contract-based financial infrastructure.[1] Tokens that adhere to this standard are referred

[1] In this context, financial infrastructure is distinct from financial market infrastructures in traditional finance (CPMI-IOSCO, 2012).

to as ERC-20 compliant.[2]

The functions exposed by an ERC-20 compliant smart contract can be invoked through transactions. For example, to transfer USDC a user broadcasts a transaction targeting the USDC smart contract that specifies a call to the contract's `transfer()` function. The function arguments indicate the recipient address and the amount to be transferred. When the transaction is included in a block and executed, the smart contract verifies that the sender holds a sufficient token balance and updates its internal state accordingly: the specified amount is deducted from the sender's balance and added to the recipient's balance.

In addition to updating its internal state, the stablecoin smart contract emits an event. Events are log entries that record specific state changes during contract execution. For the analysis of stablecoin flows, the most important events are *Transfer* events. These standardized events record all balance updates and contain three fields: the sender address, the recipient address, and the transferred amount.

To illustrate, consider two transactions: a simple stablecoin transfer and a more complex transaction involving an asset swap and multiple contracts. In the simple case, Alice transfers 10 USDC to Bob. She broadcasts a transaction targeting the USDC smart contract that calls the `transfer()` function with Bob's address as the recipient and 10 as the amount. Upon execution, the USDC contract reduces Alice's balance by 10, increases Bob's balance by 10, and emits a corresponding Transfer event recording the triple `(Alice, Bob, 10)`.

In the more complex case, suppose Alice wishes to swap 100 USDC for PYUSD. She does not transact directly with another party. Instead, she broadcasts a transaction targeting a decentralized exchange contract, specifying the input token (USDC), the output token (PYUSD), the input amount, and a minimum acceptable output that limits slip-

[2] The ERC-20 standard was proposed through the *Ethereum Request for Comments* system. It was the first formal standard for fungible tokens on Ethereum Virtual Machine-based blockchains and has since become the de facto base standard for their implementation. See `https://eips.ethereum.org/EIPS/eip-20`.

page. The exchange contract executes the swap through a USDC-PYUSD liquidity pool deployed at a third smart contract address. It invokes `transferFrom()`[3] on the USDC contract to pull 100 USDC from Alice's balance into the pool, and the pool then invokes `transfer()` on the PYUSD contract to send the corresponding amount of PYUSD back to Alice.

Three contracts are therefore involved in this single transaction – the USDC and PYUSD token contracts and the pool itself – and the transaction emits at least three event logs: a USDC Transfer event recording the triple `(Alice, Pool, 100)`, a PYUSD Transfer event recording the triple `(Pool, Alice, X)` where X is the market-determined amount of PYUSD received in the swap, and a Swap event emitted by the pool that records the trade parameters. No single event captures the full economic interpretation; only the combination of events from all three contracts reveals that Alice exchanged 100 USDC for X PYUSD at the price implied by the pool's reserves. Even this three-contract, three-event example sits at the simpler end of the spectrum: many transactions on Ethereum emit dozens or hundreds of event logs, and the most complex in our sample exceed two thousand, reflecting the deep composability that smart contract platforms enable.

While event logs convey limited additional information for simple payment transactions, they are essential for analyzing more complex interactions. Composability, a defining feature of smart contract platforms, allows a contract to invoke functions of other smart contracts during its execution. This can produce chains of nested function calls involving multiple token transfers and other state changes. In such settings, examining the original transaction message alone is often insufficient to infer the economic consequences of execution. Event logs therefore play a central role in reconstructing stablecoin flows and identifying balance changes across addresses.

[3] `transferFrom()` is effectively a direct debit: it allows the exchange contract to pull tokens from Alice's balance, subject to prior approval by Alice via the `approve()` function.

Conceptually, a transaction is an initial instruction that triggers execution, supplies the computational resources for the entire process, and provides a common frame of reference for all internal calls. Events, by contrast, constitute a detailed log of the actions taken during execution. A single transaction may cause multiple smart contracts to emit events, including token contracts as well as financial infrastructure contracts such as exchanges and lending protocols. Combining information from transaction messages with event logs allows us to reconstruct stablecoin usage within complex contract interactions, to assess the economic context of individual transfers, and to evaluate the degree of transactional complexity in which these transfers are embedded.

4 Data

We analyze three well-known U.S. dollar-denominated stablecoins: Tether (USDT), USD Coin (USDC) and PayPal USD (PYUSD). These stablecoins are selected because they account for a substantial share of stablecoin activity while spanning distinct institutional designs.

USDT is the largest stablecoin by market capitalization, representing the earliest large-scale implementation of stablecoins with off-chain reserve assets. It was issued by Tether Limited, owned by Tether Holdings Limited, incorporated and registered in the British Virgin Islands until its restructuring in January 2025. Following the restructuring, it is now issued by Tether International, incorporated and registered in El Salvador.

USDC was launched by Circle in partnership with Coinbase under the Centre Consortium and is now fully controlled by Circle. Circle is licensed in the U.S. and other advanced economies, such as the EU, UK and Singapore. Circle holds U.S. state money transmitter licenses and a Virtual Currency License granted by the New York Department of Financial Services (NYDFS). USDC positions itself as a regulated, payment-oriented stablecoin and is often the preferred choice of institutional and U.S.-regulated entities.

PYUSD marks the first large-scale entry of a traditional payment service provider into blockchain-based settlement. Paxos Trust Company was regulated by the NYDFS and, following the decision by the Office of the Comptroller of the Currency (OCC) in December 2025, it is now a fully-chartered trust company regulated by the OCC. PayPal, an incumbent global payment service provider, is licensed as a money transmitter on a state-by-state basis in the U.S. and holds a Virtual Currency License.

Our analysis focuses on the Ethereum mainnet instances of these stablecoins. Restricting attention to a single blockchain ensures methodological consistency and allows us to attribute observed differences in transaction behavior to stablecoin-specific characteristics rather than to heterogeneity in execution environments. Ethereum accounts for a substantial share of global stablecoin activity and provides a standardized setting in which all three assets are implemented as ERC-20 tokens. This standardization ensures that execution and transfer mechanics are identical across tokens, supporting the interpretation of cross-stablecoin differences as economically meaningful.

Every ERC-20 compliant token transfer emits an event log that records the sender, the recipient, and the transferred amount. We extract the complete history of transfer event logs for USDT, USDC, and PYUSD for the year 2025, spanning blocks 21,525,891 to 24,136,052. In total, the sample contains more than 241 million stablecoin transfer events. Figure 1 plots transfer-event activity by stablecoin over time.

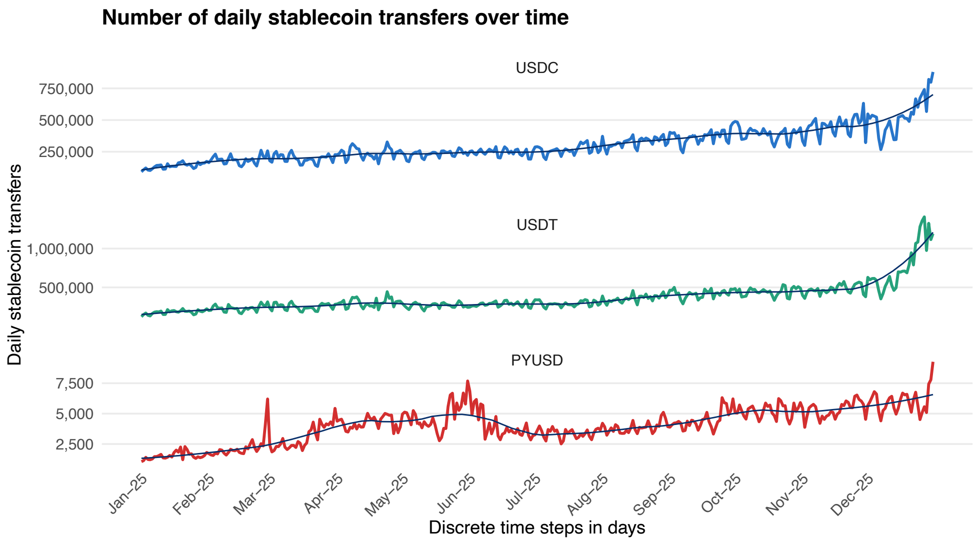

Figure 1. The figure depicts the time-series evolution of transfer-event activity for USDC, USDT, and PYUSD. Transfer counts are aggregated into equally sized block-range bins, which provide an approximation to daily transfer frequencies. The sample comprises all Ethereum transfer events observed in 2025, totaling approximately 241 million observations. The solid line corresponds to a locally weighted polynomial regression (LOESS) that smooths high-frequency variation and highlights underlying temporal trends. Importantly, transfer events capture all contract-level transfer logs and therefore should not be interpreted as peer-to-peer payment activity or user-initiated transactions.

We then use the transaction hashes associated with these transfer events to retrieve the corresponding transactions and collect all additional event logs emitted during execution. These include further token transfer events as well as events generated by smart contract-based exchanges, lending protocols, and other on-chain financial infrastructure. This approach allows us to embed each stablecoin transfer within the full set of contemporaneous contract interactions occurring within the same transaction.

The resulting dataset comprises 593 million event logs emitted by 53,919 ERC-20-compliant token contracts and 401,267 other smart contracts, across 141 million transactions. For each transaction, we additionally collect block-level information and transaction receipts, which enable the measurement of execution time, computational complexity, and transaction fees.

All blockchain data are obtained from an Ethereum archive node operated by the authors, ensuring completeness and independent verification. We further enrich the data with verified event signatures from 4byte.directory.[4] These signatures, together with the asset pool addresses we obtain directly from on-chain factory contracts (described below), are what underpin our classification of smart contract interactions. Token names are read directly from the standardized ERC-20 interface via our own node. The only use we make of Etherscan contract labels is to attach human-readable names to non-token contracts in the generalized co-usage analysis; no part of the action classification depends on them.

We query factory contracts[5] of major automated market makers (AMMs)

[4]An event signature is a deterministic cryptographic identifier of the event type emitted by a smart contract. Because all events of a given type share the same signature regardless of the emitting contract (for instance, all ERC-20 Transfer events) these signatures provide a standardized means of identifying the economic nature of an on-chain interaction. 4byte.directory is a public registry that maps these signatures, which are recorded on-chain as hashes, to their human-readable name and structure.

[5]A factory contract is a smart contract whose function is to deploy other smart contracts of a standardized type and to maintain an on-chain list of the contracts it has created. In the context of automated market makers, the factory instantiates a new pool contract each time a new trading pair is established and records the resulting address. Querying the factory therefore provides an authoritative and exhaustive list

					Obs. / Range
Panel A. Coverage and Scale					
Time span					*2025-01-01* to *2025-12-31*
Blocks					*2,610,162*
Stablecoin transactions					*141,310,373*
Stablecoin transfers					*241,423,906*
Smart contract event logs					*593,096,940*
Distinct event signatures					*13,574*
Smart contract count					*455,186*
Token contract count					*53,919*

	Min	Max	Mean	Med	SD
Panel B. Block-Level Variables					
$N = 2{,}607{,}156$					
Stablecoin transactions per block	*1*	*1,126*	*54.2*	*47*	*32.6*
Stablecoin transfers per block	*1*	*3,406*	*92.6*	*75*	*80.4*
Base fee (gwei)	*0.0087*	*826*	*2.03*	*0.55*	*7.19*
Panel C. Transaction-Level Variables					
$N = 141{,}310{,}373$					
Gas used	*37,448*	*58,605,949*	*134,458*	*62,248*	*298,993*
Effective gas price (gwei)	*0.0095*	*230,000*	*3.02*	*1.13*	*37.4*
Priority fee (gwei)	*0*	*229,993*	*1.37*	*0.20*	*36.7*
Event log count (per transaction)	*1*	*2,036*	*4.20*	*1*	*17.7*
USDT event log count (per transaction)	*0*	*1,000*	*0.925*	*1*	*4.19*
USDC event log count (per transaction)	*0*	*1,287*	*0.773*	*0*	*2.92*
PYUSD event log count (per transaction)	*0*	*67*	*0.0104*	*0*	*0.117*

Table 1. This table summarizes the constructed samples and reports descriptive statistics for the key variables used in the empirical analysis. Panel A describes the dataset's temporal coverage and overall scale; the *Blocks* row reports the total number of blocks in the sample period. Panel B reports descriptive statistics for block-level variables, computed over the subset of 2,607,156 blocks containing at least one stablecoin transaction. Panel C reports transaction-level statistics, including the distribution of event logs per transaction. The priority fee is the per-gas tip paid to the validator (or block builder), computed as the effective gas price net of the protocol-determined base fee. Gwei is a denomination of Ether, where one gwei corresponds to 10^{-9} Ether.

to obtain all relevant exchange pool addresses, and we collect addresses from the most prominent lending market contracts. This procedure allows us to categorize contract interactions into economically meaningful classes. In doing so, we attribute economic function on the basis of features that any researcher can independently reconstruct, namely standardized event signatures and the on-chain membership of protocol contracts, rather than proprietary, non-verifiable labels of externally owned (end-user) addresses, on which much of the existing on-chain literature relies. Where vendor-based studies infer the nature of activity from who is presumed to control an address, we infer it from what a transaction verifiably does on-chain.

Due to the scale and complexity of the dataset, we store and process the data in a locally managed MongoDB instance. Table 1 reports descriptive statistics for the three stablecoins over the sample period, distinguishing between block-, transaction-, and log-level observations. Figure 9 in the Appendix provides a UML-style representation of the database schema.

5 Empirical Analysis

Stablecoin transaction complexity is multi-dimensional. A transaction can be complex in how many counterparties and contracts it touches within a single atomic execution, in the type of financial action it performs, in the computational resources it consumes, or in the premium it pays for faster block inclusion. Each dimension reveals a distinct facet of stablecoin activity, varies across stablecoins, counterparty segments, and time, and admits multiple economic interpretations when read in isolation. Read jointly, the measures discipline one another. Direct on-chain measurement of complexity is therefore a precondition for the research and policy uses to which stablecoin activity is increasingly relevant, and for any meaningful comparison across stablecoins, periods, or jurisdictions.

of the trading venues that operate within a given protocol.

The empirical analysis proceeds along several complementary lines, each constructed from publicly verifiable inputs, namely standardized event signatures and protocol membership read directly from on-chain factory contracts, rather than from vendor address labels or behavioral inference, and resting on transparent structural definitions rather than ad hoc value- or activity-based filters. We characterize structural composition through token and contract co-usage, classify transactions by financial action, quantify computational burden through event-log counts and gas consumption, examine urgency through priority fee behavior, and close with intraday and intraweek timing patterns. As a baseline, we classify a transaction as a Simple Transfer if it contains exactly one stablecoin transfer event and no additional event logs. Under this definition, 68.4 percent of stablecoin transactions qualify as Simple Transfers. The remaining 31.6 percent involve additional contract interactions or multiple transfers and constitute the primary object of the subsequent analysis.

5.1 Token Co-Usage

First, we examine how stablecoins are embedded within the broader token transaction network by analyzing patterns of joint token usage at the transaction level. Rather than focusing on bilateral transfers in isolation, this approach characterizes the extent to which stablecoins are used alongside other tokens within the same transaction, capturing complementarities in on-chain activity such as trading, liquidity provision, arbitrage, and portfolio rebalancing. By quantifying the frequency with which a stablecoin co-occurs with other tokens, we obtain a granular view of the functional roles that different stablecoins play within the Ethereum ecosystem. We are therefore able to compare their degree of integration into multi-token transaction flows.

Formally, the dataset is organized as observations indexed by transaction-token pairs (tx, i), where tx denotes a transaction, i a token, and $n_{tx,i}$ the number of transfer events of token i within transaction tx. We write $i \in tx$ as shorthand for $n_{tx,i} \geq 1$, i.e., token i appears in transaction tx.

For each stablecoin in the set

$$\mathcal{S} = \{\text{USDC}, \text{USDT}, \text{PYUSD}\},$$

denoted by s, and for each other token $j \neq s$, consider the set of transactions in which both s and j appear at least once.

Within this set, the transaction-level co-usage frequency is defined as

$$C_{s,j} = \left|\{\, tx : s \in tx \ \wedge \ j \in tx \,\}\right|,$$

which counts the number of transactions in which the stablecoin s and token j are jointly used. We further denote the total number of transactions containing stablecoin s by

$$N_s \ = \ \left|\{\, tx : s \in tx \,\}\right|,$$

and report co-usage as the share $C_{s,j}/N_s$.

Figure 2 presents the top 15 relative co-usage values for each stablecoin, illustrating the distribution of token co-usage patterns for the three major USD-denominated stablecoins on Ethereum.

For USDC and USDT, Wrapped Ether (WETH) emerges as the most prominently co-used asset. The relationship is markedly stronger for USDC: transfers involving WETH appear in close to 20 percent of all USDC transactions, while the corresponding share for USDT is substantially lower, though still dominant relative to other tokens. In both cases, the second most frequently co-used asset is the counterpart stablecoin, highlighting the central role of stablecoin-stablecoin interactions. The third position is occupied by a custodial ERC-20 representation of Bitcoin, underscoring the importance of tokenized Bitcoin as a bridge asset for trading, liquidity provision, and portfolio rebalancing within the Ethereum ecosystem.

PYUSD exhibits a distinctly different co-usage profile. It does not appear among the top 15 co-used tokens for either USDC or USDT, indicating a limited presence within the transaction networks of the two dominant stablecoins. From the opposite perspective, PYUSD transactions display a high degree of stablecoin co-usage: USDC and USDT feature prominently and precede Wrapped Ether in terms of co-usage frequency. The pattern points to an asymmetric structure in which PYUSD transactions disproportionately involve USDC and USDT, while PYUSD itself remains largely peripheral from their perspective.

The apparent prominence of stablecoin co-usage within PYUSD transactions partly reflects the substantially smaller number of PYUSD transactions relative to USDC and USDT. As a consequence, PYUSD-related activity contributes only marginally to the co-usage distributions of the two dominant stablecoins, even though USDC and USDT play a central role within PYUSD transactions. Accounting for these scale differences, and benchmarking PYUSD against the co-usage patterns observed for USDC and USDT, the evidence points to a comparatively low level of integration of PYUSD into the broader token transaction network.

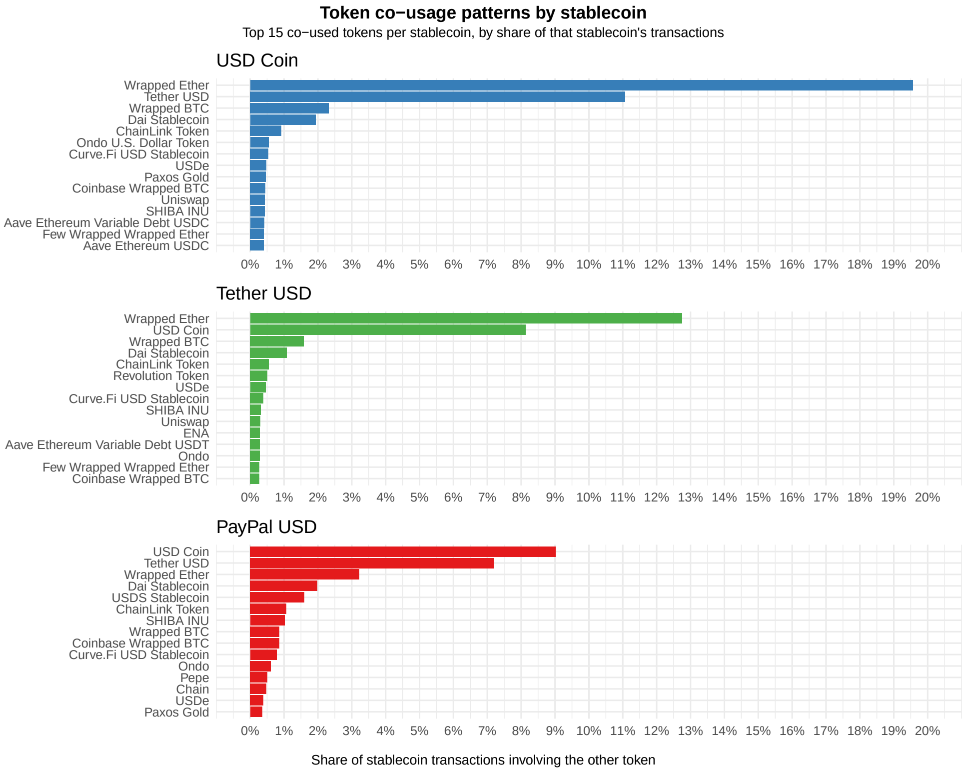

Figure 2. Co-usage patterns by stablecoin. For each stablecoin, the figure reports the most prominent co-used tokens and the share of that stablecoin's transactions in which each co-token appears within the same transaction. USDC and USDT share the same set of leading co-tokens, with the counterpart stablecoin in second place, indicating similar interaction patterns, while USDC exhibits stronger co-usage with Wrapped Ether, consistent with deeper integration into smart contract-based financial protocols. For PYUSD, stablecoins rank more prominently among the leading co-tokens, while overall co-usage is substantially lower.

5.2 Generalized Contract Co-Usage

We extend the co-usage analysis beyond token transfer events to a more general interaction setting. Rather than focusing exclusively on standardized stablecoin transfers, we consider the full set of smart contract interactions that occur within transactions involving stablecoins. This perspective captures economic usage patterns that extend beyond pure asset transfers and reflects the functional roles stablecoins play within broader on-chain activity.

In contrast to the ERC-20 token analysis, there is no standardized interface from which we can directly query contract metadata from a local node. The blockchain itself exposes only contract bytecode, not the Application Binary Interface or the underlying Solidity source code. We therefore rely on the Etherscan API to obtain human-readable names for non-token smart contracts. These labels are used only to identify the contracts reported in the co-usage results below and play no role in the action-set classification, which rests on verified event signatures and on-chain factory membership.

We then compute contract co-usage measures for stablecoin-contract pairs. The construction follows the token co-usage measure introduced above, with two modifications. First, we exclude all standardized ERC-20 transfer events from the sample. Second, the index now ranges over the full set of smart contracts rather than being restricted to token contracts. Let $m_{tx,c}$ denote the number of non-ERC-20-transfer event logs emitted by contract c in transaction tx. Then, for each stablecoin $s \in \mathcal{S}$ and each smart contract c,

$$C^{\text{ctr}}_{s,c} \;=\; \big|\{\, tx : s \in tx \;\wedge\; m_{tx,c} \geq 1 \,\}\big|,$$

and is normalized by N_s for reporting.

Figure 3 presents the smart contracts most frequently co-used with each stablecoin. Consistent with the token co-usage analysis, the Wrapped Ether contract (WETH) features prominently across all three panels. In

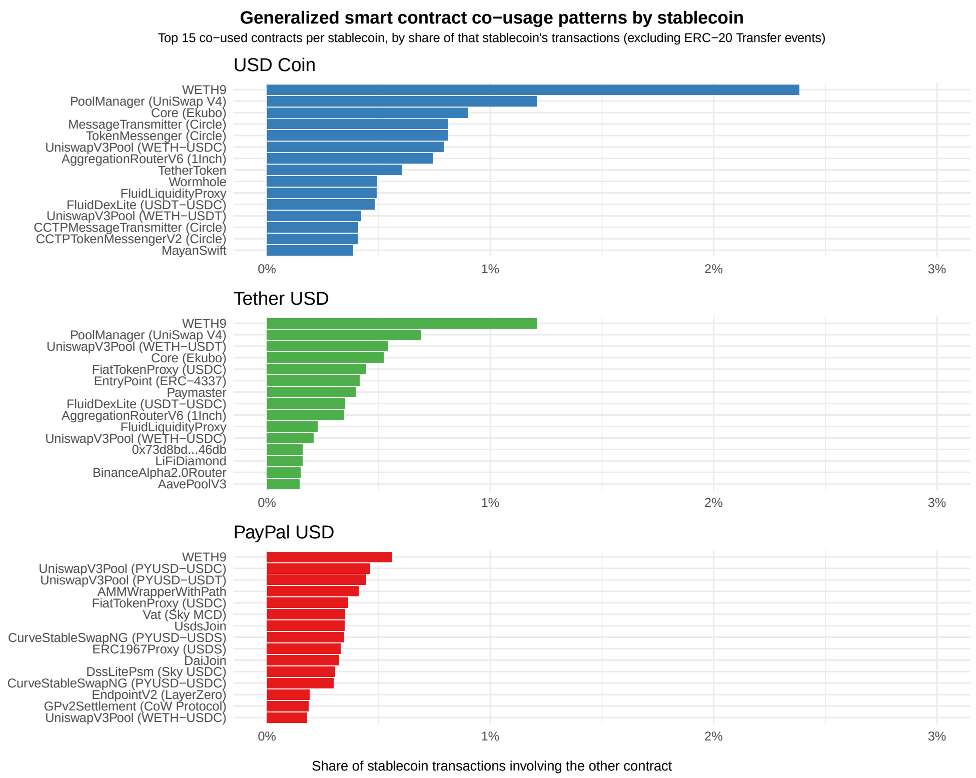

Figure 3. Generalized smart contract co-usage patterns by stablecoin. For each stablecoin, the figure reports the most prominent co-used contracts and the share of that stablecoin's transactions in which the respective contract co-appears within the same transaction. ERC-20 Transfer events have been excluded from this metric.

the present setting, this co-usage is not driven by token transfer activity, which is explicitly excluded, but by interactions related to the wrapping and unwrapping of native Ether. These operations correspond to deposits into, and withdrawals from, the WETH contract and typically occur when users introduce or remove ETH as collateral. The frequent atomic co-occurrence of such interactions with stablecoin transfers indicates that stablecoins are closely integrated with protocols that rely on Ether, which constitute some of the more decentralized components of the on-chain infrastructure.

Beyond WETH, co-usage is concentrated in smart contract-based exchange infrastructure, including automated market makers such as Uniswap and exchange aggregators such as 1inch. This pattern reflects the role of stablecoins as core trading and settlement instruments in on-chain liquidity provision and order routing. The figure also highlights the relevance of auxiliary contracts that are specific to individual stablecoins, such as messaging, bridging, and protocol-level support contracts, underscoring that stablecoin activity is embedded both in general-purpose financial infrastructure and in issuer-specific service layers.

5.3 Action Set Analysis

To obtain more granular insights into stablecoin use cases and prevailing transaction patterns, we assemble a comprehensive and systematically categorized dataset of prominent smart contract-based financial protocols and map observed on-chain activity to these contracts. To mitigate the risk of excluding economically relevant protocols, we cross-reference our protocol universe with multiple third-party databases and complement this process with a systematic inspection of event logs to identify contracts with high event-emission intensity. For each protocol, we manually identify the relevant factory contract addresses and develop scripts to programmatically extract the associated pool or pair contract addresses directly from the protocol smart contracts. This process yields a total of

555,176 distinct smart contract addresses, which we classify by economic function as either exchange or lending protocols.

Using this address set, we implement an aggregation pipeline that matches event logs to known financial protocols. We assign each transaction to one of six mutually exclusive action categories on the basis of the composition of its event logs and its interaction with known protocol addresses. Transactions classified as *Simple Transfers* contain exactly one stablecoin transfer event and no additional event logs, and reflect plain peer-to-peer payments. *Multi Transfer* transactions include more than one stablecoin transfer event but no additional event logs. The *Exchange* category captures transactions that interact with at least one smart contract-based exchange protocol, while *Lending* transactions are those that interact with at least one smart contract-based lending protocol. Transactions that involve both exchange and lending interactions are classified as *Complex Financial*. The residual category, *Other*, comprises transactions with multiple transfer and non-transfer events that do not interact with any of the identified major financial protocols and, in many cases, rely on user-specific custom smart contract code.

We then aggregate transaction-level observations into unique triplets defined by the stablecoin, the action category, and the financial protocol used. Each triplet corresponds to a distinct stablecoin-action-protocol combination. The number of transactions associated with a given triplet serves as the weight in the Sankey diagram in Figure 4, which illustrates how stablecoins are employed across economic use cases. The three layers of the diagram represent, from left to right, the choice of stablecoin, the type of economic action, and the class of protocol. The category *Multiple* captures transactions that involve more than one stablecoin at the first layer or more than one protocol at the third layer.

The figure documents substantial heterogeneity in the relative importance of Simple Transfers across stablecoins. Among more complex transactions, pure atomic asset swaps executed via smart contract-based exchange protocols dominate. Within this category, Uniswap V3 emerges as the primary venue, although transaction activity is distributed across a

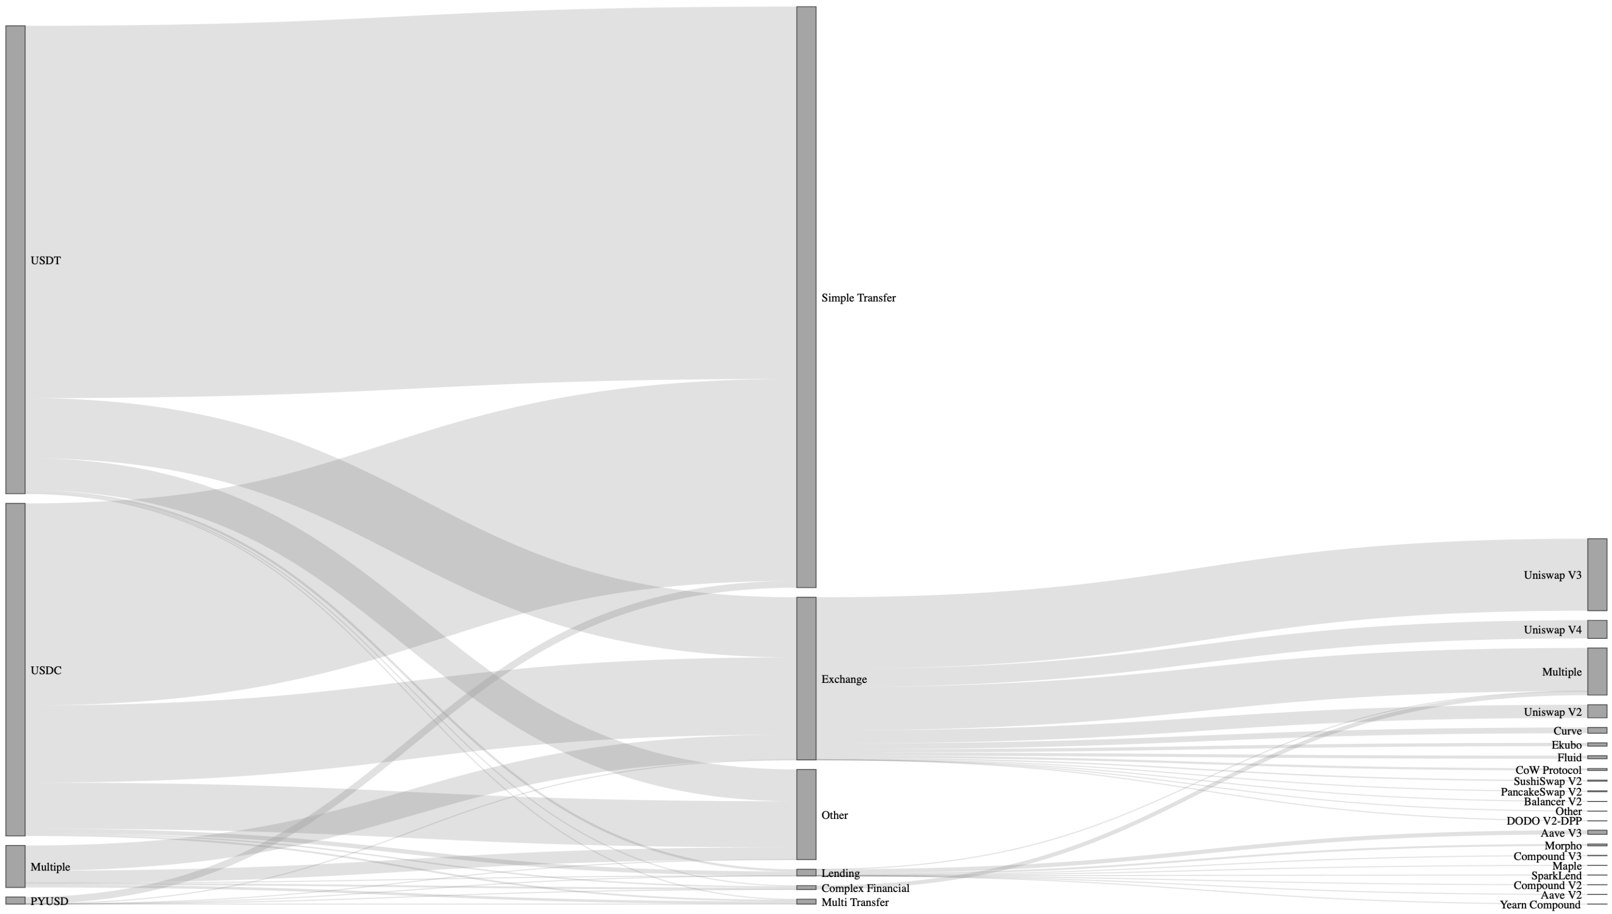

Figure 4. Transaction-level Sankey diagram illustrating dominant action categories and protocol usage by stablecoin type. The figure is based on the complete sample of more than 141 million Ethereum transactions executed in 2025 that include at least one transfer of USDT, USDC, or PYUSD. *Simple Transfer* denotes transactions with a single stablecoin transfer event and no additional event logs, corresponding to a basic payment. *Multi Transfer* refers to transactions containing multiple stablecoin transfer events, but no other interactions. *Exchange* identifies transactions interacting with at least one exchange protocol. *Lending* identifies transactions interacting with at least one lending protocol. *Complex Financial* denotes transactions that simultaneously satisfy the criteria for both lending and exchange interactions. *Other* comprises complex transactions with multiple transfer and non-transfer events that do not interact with any of the major financial protocols.

wide range of exchange protocols. Lending activity is comparatively less prominent and frequently co-occurs with exchange operations. Such cases are classified under the Complex Financial category and, by construction, necessarily involve at least one exchange and one lending protocol.

Within the subset of pure Lending transactions, Aave is the most widely used protocol. The Other category accounts for a substantial share of activity. It comprises transactions that do not rely on the dominant exchange or lending protocols but nonetheless exhibit complex smart contract interactions, including custom contract logic, bridging operations, and financial activities that extend beyond simple swapping or lending. The mean event-log count in this category is 9.67, with a maximum of 1,800 event logs observed within a single transaction, indicating a high degree of transactional complexity.

5.4 Computational Complexity

The action-set classification establishes that stablecoin transactions differ qualitatively in what they do. We now ask how they differ in how much they do, that is, in the number of computational steps they set in motion. The most immediate measure of this kind is the number of event logs a transaction emits. Each event log records a state change written during execution, so the log count of a transaction is a direct, transparent proxy for the operations bundled into a single atomically executed transaction. The measure is promising for three reasons. First, it follows mechanically from the transfer-transaction distinction we stress in this paper: a standalone payment emits a single transfer event, whereas composability manifests precisely as additional logs accumulating as a result of the same transaction. Second, it is available uniformly for every transaction in the sample and requires no protocol labeling, contract classification, or external metadata to compute. Third, it is highly informative about the shape of activity. As reported in Table 1, the event-log count per transaction has a median of one but a mean of 4.197 and a maximum of 2,036, a separation between median and mean that is the signature of a heavily right-skewed distribution in which a minority of richly composed transactions accounts for a disproportionate share of total on-chain activity.

Figure 5 develops this measure into a joint view of complexity. For each stablecoin it plots the number of that stablecoin's own token transfer events within a given transaction against the total number of event logs emitted by the same transaction, both on logarithmic scales, with shading proportional to the share of the stablecoin's transactions falling in each cell. The dashed 45-degree line marks transactions whose entire log output consists of the stablecoin's own transfers. Vertical distance above the line therefore measures additional activity executed atomically alongside the stablecoin transfer. For all three stablecoins, mass concentrates in the lower-left corner, corresponding to simple single-transfer payments that sit on or near the diagonal. The distinguishing feature is how far mass disperses upward and away from the diagonal: this dispersion is pronounced for USDC and USDT, whose transactions routinely emit tens to hundreds of logs while containing only a handful of their own transfer events, and is far more muted for PYUSD, whose mass remains close to the origin. Read through this lens, USDC and USDT are deeply embedded in composable, multi-operation transactions, while PYUSD activity is dominated by simple transfers.

The event-log count is nonetheless an incomplete measure of computational complexity. Events are emitted at the discretion of the smart contract being executed, and the decision of how many events to log, and of what kind, is an implementation-specific design choice rather than a standardized account of the work performed. Substantial computation, including internal message calls, arithmetic, control flow, as well as state reads and writes, can occur without emitting any event at all, so two transactions that consume very different computational resources may register identical log counts, and vice versa. The log count captures only the logged footprint of a transaction, not the underlying computational burden it imposes on the network. To measure that burden directly, we turn to gas usage. Gas is the unit in which the Ethereum protocol measures execution: every elementary operation carries a fixed gas cost, and the gas consumed by a transaction is the mandatory, protocol-determined sum of these costs across everything it executes (Wood, 2014). Unlike event emission, gas accounting is not optional. Given a specific set of computations, it is identical across all contracts, and is paid for by the user. It therefore provides a standardized, implementation-agnostic, and economically meaningful measure of how much computation a transaction actually performs.

Figure 6 reports the distribution of transaction-level gas consumption separately for each stablecoin. All three distributions are concentrated at low gas levels and exhibit a pronounced right skew. This shape mirrors the event-log evidence and is consistent with simple transfers dominating activity: the median gas use of 62,248 sits just above the cost of an isolated ERC-20 transfer, while the mean of 134,458 and the maximum of roughly 58.6 million reflect a long tail of computationally heavy transactions. The concentration is most extreme for PYUSD, whose transactions cluster almost entirely in the lowest gas bins, whereas USDC and USDT display visibly heavier right tails.

Figure 7 decomposes gas consumption by the action categories introduced above, plotting the distribution of gas used for each category and stablecoin. Gas consumption rises markedly and monotonically with the qualitative complexity of the action: Simple Transfers cluster tightly at the lower end, near the transfer floor; Multi Transfer, Exchange, and Lending transactions occupy an intermediate range, and Complex Financial transactions, which by construction combine exchange and lending interactions, consume the most, with medians around one million gas.

The action taxonomy is defined on the basis of event-log composition and protocol interaction, yet it orders transactions in close agreement with gas usage, a continuous measure of computational cost constructed independently of that classification. Equally notable is what the figure does not show: within any given action category, the differences across the three stablecoins are comparatively modest. Conditional on what a transaction does, the choice of stablecoin has only a limited effect on how much computation it requires, with some partial exceptions for PYUSD that we discuss next.

A modest exception is visible within the Simple Transfer category in Figure 7, where the PYUSD box sits slightly above those of USDC and USDT. This difference is unlikely to be driven by PYUSD's additional transfer checks, such as pausing or address-blocking functionality, since the other stablecoins also implement comparable compliance and control mechanisms. If such checks imposed a materially higher baseline cost for PYUSD, the lower tail of the distribution should also shift upward. Instead, the cheapest observed simple transfers are very similar across the three tokens, at 37,448 gas for USDC, 38,359 for USDT, and 38,940 for PYUSD. A more plausible explanation is compositional. ERC-20 transfers are more gas intensive when the recipient has not previously held the token, because the EVM must initialize a new balance entry rather than update an existing one. Consequently, if PYUSD simple transfers more frequently involve first-time recipients, its median or interquartile range can be slightly higher even though the minimum observed transfer cost remains close to that of USDC and USDT.

Similarly, PYUSD sits above USDC and USDT in the *Other* and *Complex Financial* categories as well. For Complex Financial the median PYUSD transaction consumes about 1.12 million gas, against roughly 0.88 million for USDC and USDT. This uptick appears to reflect composition rather than any intrinsic costliness of PYUSD interactions. PYUSD rarely enters a Complex Financial transaction on its own. When it does, the transaction is typically a larger multi-stablecoin bundle that also moves the other two assets. Among Complex Financial transactions containing PYUSD, at least three quarters also carry USDC and around half also carry USDT, whereas those built around USDC or USDT almost never contain PYUSD, with a median of zero PYUSD transfer events. The PYUSD-bearing transactions therefore involve more token transfers and protocol interactions overall, which mechanically raises their gas. This mirrors the co-usage evidence above, in which PYUSD transactions disproportionately involve the two larger stablecoins.

The aggregate heterogeneity documented above, with USDC and USDT skewed toward complexity and PYUSD toward simple transfers, is therefore primarily attributable to differences in the composition of activity across stablecoins rather than to stablecoin-specific differences in execution cost, reinforcing the interpretation of the three assets as occupying distinct functional roles.

The transaction-level measures developed in this section also pin down what complexity implies for the transfer-level perspective from which stablecoin data are typically observed. Simple Transfers account for 68.4 percent of transactions but, by definition, emit exactly one stablecoin transfer event each, contributing roughly 96.7 million of the 241.4 million transfer events in the sample. The remaining 31.6 percent of transactions emit the other 144.8 million stablecoin transfer events, so that 59.96 percent of all transfer events originate within complex transactions. The asymmetry is mechanical: complex transactions bundle many transfers each, and the most transfer-intensive of them emit hundreds. A researcher who samples stablecoin transfer events at random therefore draws an observation from a complex transaction roughly six times out of ten, even though such transactions constitute less than one third of all transactions. Interpreting each transfer event as a standalone payment thus misclassifies the majority of transfer-level observations.

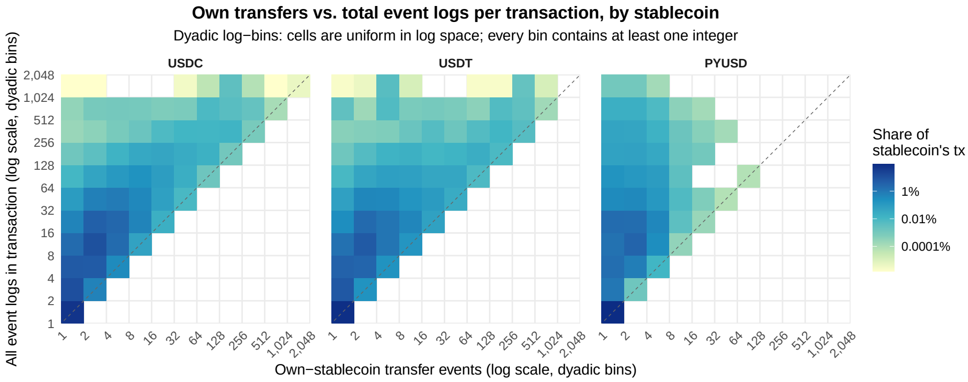

Figure 5. Joint distribution of own-stablecoin transfer events and total event logs per transaction, by stablecoin. For each stablecoin, the figure presents a heatmap in which the horizontal axis reports the number of that stablecoin's transfer events within a transaction and the vertical axis reports the total number of all event logs emitted by the same transaction, both on logarithmic scales. The plane is partitioned into dyadic cells whose edges in both axes lie at the powers of two $\{1, 2, 4, 8, 16, \ldots\}$. This binning is uniform in log space, identical across panels, and is the coarsest grid for which every integer count belongs to a non-empty cell; finer log-spaced grids would leave empty cells between consecutive integers at low counts and produce a misleading impression of missing values. Each cell is shaded by the share of the respective stablecoin's transactions that fall within it, normalized so that each panel sums to 100%, on a logarithmic color scale; darker shading indicates a higher share. The sample comprises all Ethereum transactions executed in 2025 that include at least one transfer of USDC, USDT, or PYUSD. A transaction enters the panel of every stablecoin for which it contains at least one transfer event. The dashed 45-degree line marks transactions whose entire log output consists of the stablecoin's own transfer events, so that vertical distance above the line measures the extent of additional, non-transfer activity within the transaction. Mass concentrates in the lower-left corner, corresponding to simple single-transfer payments, while the dispersion of mass upward and away from the diagonal reflects embedding within computationally complex, multi-event transactions. This dispersion is most pronounced for USDC and USDT, whereas PYUSD remains concentrated near the diagonal in the bottom-left corner.

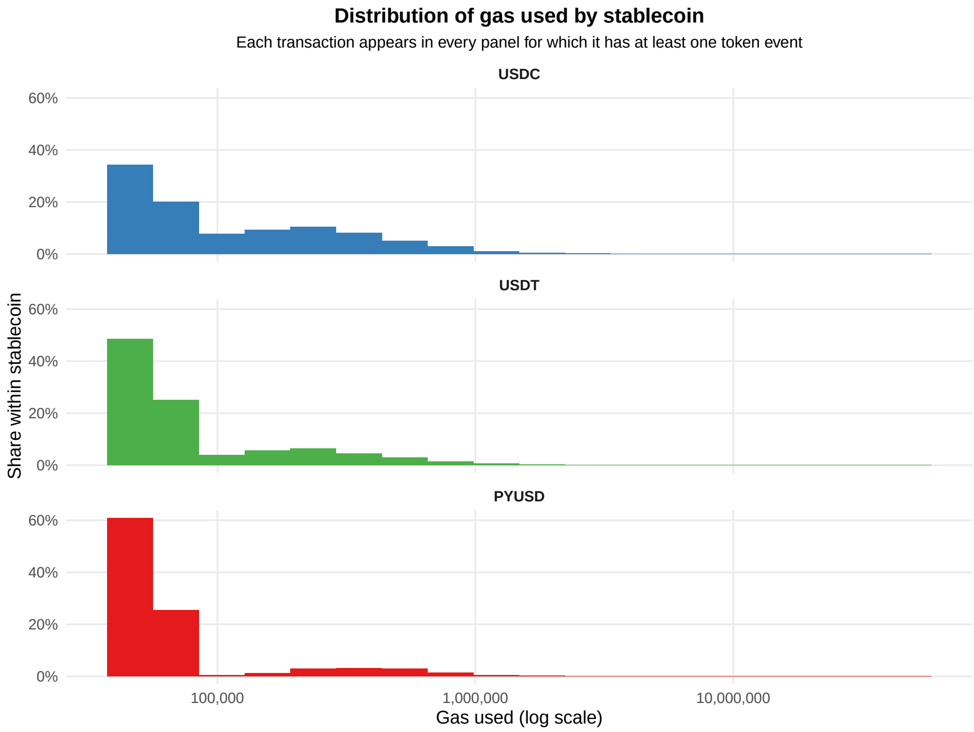

Figure 6. Distribution of gas used by stablecoin. For each stablecoin, the figure reports the relative frequency distribution of transaction-level gas consumption, with gas plotted on a logarithmic scale. The sample comprises all Ethereum transactions executed in 2025 that include at least one transfer of USDC, USDT, or PYUSD. A transaction enters the distribution of every stablecoin for which it contains at least one token transfer event, so transactions involving multiple stablecoins appear in more than one panel. Bins are constructed at uniform width in log space and are identical across the three panels. Bar heights report the share of a stablecoin's own transactions falling in each bin. All three distributions are concentrated at low gas levels and exhibit a pronounced right skew, consistent with simple transfers dominating activity. The concentration is most extreme for PYUSD, whose transactions cluster almost entirely in the lowest gas bins, whereas USDC and USDT display comparatively heavier right tails, indicating a larger share of computationally complex transactions.

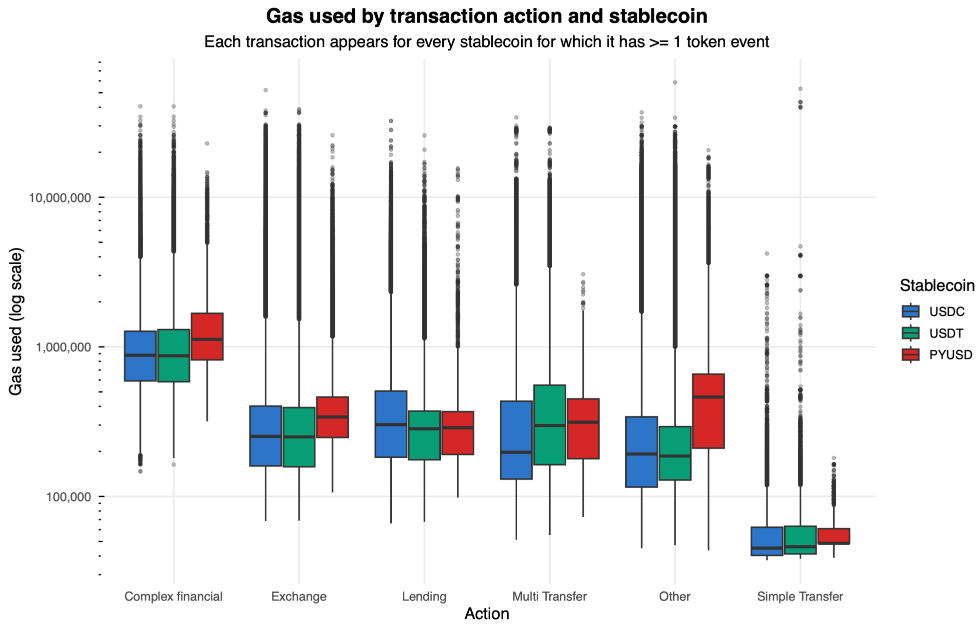

Figure 7. Gas used by transaction action and stablecoin. For each action category, the figure reports the distribution of transaction-level gas consumption separately for USDC, USDT, and PYUSD, plotted on a logarithmic scale. The sample comprises all Ethereum transactions executed in 2025 that include at least one transfer of USDT, USDC, or PYUSD. A transaction enters the distribution of every stablecoin for which it contains at least one token transfer event, so transactions involving multiple stablecoins appear in more than one box. Boxes denote the interquartile range with the median indicated by the central line, whiskers extend to 1.5 times the interquartile range, and points beyond the whiskers are plotted individually. Gas consumption increases markedly with transaction complexity, ranging from *Simple Transfer* at the lower end to *Complex Financial* at the upper end, while differences across stablecoins within a given action category are comparatively modest.

5.5 Urgency

To shed further light on economic differences across action categories, we estimate a fixed-effects regression with the priority fee, measured in gwei,[6] as the dependent variable. The priority fee is the per-unit tip paid directly to the validator (or block builder), obtained as the effective gas price net of the protocol-determined base fee. It is a sharper measure of urgency than the effective gas price, because it isolates the user's discretionary bid for execution priority from the mechanical, congestion-driven base fee that every transaction in a block pays irrespective of urgency. Being a per-unit measure, it is not mechanically related to transaction size or computational complexity. It reflects users' willingness to pay for fast inclusion. Time-sensitive transactions, such as exchange interactions, are expected to command higher priority fees than less time-critical actions.

Priority fees pose two empirical challenges that shape the estimation strategy. First, a non-trivial share of transactions, 3.4% overall, pay a zero priority fee, with the share rising to 16% for Exchange and 35% for Complex Financial transactions. A zero priority fee does not necessarily indicate low urgency. One possibility, particularly relevant for searcher and maximal-extractable-value (MEV) activity, is that the transaction compensates the block builder through an alternative channel, rather than paying through the priority fee field. Such payments are themselves recorded on-chain and could in principle be reconstructed by tracing the internal ETH transfers within the transaction, but we do not attempt that decomposition here. To the extent that some zero-tip transactions in the complex categories pay through these channels, the recorded priority fee understates their total payment for inclusion, and the complexity premia we estimate are likely lower bounds for those transactions. The implication for specification is more direct: because zero-tip transactions can be economically meaningful observations rather than missing data, we prefer estimators that retain them. Second, the priority fee distribution is heavily right-skewed. A small share of high-urgency transactions

[6] Gwei is a denomination of Ether. One gwei corresponds to 10^{-9} Ether.

pays very large tips, so estimates that target the mean of the fee differ sharply from those that target the typical paying transaction.

To accommodate both features, and to ensure that our conclusions do not hinge on any single functional form, we estimate four complementary specifications and treat the comparison across them as the main evidence. Ordinary least squares in levels delivers effects on the mean priority fee and retains zero-tip observations, but is sensitive to the heavy upper tail. The winsorized levels specification, with the priority fee capped at the 99.9th percentile, addresses that sensitivity by limiting the influence of extreme observations while preserving zeros and the levels interpretation. Ordinary least squares on the natural logarithm of the priority fee, estimated on the positive-tip subsample, instead recovers approximate proportional effects on the typical positive tip; it must drop the 3.4% of zero-tip transactions by construction. Poisson pseudo-maximum likelihood (PPML), finally, provides a multiplicative specification for the conditional mean: it retains zero-tip observations, is robust to heteroskedasticity, and its coefficients are semi-elasticities, so that a coefficient b implies a proportional change of $e^b - 1$ in the expected priority fee. The four estimators answer related but distinct questions, namely average payment, average payment robust to the upper tail, the typical positive tip, and the proportional conditional mean. Convergence across them indicates an effect that is genuinely present in the data. Divergence is itself informative about how the effect is distributed across the fee distribution, with the Complex Financial coefficient discussed below providing a leading example.

Tables 2 and 3 report the four specifications. Simple Transfers and EMEA office hours serve as the omitted reference categories for action type and time zone,[7] respectively, while USDC is the baseline stablecoin. All specifications include day fixed effects to absorb variation in aggregate economic conditions, mempool congestion, and persistent movements in the Ethereum base fee, and inference is based on heteroskedasticity-robust

[7] Office-hour buckets are defined in UTC as APAC = 00:00–07:59, EMEA = 08:00–15:59, and AMER = 16:00–23:59, and are used consistently throughout the empirical analysis.

(HC1) standard errors.

Exchange transactions command markedly higher priority fees than Simple Transfers across every specification, consistent with their time sensitivity. The PPML estimate of 1.261 in column (6) implies an expected priority fee roughly 253% above that of a Simple Transfer, and Exchange remains the largest positive coefficient in the level and winsorized specifications alike. Complex Financial transactions display a more subtle pattern. Their level and PPML coefficients are large and positive, with column (6) implying a 212% higher expected fee, but the level effect collapses under winsorization, from 1.919 to 0.444, and reverses sign in the log specification, where the typical positive Complex Financial tip is 77% below that of a Simple Transfer. This is consistent with the high average priority fee of Complex Financial transactions being a tail phenomenon: a minority of high-urgency transactions pay very large tips and lift the mean, whereas the typical complex transaction that does tip pays a below-baseline priority fee. Lending, Multi Transfer, and Other actions are associated with lower priority fees throughout, indicating systematically lower willingness to pay for immediate inclusion.

Among stablecoins, PYUSD is associated with a higher expected priority fee than USDC, by roughly 90% in the PPML specification, an effect that survives winsorization and is therefore not an artifact of extreme outliers. In the log specification, however, the typical positive PYUSD tip lies 45% below that of USDC, so PYUSD's elevated average reflects a broad right-skew across its fee distribution rather than typical behavior. USDT is economically indistinguishable from USDC in the PPML specification. These differences point to heterogeneity in user segments and transaction routing across stablecoins.

Time-of-day patterns are pronounced. Relative to EMEA hours, transactions in the APAC window pay materially lower priority fees in every specification, by approximately 19% in the PPML estimate, consistent with thinner congestion and weaker competition for inclusion during those hours, whereas AMER hours are statistically indistinguishable from EMEA. Because day fixed effects absorb day-level shifts, these are within-day intraday patterns rather than reflections of the trend.

The level specifications exhibit very low explanatory power, with an R^2 around 0.001 that rises to 0.025 under winsorization and to 0.138 in the log specification. This progression is itself informative: priority fee variation is dominated by an extreme upper tail, and the action, stablecoin, and time categories explain the typical fee far better than the mean. The coefficients should accordingly be interpreted as precisely estimated differences in central tendency, not as predictions of any individual transaction's fee. Because priority fees within a day share a common congestion and base fee environment, we additionally cluster standard errors by day. Clustering leaves all point estimates unchanged and, although it widens standard errors by a factor of three to ten, the economically large effects, Exchange, Complex Financial, PYUSD, APAC, Lending, and Multi Transfer, all remain statistically significant. Only the small Other and Multiple coefficients weaken materially, with Other no longer significant in the level specification, while USDT remains indistinguishable from USDC.

Dependent Variable:	Priority Fee (gwei)					
Estimator:		OLS (levels)			Poisson PML	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
action: Complex Financial	1.836***	1.903***	1.919***	1.114***	1.125***	1.139***
	(0.1277)	(0.1305)	(0.1305)	(0.0469)	(0.0480)	(0.0480)
action: Exchange	2.342***	2.369***	2.383***	1.240***	1.250***	1.261***
	(0.0125)	(0.0130)	(0.0130)	(0.0038)	(0.0040)	(0.0040)
action: Lending	-0.2703***	-0.2707***	-0.2605***	-0.3000***	-0.2969***	-0.2884***
	(0.0112)	(0.0113)	(0.0113)	(0.0154)	(0.0155)	(0.0155)
action: Multi Transfer	-0.5562***	-0.4821***	-0.4784***	-1.288***	-1.276***	-1.270***
	(0.0025)	(0.0114)	(0.0114)	(0.0088)	(0.0103)	(0.0103)
action: Other	-0.0910***	-0.0683***	-0.0583***	-0.1180***	-0.1091***	-0.1017***
	(0.0164)	(0.0166)	(0.0167)	(0.0204)	(0.0201)	(0.0202)
stablecoin: USDT		-0.0177**	-0.0119		-0.0044	-0.0021
		(0.0061)	(0.0061)		(0.0045)	(0.0045)
stablecoin: PYUSD		0.8075***	0.8183***		0.6344***	0.6412***
		(0.0048)	(0.0046)		(0.0035)	(0.0033)
stablecoin: Multiple		-0.1438***	-0.1397***		-0.0074	-0.0059
		(0.0232)	(0.0231)		(0.0106)	(0.0106)
tz_office: AMER			-0.0051			-0.0066
			(0.0077)			(0.0054)
tz_office: APAC			-0.2751***			-0.2113***
			(0.0078)			(0.0058)
Fixed-effects						
day	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	141,310,373	141,310,373	141,310,373	141,310,373	141,310,373	141,310,373
R^2	0.00115	0.00116	0.00117	–	–	–

Heteroskedasticity-robust (HC1) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table 2. Priority fee regressions. The dependent variable is the priority fee (the per-gas tip paid to the validator, in gwei). Columns (1)-(3) report OLS estimates in levels; columns (4)-(6) report Poisson pseudo-maximum-likelihood (PPML) estimates, which retain zero-tip transactions and whose coefficients are semi-elasticities, so that a coefficient b implies a proportional change of $e^b - 1$ in the expected priority fee. The omitted categories are Simple Transfer (action), USDC (stablecoin), and EMEA (office-hours time zone, defined in UTC). All specifications include day fixed effects and are estimated on all 2025 Ethereum transactions containing at least one USDC, USDT, or PYUSD transfer. A zero priority fee need not indicate the absence of urgency, as MEV and searcher transactions often pay validators through direct coinbase transfers rather than the priority fee field; the estimated complexity premium is therefore a lower bound.

Dependent Variable:	log(Priority Fee), positive tips			Priority Fee (gwei), winsorized 99.9%		
Estimator:	OLS			OLS		
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
action: Complex Financial	-1.572***	-1.477***	-1.454***	0.3656***	0.4312***	0.4437***
	(0.0085)	(0.0085)	(0.0085)	(0.0094)	(0.0094)	(0.0094)
action: Exchange	0.3299***	0.3958***	0.4150***	1.347***	1.383***	1.393***
	(0.0009)	(0.0010)	(0.0010)	(0.0019)	(0.0020)	(0.0020)
action: Lending	-0.7435***	-0.6542***	-0.6409***	-0.2829***	-0.2636***	-0.2562***
	(0.0041)	(0.0041)	(0.0041)	(0.0019)	(0.0019)	(0.0019)
action: Multi Transfer	-1.444***	-1.334***	-1.332***	-0.5835***	-0.5086***	-0.5061***
	(0.0036)	(0.0037)	(0.0037)	(0.0011)	(0.0020)	(0.0020)
action: Other	-1.045***	-0.9724***	-0.9582***	-0.1963***	-0.1625***	-0.1550***
	(0.0012)	(0.0013)	(0.0013)	(0.0008)	(0.0009)	(0.0009)
stablecoin: USDT		0.3011***	0.3065***		0.0493***	0.0533***
		(0.0007)	(0.0007)		(0.0008)	(0.0008)
stablecoin: PYUSD		-0.6365***	-0.6005***		0.8575***	0.8679***
		(0.0050)	(0.0050)		(0.0023)	(0.0023)
stablecoin: Multiple		0.0622***	0.0680***		-0.0793***	-0.0764***
		(0.0019)	(0.0019)		(0.0032)	(0.0032)
tz_office: AMER			-0.1710***			-0.0226***
			(0.0008)			(0.0010)
tz_office: APAC			-0.4552***			-0.2106***
			(0.0008)			(0.0010)
Fixed-effects						
day	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	136,491,405	136,491,405	136,491,405	141,310,373	141,310,373	141,310,373
R^2	0.13511	0.13658	0.13848	0.02422	0.02449	0.02483

Heteroskedasticity-robust (HC1) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table 3. Priority fee regressions, robustness. Columns (1)-(3) report OLS on the natural logarithm of the priority fee, estimated on the positive-tip subsample (zero-tip transactions, 3.4% of the sample, are dropped because the logarithm is undefined at zero); coefficients are approximate proportional effects on the typical positive tip. Columns (4)-(6) report OLS in levels with the priority fee winsorized at the 99.9th percentile, which limits the influence of the extreme upper tail. Omitted categories, day fixed effects, and standard errors are as in Table 2.

5.6 Business-Hour Alignment

Transaction timing provides an additional lens through which to study stablecoin activity. Variation across hours of the day and days of the week offers insight into the interaction of stablecoins with geographically segmented markets, heterogeneous user groups, and market infrastructures that operate on conventional business calendars. Considered jointly with measures of transaction complexity and use-case classifications, temporal patterns help clarify how different forms of stablecoin usage map into distinct timing profiles.

To isolate genuine payment usage, we conduct a subsample analysis based exclusively on Simple Transfers, excluding all more complex transactions. Figure 8 presents descriptive evidence on intraday and intraweek variation in Simple Transfer activity. Each heatmap reports the share of total Simple Transfer transactions occurring in a given hour-of-day and weekday, with all timestamps expressed in UTC. This visual representation facilitates direct comparisons of timing patterns both within and across stablecoins.

Two striking patterns emerge. First, USDT and USDC exhibit pronounced weekday effects. Transaction intensity is substantially higher from Monday through Friday and declines markedly over the weekend. Intraday activity for both stablecoins is concentrated in overlapping windows that correspond closely to standard business hours in Europe and, to a lesser extent, North America. Within weekdays, activity increases during the late morning UTC hours, peaks in the early to mid-afternoon, and gradually declines thereafter. Second, PYUSD displays a distinctly different temporal structure. Weekend effects are considerably weaker, and intraday activity is tightly concentrated in hours that align with business hours in the Americas.

Whether these temporal regularities are specific to payment activity or extend to more complex transactions is itself a question worth asking; Appendix Figures 10 and 11 present the corresponding descriptive heatmaps for two other action categories. To formalize these descriptive observa-

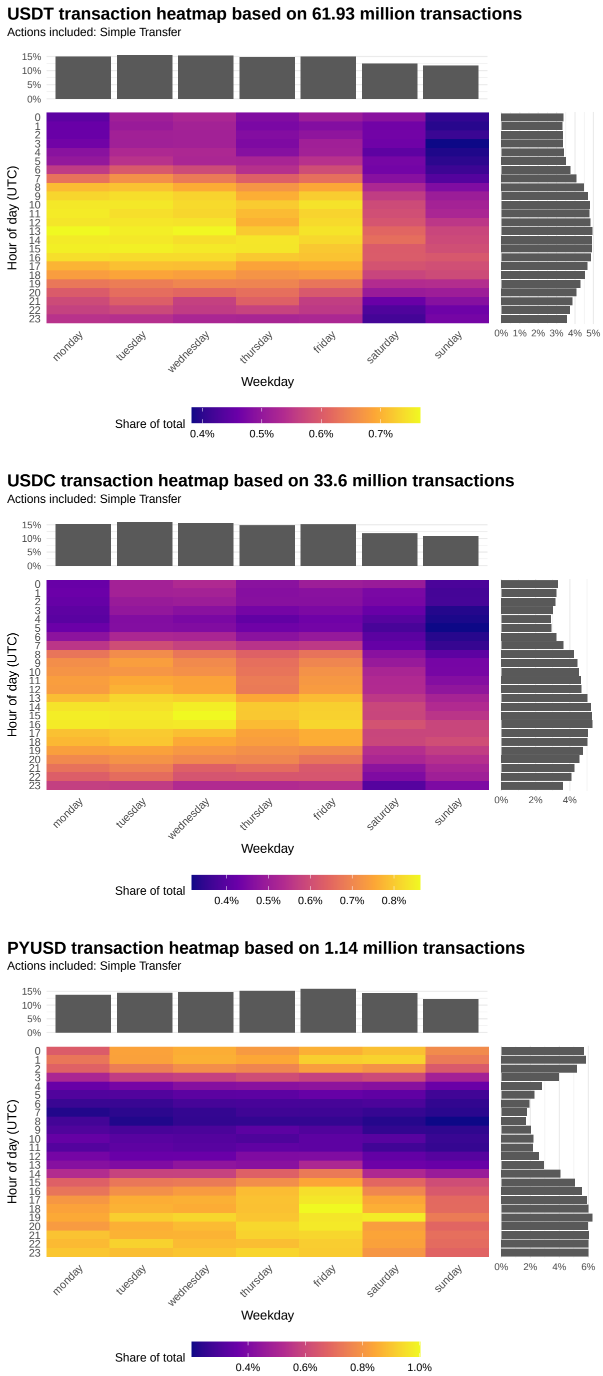

Figure 8. Heatmap of simple transfer activity by stablecoin. The figure documents pronounced heterogeneity in intraday and intraweek timing patterns across stablecoins, with particularly distinct behavior for PYUSD. Transfer activity in USDC and USDT exhibits a marked decline during weekends, consistent with lower transactional intensity outside regular business days. PYUSD displays a strong tilt toward American business (AMER) hours. In contrast, USDT and USDC are predominantly aligned with EMEA business hours, with USDC exhibiting a modest shift toward AMER relative to USDT. All times are reported in UTC.

tions and assess their statistical significance, we estimate per-stablecoin Poisson regressions of the daily transaction count on time-zone and weekend dummies, with week-of-year fixed effects to absorb slow trends such as adoption growth and longer-run shifts in network conditions, and day-clustered standard errors to accommodate correlated day-level shocks across the three time-zone cells of the same day. For each stablecoin $s \in \{\text{USDC, USDT, PYUSD}\}$ the headline specification, restricted to the Simple Transfer subsample, is

$$\log \mathbb{E}[N^s_{t,d}] = \alpha^s_{\text{week}(d)} + \beta^s_1 \cdot \text{AMER} + \beta^s_2 \cdot \text{APAC} + \beta^s_3 \cdot \text{weekend}(d),$$

where $N^s_{t,d}$ is the count of stablecoin-s Simple Transfer transactions on day d in time-zone bucket t, with EMEA-weekday as the reference cell. Multi-stablecoin transactions are excluded so each row is attributable to a single stablecoin. We use week-of-year rather than day fixed effects because day fixed effects would absorb the weekend dummy exactly; week-of-year FE absorb slow trends while leaving both tz_office and weekend identified within week. Table 4 reports the estimates.

Dependent Variable:	Daily Simple Transfer count $N^s_{t,d}$ (Poisson)		
Sample:	USDC	USDT	PYUSD
Model:	(1)	(2)	(3)
Variables			
tz_office: AMER	-0.0384^{***}	-0.1340^{***}	$+0.7400^{***}$
	(0.0069)	(0.0051)	(0.0222)
tz_office: APAC	-0.4084^{***}	-0.3116^{***}	$+0.2572^{***}$
	(0.0110)	(0.0085)	(0.0246)
weekend	-0.2808^{***}	-0.2018^{***}	-0.0988^{***}
	(0.0152)	(0.0130)	(0.0188)
Fixed effects			
week-of-year	Yes	Yes	Yes
Fit statistics			
Observations	1,095	1,095	1,095
Clusters (days)	365	365	365
Pseudo R^2	0.897	0.892	0.789

Day-clustered standard errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table 4. Headline temporal-pattern regressions per stablecoin on the Simple Transfer subsample. The dependent variable is the count of stablecoin-s Simple Transfer transactions in each (tz_office × day) cell. EMEA-weekday is the reference cell. The exponent of each coefficient is the proportional change in expected Simple Transfer count: PYUSD's AMER coefficient implies $e^{0.74} - 1 = 109.6\%$ higher AMER-hours intensity than EMEA, and the weekend coefficients translate to drops of -24% (USDC), -18% (USDT), and -9% (PYUSD). Multi-stablecoin transactions are excluded so each observation is attributable to a single stablecoin.

PYUSD is the dramatic outlier on the business-hour-alignment dimension. Its AMER coefficient of +0.740 implies an AMER-hours Simple Transfer intensity roughly 110% higher than its EMEA-hours intensity, against the opposite tilt for both incumbents: USDC's AMER intensity is −3.8% relative to EMEA and USDT's is −12.5%. Within the incumbents, USDT is more strongly EMEA-concentrated than USDC, confirming the descriptive caption's observation of USDC's modest shift toward AMER relative to USDT. Weekend activity is lower for all three (−24% for USDC, −18% for USDT, −9% for PYUSD), with PYUSD's drop the smallest of the three. APAC activity is well below EMEA for the two incumbents (USDC −34%, USDT −27%) but elevated for PYUSD (+29%).

To establish that these cross-stablecoin orderings are not artifacts of sampling variation, Table 5 reports a pooled regression with stablecoin × time interactions; the interaction coefficients are direct difference tests against the USDC reference. All are highly significant under day-clustered standard errors: USDT is 0.10 log-points more AMER-negative than USDC, PYUSD is 0.78 log-points more AMER-positive, and PYUSD's weekend drop is 0.19 log-points milder than USDC's.

Extending the sample from Simple Transfers to all action types with action as a fixed effect leaves the headline coefficients essentially unchanged (USDC AMER moves from −0.0384 to −0.0268, PYUSD AMER from +0.740 to +0.719, weekend effects within 0.05 log-points), because Simple Transfers are a substantial share of each stablecoin's transaction volume and therefore dominate the volume-weighted average. This Simple Transfer dominance, however, masks the question of whether the business-hour and weekly patterns are themselves specific to payment activity or extend across action categories. To address this we add action × tz_office and action × weekend interactions to the per-stablecoin regressions, with Simple Transfer as the reference action. Table 6 reports the implied within-action coefficients, computed as $\beta_t + \delta_{k,t} - \delta_{k,\mathrm{EMEA}}$ for tz dimension t and action k and as $\beta_{\mathrm{weekend}} + \psi_{k,\mathrm{weekend}}$ for the weekend coefficient, converted to percentage changes vs the EMEA-weekday

Dependent Variable:	Daily Simple Transfer count (Poisson, pooled)
Model:	(1)
Stablecoin main effects (vs USDC)	
stablecoin: USDT	$+0.5996^{***}$
	(0.0093)
stablecoin: PYUSD	-3.947^{***}
	(0.0291)
Time main effects (eval. at USDC; tz ref.: EMEA-weekday)	
tz_office: AMER	-0.0384^{***}
	(0.0068)
tz_office: APAC	-0.4084^{***}
	(0.0109)
weekend	-0.2835^{***}
	(0.0159)
Cross-stablecoin difference tests	
AMER × USDT	-0.0957^{***}
	(0.0049)
AMER × PYUSD	$+0.7784^{***}$
	(0.0221)
APAC × USDT	$+0.0968^{***}$
	(0.0078)
APAC × PYUSD	$+0.6656^{***}$
	(0.0243)
weekend × USDT	$+0.0832^{***}$
	(0.0120)
weekend × PYUSD	$+0.1872^{***}$
	(0.0390)
Fixed effects	
week-of-year	Yes
Fit statistics	
Observations	3,285
Clusters (days)	365
Pseudo R^2	0.981

Day-clustered standard errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table 5. Cross-stablecoin difference tests in temporal patterns. Pooled Poisson regression with stablecoin × tz_office and stablecoin × weekend interactions, estimated on the Simple Transfer subsample. The interaction coefficients are direct difference tests against the USDC reference: an AMER × stablecoin$_s$ coefficient is the additional AMER tilt that stablecoin s exhibits beyond USDC's. All cross-stablecoin differences are highly significant. Sample, fixed effects, and standard-error treatment as in Table 4.

baseline within each action.

Action	USDC	USDT	PYUSD
Panel A. AMER tilt (% change vs EMEA, weekday)			
Simple Transfer	−3.8%	−12.5%	**+109.6%**
Complex Financial	+14.7%	+11.5%	+15.5%
Exchange	+3.6%	−8.3%	−10.1%
Lending	−4.5%	−14.2%	+2.9%
Multi Transfer	−18.0%	−13.6%	+22.2%
Other	−6.7%	−6.0%	+24.0%
Panel B. APAC tilt (% change vs EMEA, weekday)			
Simple Transfer	−33.5%	−26.8%	+29.3%
Complex Financial	+8.7%	−1.0%	+9.3%
Exchange	−4.8%	−9.9%	−2.9%
Lending	−14.4%	−20.6%	−20.8%
Multi Transfer	−62.0%	+14.7%	+144%†
Other	−17.1%	−14.0%	+21.6%
Panel C. Weekend effect (% change vs weekday)			
Simple Transfer	−24.7%	−18.4%	−9.4%
Complex Financial	−23.2%	−19.6%	−58.8%
Exchange	−12.7%	−13.5%	−4.4%
Lending	−21.1%	−22.4%	−21.2%
Multi Transfer	−34.0%	−3.4%	−49.2%†
Other	−21.6%	−15.7%	−24.9%

Implied from per-stablecoin Poisson regressions with action × tz_office and action × weekend interactions; week-of-year FE; day-clustered SEs.
† Based on only 42 PYUSD Multi Transfer transactions in 2025; noisy.

Table 6. Action heterogeneity in temporal patterns. Implied within-action coefficients from per-stablecoin Poisson regressions with action × tz_office and action × weekend interactions, converted to percentage changes vs the EMEA-weekday baseline within each action. The bolded PYUSD Simple Transfer AMER entry highlights the dramatic business-hour concentration of PYUSD payment activity, which collapses to between −10% and +24% across non-Simple action categories. For USDC and USDT, the business-hour patterns are smaller in magnitude and broadly similar across action categories rather than concentrated in any specific use case.

The action-by-action breakdown reveals two distinct stories. First, PYUSD's strong AMER concentration is overwhelmingly a Simple Transfer phenomenon. Its Simple Transfer AMER coefficient of +110% collapses to between −10% and +24% across every other action category. When PYUSD is used for complex financial activity, exchange interactions, lending, or any non-payment category, the distinctive American-business-hours alignment is absent, and in the case of Exchange transactions PYUSD tilts mildly toward EMEA. The aggregate "PYUSD is an American-hours stablecoin" observation is therefore most precisely a statement about PYUSD payment activity, not about PYUSD activity in general.

Second, the incumbents do not display analogous dilution. For USDC and USDT, the business-hour patterns are smaller in magnitude (mostly within ±20% AMER) and broadly similar across action categories rather than concentrated in any specific use case. USDT's pattern is particularly uniform: every non-Simple action category remains EMEA-tilted, with Exchange at −8.3%, Lending at −14.2%, Multi Transfer at −13.6%, and Other at −6.0%, all reinforcing the EMEA orientation rather than diluting it; only Complex Financial (+11.5%) shows a mild AMER lean. USDC's pattern is more heterogeneous, spanning Multi Transfer (−18.0%) to Complex Financial (+14.7%), but again does not collapse into a single "USDC time profile" that fades for complex activity. For the incumbents, the temporal patterns documented in the descriptive heatmaps extend beyond payment activity and reflect a more general feature of stablecoin use; for PYUSD they are specifically a payment phenomenon.

Weekend effects show only limited action heterogeneity. For USDC and USDT the within-action weekend drops are modest and broadly similar across categories, spanning roughly −3% to −34%, and most weekend-by-action interactions are statistically insignificant. The clearest departures are for PYUSD, whose Complex Financial (−59%) and Multi Transfer (−49%) drops exceed its −9% Simple Transfer baseline; both rest on few observations, the latter on only 42 PYUSD Multi Transfer transactions across the year, and are correspondingly noisy.

6 Conclusion

This paper examines the structure of stablecoin transactions in 2025 using a novel dataset of 593 million Ethereum event logs covering the three well-known U.S. dollar-denominated stablecoins: Tether (USDT), USD Coin (USDC), and PayPal USD (PYUSD). These three account for a substantial share of stablecoin activity and span distinct institutional designs. We formalize the distinction between a transaction and the transfer events it emits, and develop transparent, vendor-independent measures of transactional complexity along several conceptually independent dimensions: token and contract co-usage, financial-action composition, computational burden, urgency, and intraday and intraweek timing.

The framework rests on publicly verifiable inputs: event signatures and protocol membership read from on-chain factory contracts. Third-party labeling is minimal and purely descriptive, and the classifications follow transparent structural definitions rather than proprietary address tagging or ad hoc filters. The framework complements existing vendor- and AI-based approaches with a reproducible lens on on-chain stablecoin activity.

Two main findings emerge. First, a substantial share of stablecoin activity extends beyond simple payments: 31.6 percent of stablecoin transactions involve additional contract interactions or multiple transfers, ranging from basic atomic asset swaps to intricate operations that span dozens of counterparties and smart contract protocols. Second, the three stablecoins are not used interchangeably. We document meaningful differences along several of the dimensions we measure, including the tokens and contracts they co-occur with, the urgency with which their transactions are executed, and the alignment of their activity with regional business hours and weekly calendar patterns. The differences are most pronounced for PYUSD, while USDC and USDT exhibit broadly similar patterns with smaller, yet still meaningful differences between them. These differences are economically substantive and reflect heterogeneity in institutional design, regulatory status, user base, and functional role rather than incidental variation.

A practical implication for empirical work follows directly from these findings, and concerns the lens through which stablecoin data are typically observed. Because stablecoin activity is most commonly accessed through transfer events, the natural empirical starting point is a transfer-level dataset. Yet the perspective is decisive: while 31.6 percent of stablecoin transactions involve additional contract interactions or multiple transfers, almost 60 percent of all transfer events occur within such composite transactions, because complex transactions emit many transfer events each. Many transfers therefore correspond to intermediate accounting steps, internal routing legs of exchange or lending operations, or technical flows with no direct interpretation as economic value transfers between independent counterparties. Treating each transfer as a standalone payment overstates measured activity and transferred volumes, distorts concentration measures, and can yield misleading inferences about the economic role of stablecoins. Using transfer events without additional categorization and complexity analysis is therefore not a neutral choice of unit of observation but a substantive modeling assumption that the evidence here suggests is rarely warranted. A precise characterization of transaction structure, and of the function stablecoins fulfill within those transactions, is a prerequisite for credible empirical research on stablecoin activity.

Several policy implications follow. *First*, the transparent, reproducible metrics developed here are well suited to inclusion in supervisors' and central banks' analytical toolkits alongside vendor data and survey evidence; combining these sources can reduce model risk and improve comparability across studies. *Second*, the action-set classification developed here separates peer-to-peer payments from settlement-like composite transactions and offers supervisors an empirical basis for activity-based regulation, consistent with the CPMI-IOSCO guidance on stablecoins used as settlement assets. *Third*, the distinct business-hour profiles documented across the three stablecoins, with PYUSD concentrated during American business hours and USDC and USDT during European business hours, underscore the inherently cross-jurisdictional nature of on-chain stablecoin activity and the corresponding need for international coordination and data-sharing arrangements among supervisory authorities.

The framework has limitations. Protocol design and market microstructure shape what we observe, and off-chain arrangements such as custodial netting and settlement batching can blur the interpretation of on-chain flows. Our analysis is also limited to the Ethereum mainnet and to a single calendar year. Natural extensions include expanding chain coverage and temporal granularity, triangulating the complexity metrics with independent ground-truth audits, exchange disclosures, and household or firm surveys, integrating geolocation methods to map transactional complexity to regions and corridors, and examining how regulatory and macroeconomic shocks affect the structural patterns identified here.

Nevertheless, the evidence provided in this paper supports a conceptual reframing. Stablecoins are not interchangeable digital representations of conventional payment instruments; they are constitutive elements of a programmable financial platform, structurally differentiated across issuers and embedded in composable smart contract infrastructure in ways that differ materially from traditional payment systems. Interpretations of stablecoin activity that abstract from transaction-level structure, atomicity, and composability risk misrepresenting both the scale and the nature of that activity, with consequences for empirical research and for relevant supervisory and regulatory frameworks. Our methodology offers a transparent, vendor-independent lens that complements existing data sources and supports comparability across studies and jurisdictions.

References

Ahmed, R. and Aldasoro, I. (2025), Stablecoins and safe asset prices, BIS Working Papers 1270, Bank for International Settlements. Revised in February 2026.

Aldasoro, I., Beltran, P. and Grinberg, F. (2026), Stablecoin flows and spillovers to FX markets, BIS Working Papers 1340, Bank for International Settlements.

Aldasoro, I., Cornelli, G., Ferrari Minesso, M., Gambacorta, L. and Habib, M. M. (2024), Stablecoins, money market funds and monetary policy, BIS Working Papers 1219, Bank for International Settlements.

Auer, R., Lewrick, U. and Paulick, J. (2025), DeFiying gravity? An empirical analysis of cross-border Bitcoin, Ether and stablecoin flows, BIS Working Papers 1265, Bank for International Settlements.

Batra, I., Zevin, M., Mathur, A., Quitete, C., Bravo, C. and Li, D. (2026), Stablecoin payments – the truth behind the numbers, BCG White Paper, Boston Consulting Group.

Cerutti, E. M., Chen, J. and Hengge, M. (2024), A primer on bitcoin cross-border flows: measurement and drivers, IMF Working Paper 24/85, International Monetary Fund.

CPMI-IOSCO (2012), 'Principles for financial market infrastructures'. Accessed: 27 May 2026.
URL: *https://www.bis.org/cpmi/publ/d101.htm*

CPMI-IOSCO (2022), Application of the principles for financial market infrastructures to stablecoin arrangements, Technical report, Committee on Payments and Market Infrastructures (CPMI) and the International Organization of Securities Commissions (IOSCO).

Financial Stability Board (2019), Regulatory issues of stablecoins, Technical report, Financial Stability Board.

Financial Stability Board (2023), High-level recommendations for the regulation, supervision and oversight of global stablecoin arrangements: Final report, Technical report, Financial Stability Board.

G7 Working Group on Stablecoins (2019), Investigating the impact of global stablecoins, CPMI Papers No 187, Committee on Payments and Market Infrastructures (CPMI).

Higginson, M., Zorrilla, A., Madden, J. and Kirchner, M. (2026), Stablecoins in payments: What the raw transaction numbers miss, Report, McKinsey & Company.

Kosse, A., Glowka, M., Mattei, I. and Rice, T. (2023), Will the real stablecoin please stand up?, BIS Papers 141, Bank for International Settlements.

Reuter, M. (2025), Decrypting crypto: How to estimate international stablecoin flows, IMF Working Paper 25/141, International Monetary Fund.

Sheffield, C. (2025), 'Making sense of stablecoins: Onchain analytics can help financial institutions cut through complexity', Visa Trends and Insights.

Wood, G. (2014), 'Ethereum: A secure decentralised generalised transaction ledger'. Ethereum Project Yellow Paper. Berlin Version.
URL: *https://ethereum.github.io/yellowpaper/paper.pdf*

7 Appendix

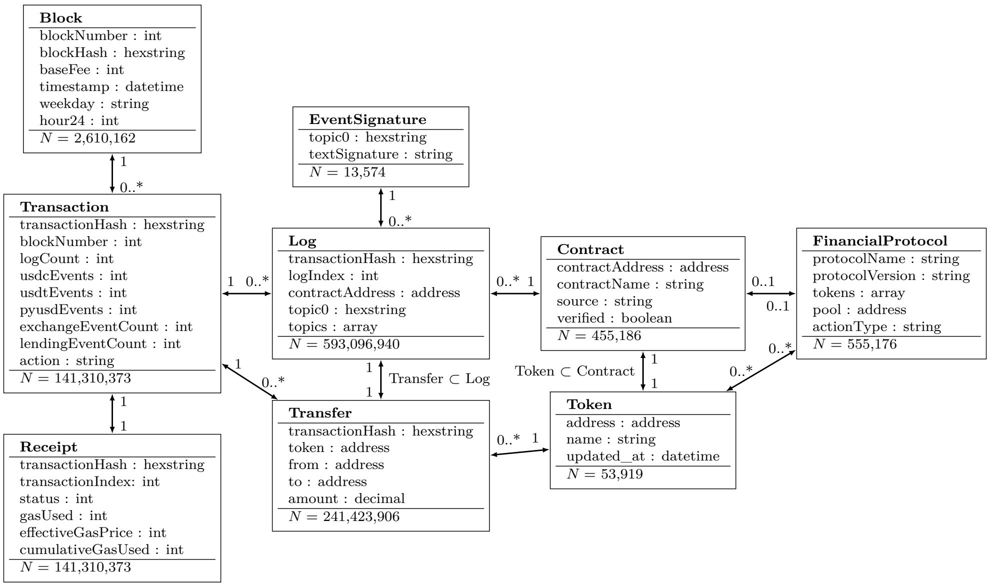

Figure 9. Logical UML-style representation of the `stablecoins` database schema used for this research project. The `Contract` entity comprises the smart contracts observed in the stablecoin transaction sample, whereas `FinancialProtocol` is a comprehensive registry of identified pool deployments across the entire Ethereum contract universe; the optional one-to-one mapping between them therefore captures the intersection, namely contracts that both appear in the stablecoin sample and are recognized as pools of a tracked protocol. For the same reason, the `tokens` array on `FinancialProtocol` can reference token addresses that have no corresponding row in the `Token` entity, which is restricted to ERC-20 contracts observed in the stablecoin sample; the `Token`–`FinancialProtocol` mapping is therefore partial.

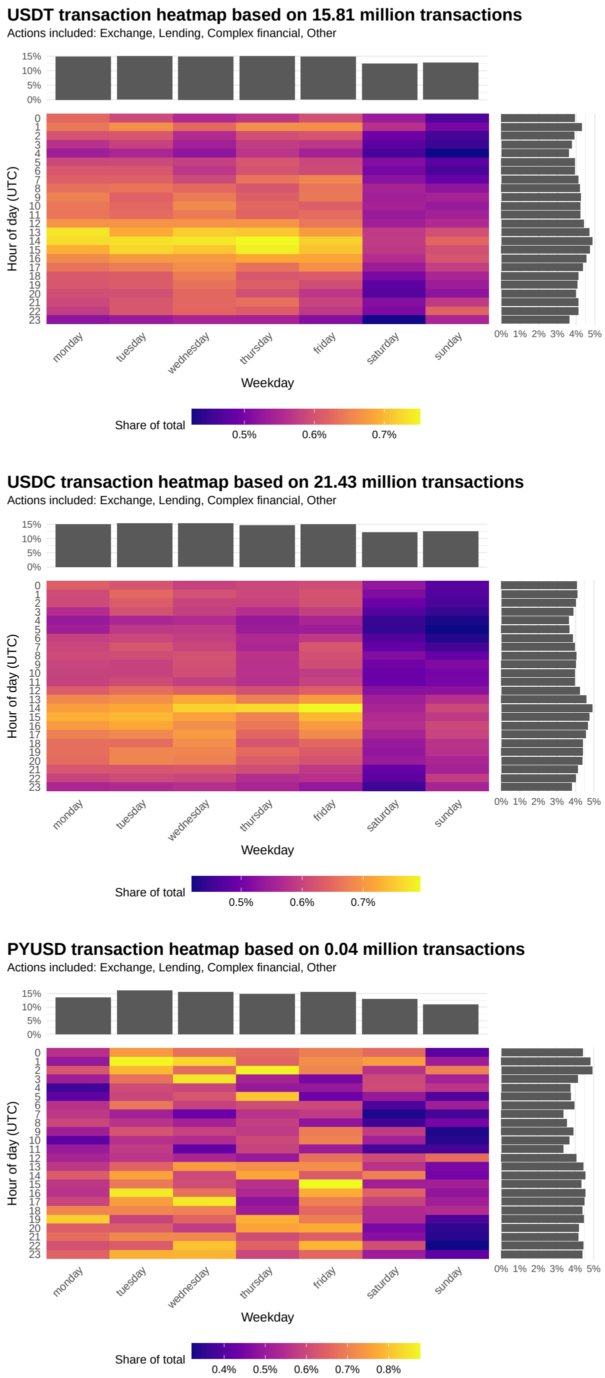

Figure 10. Heatmap of all complex transaction activity by stablecoin. Note more even distribution, when compared to Simple Transfers. Weekend effect still visible.

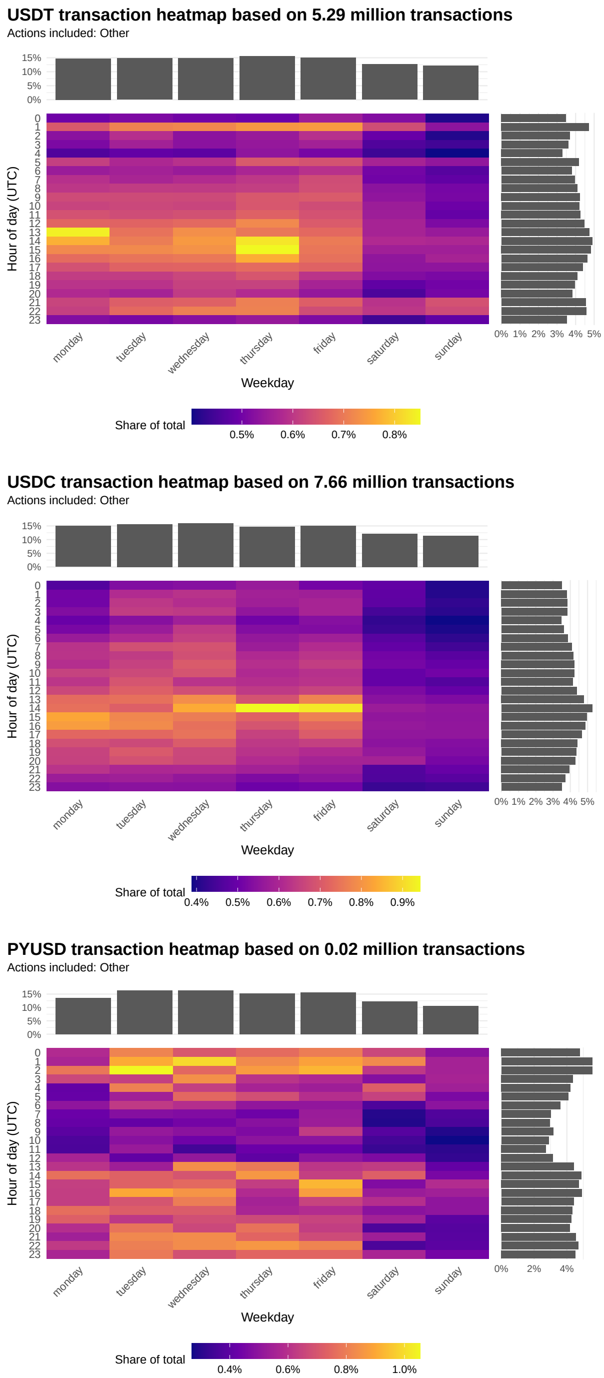

Figure 11. Heatmap of all *Other* transaction activity by stablecoin. Note repeated patterns during specific weekday/time-slots.