Blockchain Currency Markets

Angelo Ranaldo [*], Ganesh Viswanath-Natraj [†], Junxuan Wang [‡§]

This version: November 18, 2025

Abstract

We conduct the first comprehensive study of blockchain currencies—stablecoins pegged to fiat currencies and traded on decentralized exchanges. Using transaction-level data linked to wallet characteristics, we show that prices in these markets are generally efficient, though constrained by blockchain-specific frictions such as gas fees and Ether volatility. Decentralized exchange rates closely track traditional currency markets through arbitrage and informed trading. Traders with significant market share and access to primary markets have greater price impact, reflecting informational advantages. While blockchain markets may improve access for customers excluded from traditional venues, their scalability depends on addressing frictions inherent to decentralized trading.

Keywords: Stablecoins, foreign exchange, blockchain, price efficiency, market resilience, microstructure.
JEL Classifications: D53, E44, F31,G18, G20, G28

[*]angelo.ranaldo@unibas.ch, University of Basel and Swiss Finance Institute

[†]ganesh.viswanath-natraj@wbs.ac.uk, Warwick Business School and Gillmore Centre for Financial Technology

[‡]junxuanwang@hkust-gz.edu.cn, Hong Kong University of Science and Technology (Guangzhou)

[§]For comments, we would like to thank Andrea Barbon, Jean Barthelemy, Dirk Baur, Rhys Bidder, Mathieu Bouvard, Ajit Desai (Discussant), Barry Eichengreen, Thierry Foucault, Cam Harvey, Wenqian Huang, Kim-Sau Chung, Alfred Lehar, Gordon Liao, Rich Lyons, Igor Makarov (Discussant), Magali Marx (Discussant), Benoit Nguyen, Andreas Park, Christine Parlour, Khai Uy Pham, Lorenzo Schoenleber (Discussant), Gustavo Schwenkler (Discussant), Olav Syrstad, Tammaro Terracciano (Discussant), Alexander Wehrli (Discussant), Anthony Lee Zhang, and seminar participants at the Banque de France, the Digital Economy Seminar (HK Baptist, Monash, NTU, and RUC), the Banque de France Digital Currency and Financial System Conference, the Edinburgh Fintech Conference, the EDHEC Seminar Series on the Future of Finance, HKUST Guangzhou, the Imperial College Hedge Fund conference, the Deutsche Bundesbank International Conference on Payments and Securities Settlement, the Lund Cryptocurrency Conference, the SNB-CIF Conference on Cryptoassets and Financial Innovation, the Swiss National Bank, University of Reading, University of Toulouse Sustainable Finance conference, and the University of Western Australia. Angelo Ranaldo acknowledges financial support from the Swiss National Science Foundation (SNSF grant 204721). Ganesh Viswanath-Natraj acknowledges support from the Berkeley Ripple University Blockchain Research Initiative. We would like to thank Edouard Mattille and Peteris Kloks for assistance with the CLS Market Data.

1 Introduction

Decentralized finance (DeFi) refers to financial services built on public blockchains, where transactions are executed automatically through smart contracts rather than intermediated by banks or dealers.[1] These systems create new venues for global foreign exchange (FX) trading by allowing fiat-referenced stablecoins to be exchanged on decentralized exchanges (DEXs). Stablecoins are digital tokens designed to maintain parity with traditional currencies through collateral backing and redemption mechanisms. Examples include EURC, a euro-referenced stablecoin, and USDC, a dollar-referenced stablecoin, both issued by Circle and backed by liquid reserves.[2] Blockchain markets enable these stablecoins to trade on-chain, with pairs such as EURC/USDC tracking conventional FX pairs like EUR/USD. Although small relative to global FX turnover, they have attracted growing attention as potential infrastructure for cross-border payments and settlement.[3]

This paper makes two contributions. First, it provides the first microstructure-based analysis of a blockchain currency market that directly parallels a traditional FX pair. Second, it uses granular blockchain transaction data to test whether order flow, defined as the net signed volume of buyer- and seller-initiated trades, reflects arbitrage activity or conveys private information about fundamentals. We show that participants with large market share and direct access to stablecoin issuance and redemption exhibit informational advantages consistent with the asymmetric information paradigm in FX markets (Ranaldo and Somogyi, 2021).

[1] Smart contracts are self-executing programs that enforce agreed terms automatically on the blockchain and are publicly verifiable by all network participants.

[2] USDC is backed by U.S. dollars and short-term U.S. Treasury bills, while EURC is backed by euro deposits and high-quality euro assets. Both are issued by Circle under a full-reserve model, with reserves held at regulated institutions and verified through monthly attestations. Each token is redeemable 1:1 for the corresponding fiat currency.

[3] See BIS Innovation Hub (2023), "Project Mariana: Cross-border trading and settlement using wholesale CBDC and DeFi infrastructure." https://www.bis.org/about/bisih/topics/cbdc/mariana.htm.

We focus on the EURC/USDC pair on Uniswap V3, the dominant DEX for stablecoin trading and an on-chain analogue to the EUR/USD market. A key advantage of blockchain data is its transparency and granularity. Every transaction can be traced to a unique wallet address, allowing participants to be classified by trading behavior. We distinguish three economically meaningful groups. *Sophisticated traders* are large, high-frequency participants who act as informed arbitrageurs. *Primary dealers* have privileged access to fiat issuance and redemption channels through the stablecoin issuer. *Liquidity providers (LPs)* supply liquidity to the pool but typically play a passive role. Such identification is rarely possible in traditional FX microstructure data, where dealer and client functions must be inferred indirectly (Hortaçsu and Sareen, 2005; Hagströmer and Menkveld, 2019). The transparency of blockchain data therefore provides a novel setting to study how information is incorporated into currency prices in real time.

We start by documenting three stylized facts. First, EURC/USDC prices are closely connected to traditional FX benchmarks, with price deviations from the Continuous Linked Settlement (CLS) EUR/USD rate[4] averaging 24 basis points and trading volumes strongly correlated with interbank activity during core trading hours. Second, these deviations correlate with blockchain-specific frictions such as gas fees and Ether volatility rather than traditional balance sheet constraints. Third, transaction costs vary substantially across traders, with gas fees dominating for smaller accounts and slippage for larger traders.[5] These costs, although higher than those faced by inter-dealer market participants, are comparable to those incurred by less privileged OTC clients (Hau et al., 2021).

[4]CLS is the primary settlement system for wholesale FX transactions, widely used as a benchmark for interbank EUR/USD prices.

[5]Gas fees are transaction costs paid to blockchain validators for confirming and recording trades on the network. They depend primarily on network congestion and are largely fixed per transaction. Slippage refers to the price impact arising when a trade moves the market price due to limited liquidity and increases with trade size.

Building on these facts, we conduct microstructure tests to identify how blockchain trading influences traditional FX prices. We focus on two mechanisms: a *feedback trading* channel, where traders arbitrage price discrepancies between venues, and an *asymmetric information* channel, where order flow conveys private information about fundamentals. The evidence supports both mechanisms. Blockchain order flow frequently corrects price gaps between EURC/USDC and EUR/USD, consistent with active arbitrage by sophisticated traders who, owing to higher trading frequency and larger capital, are better positioned to exploit such opportunities than LPs or primary dealers, whose smaller trade sizes and higher proportional costs constrain their activity. Blockchain markets also process fundamental information. During the USDC de-pegging episode on 11 March 2023, when concerns about USDC reserves at Silicon Valley Bank (SVB) drove its price down to 87 cents, sophisticated traders predominantly bought EURC and sold USDC, reflecting their awareness of underlying reserve risks. In contrast, LPs showed no notable change in swap or mint–burn behavior, consistent with their role as uninformed participants focused on inventory management rather than reacting to information about fundamentals.

To assess whether blockchain order flow predicts traditional FX rates and reveals asymmetric information among participants, we estimate a structural VAR (SVAR) model that captures the dynamic feedback between blockchain order flow and price movements. Our baseline identification follows Hasbrouck (1991), assuming that order flow affects prices contemporaneously while prices affect subsequent trading with a lag. By disaggregating order flow into distinct trader categories, we differentiate the price impact of sophisticated traders, primary dealers, and LPs. We find that sophisticated traders and primary dealers exhibit significant and persistent price impacts, reinforcing their role in incorporating fundamental information into prices, while LPs display an insignificant or weakly negative impact.

To understand whether these permanent price impacts reflect private information rather than mechanical arbitrage, we employ two complementary identification strategies. First, we separate transactions based on routing, distinguishing private transactions that bypass the public mempool from public transactions.[6] Private trades, which are concentrated among sophisticated traders and primary dealers, generate larger and more persistent price impacts, consistent with their containing more information about fundamentals. Second, we decompose order flow into a predicted feedback component, driven by lagged DEX–CLS price differences, and a residual component, which we interpret as informational order flow. Only the residual component produces significant price impact with respect to benchmark FX returns, whereas feedback-based arbitrage does not. This evidence supports the asymmetric information paradigm in blockchain FX markets.

Finally, we conduct a series of robustness tests to verify that our results are not driven by alternative mechanisms. We re-estimate the SVAR model using specifications that control for net liquidity provision, trading frequency and trade size. The estimated price impacts remain close to the baseline in all alternative specifications. We then examine the intraday distribution of price impacts and find that the effects of sophisticated investors and primary dealers are concentrated between 13:00 and 15:00 UTC, overlapping with traditional trading hours and macroeconomic news releases. In contrast, LPs show no systematic impact during this window, consistent with hedging rather than informational trading. Finally, we consider strategic liquidity provision such as just-in-time (JIT) strategies. Although one wallet exhibits this behavior, it is an isolated case rather than a systematic feature of the EURC/USDC market. Overall, these tests confirm that the observed price impacts primarily reflect informed trading rather than alternative factors such as

[6] The mempool is a temporary holding area where pending transactions are stored before being confirmed on-chain. Private transactions are submitted directly to validators and do not appear in the public mempool prior to confirmation.

liquidity provision.

Related Literature. This paper contributes to several strands of research on DEXs, stablecoins, and FX market microstructure.

On the design and functioning of DEXs, Barbon and Ranaldo (2024) assess DEX efficiency relative to traditional markets, Capponi and Jia (2021) examine the incentives for liquidity provision in AMMs, and Lehar and Parlour (2025) analyze theoretically the incentives for informed trading on DEX venues and how these depend on the characteristics of liquidity pools. Malinova and Park (2024) and Foley et al. (2023) study the role of AMMs in FX and equity markets. Using transaction-level blockchain data, Liu et al. (2023) analyze the TerraLuna collapse, while Adams et al. (2023) evaluate DEX trading costs for the EURC/USDC market and compare them to those in traditional OTC trading and cross-border payments. Relative to this literature, we use blockchain data to identify informational advantages across trader types and to show how participants link decentralized and traditional FX markets through feedback trading and the incorporation of fundamental information.

Our work also contributes to the growing literature on stablecoins and their price stability. Prior studies (Lyons and Viswanath-Natraj, 2023; Ma et al., 2025) show that arbitrage stabilizes on-chain exchange rates across tokens and markets, while Gorton et al. (2022); Aldasoro et al. (2023); Eichengreen et al. (2023) highlight the fragility and run risk inherent in collateralized stablecoins. We extend this literature by examining how stablecoins operate within the informational efficiency of blockchain currency markets and by assessing their potential role in foreign exchange trading. In particular, we provide evidence on the price and informational efficiency of blockchain currency markets and their ability to compete with traditional OTC infrastructure.

We also contribute to the literature on FX microstructure and price formation using order flow (Evans and Lyons, 2002; Ranaldo and Somogyi, 2021; Huang et al., 2025). Traditional FX pricing models emphasize portfolio shifts and inventory management (Evans and Lyons, 2002), whereas pricing on Uniswap V3 is determined algorithmically through bonding curves that clear trades without intermediaries. In both settings, sophisticated traders and primary dealers hold informational advantages, while LPs act more passively by rebalancing inventory in ways similar to market makers in limit-order book markets. Our results show that blockchain order flow predicts EUR/USD returns, consistent with the asymmetric information paradigm in FX markets (Ranaldo and Somogyi, 2021) and conceptually related to the microstructure tests of Lyons (1995), which distinguish price adjustments driven by inventory control from those driven by private information. Inventory control plays a limited role in blockchain currency markets because there is no inter-dealer tier through which liquidity providers can share risk. Instead, the feedback trading channel, in which arbitrageurs benchmark to traditional prices, and the private information channel together explain how blockchain currency markets benchmark to traditional FX markets.

The remainder of the paper is structured as follows. Section 2 describes the institutional setting and data. Section 3 analyzes market efficiency and transaction costs in decentralized currency markets, highlighting their connection to traditional FX markets and trader-level frictions. Section 4 examines the relationship between blockchain trading and traditional markets and evaluates the information content of different types of market participants. Section 5 concludes.

2 Data and Institutional Background

2.1 Market Structure of Blockchain and Traditional Currency Markets

Figure 1 illustrates the structural differences between traditional and blockchain-based currency markets, focusing on liquidity provision and price stabilization. In traditional FX markets, the inter-dealer market is central to price discovery. Dealer banks act as market makers, providing liquidity to customers while trading among themselves to manage inventories and facilitate price formation. Corporates, investment funds, and non-bank financial institutions typically access FX liquidity through these dealer banks. Since the early 1990s, electronic trading platforms such as Refinitiv and EBS have supported this structure by allowing dealers to post bid and ask quotes on centralized limit order books (see King et al., 2012; Chaboud et al., 2023). Dealer banks contribute to FX price formation through the information embedded in customer and inter-dealer order flow, consistent with the inventory and portfolio-shift models of Evans and Lyons (2002) and related work (Bjønnes and Rime, 2005; Ranaldo and Somogyi, 2021; Huang et al., 2023).

By contrast, blockchain-based currency markets operate in a decentralized structure in which secondary market trading is disintermediated, while primary issuance still involves a stablecoin intermediary. In the primary market, the stablecoin treasury managed by Circle mints EURC and USDC and distributes tokens to users who interact directly with the issuer, referred to as *primary dealers*. In the secondary market, these tokens circulate across a range of applications, including trading on decentralized (DEX) and centralized (CEX) exchanges, depositing in lending and liquidity protocols, or facilitating cross-border payments (Adams et al., 2023). On DEXs, trades are executed through *automated market makers (AMMs)*—algorithmic smart contracts that determine prices and execute transactions according to pre-specified mathematical rules. Market participants include LPs, arbitrageurs, and algorithmic traders, each playing distinct roles in price discovery and liquidity formation.

[INSERT FIGURE 1 and TABLE 1 ABOUT HERE]

Primary dealers play a key role in arbitraging price deviations between primary and secondary markets. A defining feature of EURC and USDC is the issuer's commitment to maintain the peg at par (1 EURC = 1 EUR, 1 USDC = 1 USD). This peg is enforced through arbitrage. When USDC trades above one dollar, primary dealers can deposit USD with Circle to mint new USDC and sell it at a premium, expanding supply and restoring parity. When USDC trades below one dollar, dealers buy the token at a discount and redeem it for USD, reducing supply and pushing the price back toward par. Further details on issuance and arbitrage are provided in Appendix A.

While both market designs facilitate currency exchange and liquidity provision, Table 1 summarizes the main institutional differences between traditional FX markets and blockchain-based systems such as Uniswap. Algorithmic market making, atomic settlement, and wallet-level transparency distinguish decentralized markets from dealer-based systems. For instance, traditional FX settlement typically occurs two business days after the trade date ($T + 2$), whereas blockchain transactions settle *atomically*, meaning that execution and settlement occur within the same block. This usually takes seconds rather than days and therefore reduces counterparty and settlement risk in principle.[7]

[7]Under the $T + 2$ convention, counterparties exchange currencies two business days after the trade date, allowing time for payment processing and liquidity coordination. Settlement risk is mitigated through the CLS system, which provides payment-versus-payment (PvP) settlement via a centralized intermediary. In contrast, atomic settlement on blockchains embeds PvP directly in the transaction layer, so both legs of a trade either execute together or fail together within a single block. On the Ethereum network, a new block is confirmed roughly every 12–15 seconds, implying near-instant finality relative to traditional settlement cycles. While this removes timing risk, it introduces potential smart-contract and operational risks.

These institutional differences imply that blockchain-based markets offer algorithmic transparency and continuous pricing, whereas traditional FX markets rely on discretionary intermediation and balance-sheet capacity. We now turn to the price-setting mechanics of Uniswap to formalize how decentralized liquidity and automated pricing function in practice.

2.2 Uniswap V2 and the Bonding Curve

Uniswap is a decentralized AMM that replaces centralized order books with on-chain liquidity pools managed by smart contracts. The first version (V1) launched in 2018, followed by V2 in 2020.[8] In Uniswap V2, each pool holds reserves of two tokens, EURC and USDC, denoted by x and y, representing the base and quote tokens respectively. The pool satisfies the constant-product invariant

$$k = x \times y, \tag{1}$$

which must hold before and after each trade. The marginal price of the base token in terms of the quote token is

$$P_{x/y} = \frac{y}{x}. \tag{2}$$

Because both reserves and transactions are public on-chain, price discovery is algorithmic and liquidity is continuously available.

Swaps and liquidity provision. Panel (a) of Figure 2 illustrates how swaps and liquidity mints operate along and across the bonding curve defined by (1). A swap moves reserves along the curve. If a trader purchases $\Delta x > 0$, meaning a buy of the base token (EURC) in exchange for USDC, the

[8] See Adams et al. (2023) and the Uniswap V2 white paper (`https://uniswap.org/whitepaper-v2.pdf`).

post-trade reserves and price satisfy

$$x' = x - \Delta x, \quad y' = \frac{k}{x'}, \quad P'_{x/y} = \frac{k}{(x - \Delta x)^2}. \tag{3}$$

This preserves k and implies that, because $x' < x$, the marginal price increases with each purchase, generating slippage along the curve. Liquidity mints differ in that they shift the entire bonding curve outward by scaling up k. To keep $P_{x/y}$ constant, LPs add both tokens in proportion to the prevailing price so that y/x remains unchanged. Burns reverse this process, shifting the curve inward.

Numerical example. Panel (b) of Figure 2 illustrates how price impact depends on liquidity. Each pool starts at $P_{x/y} = 1.10$ but differs in scale: $k_{\text{low}} = 11{,}000$, $k_{\text{medium}} = 275{,}000$, $k_{\text{high}} = 1.1 \times 10^6$. For the low-liquidity pool with $(x, y) = (100, 110)$, a purchase of $\Delta x = 5$ gives $x' = 95$, $y' = 11{,}000/95 \approx 115.79$, $P'_{x/y} = 115.79/95 \approx 1.219$. For the medium- and high-liquidity pools, the same trade results in $P'_{x/y} \approx 1.122$ and 1.111, respectively.[9]

Price impact is therefore largest in smaller pools. Slippage, defined as the percentage deviation between post-trade and pre-trade prices, can be expressed as

$$\text{Slippage} = 100 \times \left[\frac{1}{\left(1 - \frac{\Delta x}{x}\right)^2} - 1 \right]. \tag{4}$$

Price impact depends on trade size relative to pool depth, with larger pools exhibiting flatter slippage curves and smaller price responses for a given trade.

[9] For the medium pool ($k_{\text{medium}} = 275{,}000$), initial reserves are $(x, y) = (500, 550)$. After a 5 EURC purchase, $x' = 495$, $y' = 275{,}000/495 \approx 555.56$, $P'_{x/y} = 555.56/495 \approx 1.122$. For the high pool ($k_{\text{high}} = 1.1 \times 10^6$), reserves are $(1000, 1100)$, giving $x' = 995$, $y' = 1.1 \times 10^6/995 \approx 1105.53$, $P'_{x/y} = 1105.53/995 \approx 1.111$.

[INSERT FIGURE 2 ABOUT HERE]

2.3 Uniswap V3 Concentrated Liquidity and Tick-Based Pricing

Uniswap V3 generalizes the constant-product rule of V2 by allowing liquidity to be concentrated within a user-defined price range rather than distributed uniformly. In V2, LPs deposited equal values of two tokens across the entire price range $[0, \infty]$, satisfying $x \times y = k$. This allocation is capital-inefficient, as liquidity placed far from the current market price remains idle and does not contribute to meeting trading demand.

Uniswap V3 improves efficiency through two innovations. LPs can allocate liquidity within a chosen price interval $[p_a, p_b]$, making capital active only when the market price lies inside this range. It also introduces multiple fee tiers (0.01%, 0.05%, 0.3%, 1%) that segment liquidity across pools, allowing LPs to tailor risk and return preferences.[10]

Within any active range $[p_a, p_b]$, Uniswap V3 prices trades using *virtual reserves* that satisfy

$$\left(x + \frac{L}{\sqrt{p_b}}\right)\left(y + L\sqrt{p_a}\right) = L^2, \tag{5}$$

where L is the *liquidity parameter*. A higher L implies deeper liquidity and smaller price impact. When the range expands to $[0, \infty)$, the terms $L/\sqrt{p_b}$ and $L\sqrt{p_a}$ vanish, and (5) simplifies to $x \times y = L^2$. Hence L is the V3 analogue of $\sqrt{k}$ and governs price responsiveness near the market price. Further details on liquidity aggregation and price setting are provided in Appendix B.1 and B.2.

[10] For example, the USDC/ETH pair trades in both 0.05% and 0.3% fee tiers. Professional LPs often supply liquidity to low-fee pools to capture higher turnover, while smaller LPs prefer higher-fee pools requiring less rebalancing.

Tick-based pricing. Uniswap V3 discretizes prices into ticks indexed by $i \in \mathbb{Z}$, where $p_i = 1.0001^i$. Each tick represents about one basis point. The pool's tick spacing determines which ticks can host liquidity. For instance, the EURC/USDC 0.05% pool has tick spacing 10, so only every tenth tick (e.g., −20, −10, 0, 10, 20) can be initialized. Each position corresponds to one or more of these 10-tick intervals defining a range $[p_i, p_{i+10}]$ within which liquidity is active.

Liquidity snapshot. Figure 3 shows the tick-level liquidity distribution for the EURC/USDC 0.05% pool at block 19,771,559 (April 30, 2024). The horizontal axis measures distance from the current tick, normalized so that tick 0 corresponds to the market price p_M. Because tick spacing is 10, each step represents one 10-tick interval. Liquidity is expressed in thousands of USDC for comparability. Ticks left of zero correspond to prices below p_M (buy limit orders for EURC); ticks right of zero correspond to prices above p_M (sell limit orders). This tick-based structure resembles a tokenized limit order book, with liquidity concentrated near the mid-price and tapering toward the tails.

[INSERT FIGURE 3 ABOUT HERE]

Liquidity measurement. Following Klein et al. (2024), we construct on-chain measures of net liquidity from Uniswap V3 "mint" and "burn" events, which record additions and withdrawals of liquidity over price intervals $[P_a, P_b]$ at the block level. Mints represent new liquidity added to the pool, and burns represent liquidity removed. For each block k:

$$mint^{net}_{(k)} = mint^{ask}_{(k)} - mint^{bid}_{(k)}, \quad burn^{net}_{(k)} = burn^{ask}_{(k)} - burn^{bid}_{(k)}, \quad Liquidity^{net}_{(k)} = mint^{net}_{(k)} - burn^{net}_{(k)}.$$

A positive $Liquidity^{net}_{(k)}$ indicates that more ask-side (EURC-side) liquidity has been added than withdrawn in block k. To distinguish between active and passive liquidity, we classify positions within 100 basis points of the mid-price as *best*, and those beyond as *away*. The block-level $Liquidity^{net}_{(k)}$ is aggregated to hourly frequency to align with measures of order flow and price changes. Further details, including disaggregation into *best* and *away* ranges, are provided in Appendix B.3.

2.4 Data

2.4.1 CLS EUR/USD Benchmark and Uniswap EURC/USDC Price

We source a benchmark EUR/USD rate from CLS. This provides a volume-weighted average price of interbank quotes at five-minute intervals, which we aggregate to hourly and daily frequencies for our analysis. The data on EURC/USDC is constructed as the last price (both hourly and daily UTC time) using the history of DEX transactions collected from the Uniswap V3 EURC/USDC pool, which is obtained from the Subgraph API.[11]

Our CLS benchmark rate provides an effective benchmark for the EURC/USDC rate from the Uniswap V3 pool. Panel (a) of Figure 4 plots EURC/USDC and EUR/USD prices, as well as the price difference between the EURC/USDC and EUR/USD price. Consistent with Adams et al. (2023), the EURC/USDC market tracks the traditional market and the average (absolute) deviation is 24 basis points. There is more volatility during the early period, which corresponds to low liquidity in the EURC/USDC pool. For this reason, our empirical analysis in Sections 3 and 4 begins on August 15, 2022. Another significant event is the de-pegging of USDC which occurred

[11] API available at https://thegraph.com/hosted-service/subgraph/uniswap/uniswap-v3

in March 2023. This event led to USDC trading at a discount due to concerns on the backing of USDC reserves that were held with Silicon Valley Bank. EURC/USDC traded at a relative premium compared to EUR/USD rates during the days of March 11-12 2023.

[INSERT FIGURE 4 ABOUT HERE]

2.4.2 DEX Trading Volume and Liquidity Provision

The Uniswap V3 dataset contains the complete history of swap transactions, which represent all trades involving the purchase of EURC against the sale of USDC or vice versa. These transactions are recorded at the wallet level, where each wallet corresponds to an Ethereum address that securely holds and manages the tokens associated with that address.[12]

We complement this with data from Kaiko, a cryptocurrency market data provider offering regulatory-compliant, institutional-grade information. This dataset records all liquidity transactions executed by LPs, including the amounts of EURC and USDC added to or removed from the pool and the price range over which liquidity is allocated.[13] A key feature of our analysis is the use of blockchain-level granularity to examine the heterogeneity of market participants. We classify addresses into three main categories. The first group consists of traders with significant market share based on cumulative trading volume. The second group includes traders who also act as LPs. The third group contains addresses that transact directly with the stablecoin Treasury.

Sophisticated Traders. In each month, we aggregate trading volume by wallet and select those that feature in the top 10. The share of top 10 addresses, including any intersection with other

[12] Technically, a wallet stores the private keys required to access and authorize transactions from a specific Ethereum address.

[13] For example, if the current market price of EURC is 1.10 USDC, an LP may supply only EURC if the price range lies above 1.10, only USDC if it lies below 1.10, or both tokens if the range includes the current price. The precise token amounts are determined by the Uniswap V3 AMM pricing algorithm.

categories, averages 52 percent of aggregate trading volume over our sample from 15 August 2022 to 30 April 2024. In the empirical analysis, the Top 10 category consists only of addresses that do not overlap with primary dealers or LPs, ensuring mutually exclusive classifications.

Primary Dealers. Primary dealers are identified as wallets that have transacted with either the EURC or USDC Treasury.[14] We cross-reference all DEX participants with the complete set of addresses that have traded with these Treasuries. Primary dealers interact with the Treasury by exchanging fiat currency for stablecoins at the pegged rate of one to one or by redeeming stablecoins for fiat through token returns. Including any intersection with other categories, primary dealers account for 7 percent of aggregate trading volume. In all regressions, primary dealers are defined as a separate, non-overlapping group.

Liquidity Providers (LPs). LPs are wallets that both trade and provide liquidity, depositing or withdrawing EURC and USDC from the pool. LPs, including any intersections with other groups, account for 7 percent of aggregate trading volume. Like the other classifications, the LP group in our analysis is defined to be mutually exclusive.

Table 2 presents summary statistics on the number of transactions and volume per transaction for seven trader groups. The three primary groups consist of sophisticated traders, primary dealers, and LPs, with 76, 68, and 90 unique addresses identified for each group respectively. To ensure clarity, categories are defined to be mutually exclusive. Top 10 includes only traders ranked by trading volume who are not classified as primary dealers or LPs, while the other two groups are likewise distinct.

[14] The USDC Treasury address used to retrieve transaction history is `0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48`, and the EURC Treasury address is `0x1abaea1f7c830bd89acc67ec4af516284b1bc33c`.

Beyond these main categories, we identify small intersections across groups. Six traders are both sophisticated traders and primary dealers, seven are both sophisticated traders and LPs, and three are both primary dealers and LPs.[15] The majority of addresses, 2,342 in total, fall into a residual category not classified as any of the three groups. Transaction frequency varies considerably across categories. Sophisticated traders average 58 transactions per address, while those classified as both sophisticated traders and primary dealers average 89 transactions per address.

[INSERT TABLE 2 ABOUT HERE]

Blockchain Order Flow. We construct signed measures of blockchain order flow to capture the net direction of trading. Each swap in the EURC/USDC pool records amounts in the base and quote currencies, denoted `amount0` and `amount1`, extracted from the Ethereum blockchain API. These are reported from the pool's perspective. Since EURC is the base currency, a negative value indicates that a trader buys EURC and removes it from the pool, which we classify as a buyer-initiated trade. A positive value indicates a sale of EURC to the pool, classified as a seller-initiated trade.

The blockchain order flow for each interval t is defined as the net buyer-initiated volume,

$$OF_t = \sum_{k \in \mathcal{K}(t)} (\mathbb{1}[T_k = B] - \mathbb{1}[T_k = S]) \times V_k,$$

where $\mathcal{K}(t)$ indexes all swaps in interval t, T_k indicates buyer or seller initiation for EURC, and V_k denotes the EURC-equivalent trade size.

Panel (b) of Figure 4 plots cumulative blockchain order flow alongside prices. There is positive

[15]This latter group, which recorded only six transactions, is excluded from our heterogeneous trading analysis.

co-movement between order flow and the EURC/USDC price. The left plot combines aggregate order flow and liquidity provision, while the right plot disaggregates order flow by trader group, including sophisticated traders, primary dealers, LPs, intersection groups, and the residual category.

Subdividing order flow in this way reveals a clear pattern. The Top 10 group typically takes the opposite side of trades relative to LPs and residual traders. When the cumulative order flow of sophisticated traders is positive, LPs and residual traders are net sellers, and vice versa. While aggregated, this pattern is consistent with an asymmetric information interpretation in which sophisticated traders act as informed participants who exploit short-term mispricing, while LPs and residual wallets provide liquidity on the other side. We test this channel formally in Section 4.

Turning to LPs, their behavior differs from that of informed traders. LPs consistently accumulate EURC when prices rise, suggesting that they do not actively rebalance to remain inventory neutral, unlike traditional FX dealers who offset order-flow shocks (Lyons, 1995; Bjønnes and Rime, 2005). The token-level liquidity series in Panel (b) support this pattern: as EURC appreciates, LPs reallocate liquidity toward EURC positions. This shift coincides with EURC-demanding LP order flow, consistent with rebalancing toward the appreciating token.

Appendix C provides additional detail on the transaction data and participant characteristics used in our analysis. Appendix C.1 introduces key variables and presents example swap and mint transactions from the EURC/USDC pool, showing how execution prices, pool prices, and slippage are computed from on-chain data. Appendix C.2 reports statistics on trading activity and wallet characteristics, including the number of active addresses, trading volumes, and the share of volume accounted for by sophisticated traders. It also examines the persistence of wallet classifications over time and the identification of private transactions using Blocknative mempool archives, which record transactions submitted directly to validators without appearing in the public mempool. Appendix C.3 focuses on liquidity provision, reporting the number of LPs, aggregate liquidity, concentration among the largest providers, intra-day patterns, and evidence of inventory management behavior by large LPs. Appendix C.3.3 further documents these behaviors at the transaction level, showing that large LPs respond to EURC imbalances by withdrawing USDC, purchasing EURC through swaps, and redepositing liquidity to restore balance around the current price. Their role in absorbing inventory and responding to stress events such as the USDC de-pegging is analyzed in Section 4.

2.4.3 Additional Data

CLS Volume. To study the transaction volumes in the traditional currency market, we utilize the CLS FX dataset. CLS Group handles around 40% of global FX transaction volume, including spot, swap, and forward transactions, for up to 18 currencies.[16] CLS data provides aggregated spot FX volume at an hourly frequency, and has been used in a number of papers analyzing the microstructure of the FX spot and swap markets (Ranaldo and Somogyi, 2021; Kloks et al., 2023; Huang et al., 2023). We focus on the CLS FX Spot Flow dataset to construct sector-level volumes and order flows. This dataset records transaction buy and sell volumes between price-takers and market-makers (banks), with price-takers further divided into three categories: funds, non-bank financials, and corporates.

Consequently, we use the Flow dataset to construct sector-level volumes, which include: (i) interbank, (ii) bank-funds, (iii) bank-non-bank financials, and (iv) bank-corporates. To measure interbank volume, we take the aggregate volume from the Flow dataset and subtract the bilateral volume involving banks and other participants, such as funds, non-bank financial institutions, and

[16] The 18 currencies are AUD, CAD, DKK, EUR, HKD, HUF, ILS, JPY, MXN, NZD, NOK, SGD, ZAR, KRW, SEK, CHF, GBP, and USD. In total, 33 currency pairs are settled by CLS.

corporates.

Figure 5 plots hourly trading volume for decentralized and traditional FX markets. Panel (a) reports trading on Uniswap V3 in the EURC/USDC market, while Panel (b) shows CLS volumes for the EUR/USD market by sector. Traditional trading is concentrated between 13:00 and 16:00 UTC, driven mainly by interbank and fund-bank flows. This period coincides with overlapping trading hours in London, Frankfurt, and New York and culminates at the 16:00 UTC WMR fix, a key benchmark for institutional order execution (Krohn et al., 2024).

By contrast, blockchain trading occurs throughout the day, with activity peaks both in the afternoon and around 09:00 UTC. This reflects the 24/7 nature of decentralized markets, where trading can be driven by retail participation and occurs independently of macroeconomic news releases or benchmark fixing windows.[17] Despite this different timing structure, the distribution of blockchain trading activity resembles that of the inter-dealer spot FX segment, indicating integration in how price signals are processed across venues. In scale, average daily CLS volume is 28.42 billion EUR compared to 0.423 million EURC on Uniswap, equivalent to roughly 0.0015% (0.15 bps) of total EUR/USD trading.[18]

[INSERT FIGURE 5 ABOUT HERE]

Gas Fees. Gas fees are payments to Ethereum validators for confirming transactions. We use two sources. First, actual gas fees (in ETH, converted to USDC) for each swap transaction are retrieved from Etherscan to compute arbitrage bounds and trader-level transaction costs. Second, we use the CoinMetrics daily index of average gas fees (USD) to examine price efficiency at the

[17] Appendix C.3.2 documents intraday patterns in liquidity provision, showing fewer mints and burns during peak trading hours but no systematic volume pattern.

[18] See Table 3 for summary statistics of blockchain and CLS volumes.

daily frequency (https://coinmetrics.io).

Market Volatility. We use the EthVol index from T3 Index (https://t3index.com/) as a measure of expected 30-day implied volatility for Ether. This model-free index interpolates across option expirations to capture forward-looking volatility based on investor expectations.

Intermediary Constraints. We include two measures of dealer constraints. The first is the *intermediary capital risk factor* from He et al. (2017), defined as AR(1) innovations to the market-based capital ratio of U.S. primary dealers. The second follows Huang et al. (2025) and measures violations of the law of one price (VLOOP) in G10 FX markets, constructed from minute-level LSEG quotes on EUR/USD, USD/X, and EUR/X cross rates. The first principal component (PC1) of standardized VLOOP series explains about 46% of total variation and captures global inter-dealer pricing distortions. Summary statistics of trading volume, prices, and blockchain variables are reported in Table 3.

[INSERT TABLE 3 ABOUT HERE]

3 Stylized Facts on Blockchain Prices, Volumes, and Costs

3.1 Price and Volume Connection

Fact #1: *DEX Prices and Volumes are Closely Connected to Traditional FX Markets*

We first examine how blockchain markets connect to traditional FX benchmarks through both prices and trading volumes. Our baseline measure of market efficiency is the absolute price deviation between EURC/USDC and the EUR/USD benchmark, defined as

$$\Delta_t = \left| p_{\text{EURC/USDC},t} - p_{\text{EUR/USD},t} \right| . \tag{6}$$

Across the sample, Δ_t averages 24 basis points with a median of 16 basis points, rising above 200 basis points during the March 2023 USDC de-pegging episode. These descriptive statistics show that while the DEX rate broadly tracks the traditional benchmark, short-lived but substantial dislocations can occur under stress.

Turning to volumes, blockchain activity is systematically linked to traditional EUR/USD markets. DEX and CLS volumes display similar intra-day patterns, peaking during 13:00 to 16:00 UTC when European and U.S. markets overlap. Using CLS data disaggregated by sector, including interbank, non-bank financial, and corporate participants, we examine how DEX volumes by participant type—sophisticated traders, primary dealers, and LPs—align with traditional market activity. The literature finds that interbank volumes are particularly responsive to fundamental news through both public and private information (Ranaldo and Somogyi, 2021; Huang et al., 2023). Consistent with this, our analysis shows that DEX trading patterns, especially among sophisticated traders and primary dealers, closely follow those of the informed interbank segment. Appendix D.1 reports the full set of volume correlations across participant groups and trading hours, confirming that correlations are strongest with the interbank segment and during peak trading hours when New York, London, and Frankfurt are open.

Formally, we estimate the relationship between blockchain and traditional FX volumes at the hourly frequency,

$$V_{N_{DEX},t} = \alpha + \sum_{i \in N_{CLS}} \beta_i V_{N_{CLS},t} + \varepsilon_t, \tag{7}$$

where $V_{N_{DEX},t}$ represents DEX trading volumes by participant type and $V_{N_{CLS},t}$ denotes sector-level CLS volumes.

Table 4 reports a strong and statistically significant correlation between blockchain and traditional activity, particularly with interbank volume. Column (1) indicates that a one million euro increase in interbank trading volume corresponds to roughly a 4.35 euro increase in DEX activity for sophisticated traders, with similar magnitudes across trader types. These results suggest that informed interbank activity correlates with blockchain trading behavior across participants, reinforcing the link between decentralized and traditional FX markets.

Panel (a) of Figure 6 illustrates that average DEX volumes peak during 13:00 to 16:00 UTC for all groups, declining by about 50 percent for sophisticated traders and 37 percent for primary dealers outside these hours. For traders classified as both sophisticated and primary dealers, the drop reaches 74 percent. Panel (b) compares weekday and weekend activity, showing that volumes fall sharply on weekends for all groups, with the steepest decline of 87 percent among sophisticated primary dealers. These patterns confirm that DEX activity is most intense during core trading hours and that blockchain markets move closely with traditional FX markets in both price and quantity dimensions.

[INSERT TABLE 4 and FIGURE 6 ABOUT HERE]

3.2 Price Efficiency and Blockchain Frictions

Fact #2: *Peg Efficiency is Driven by Blockchain-Specific Limits to Arbitrage*

In an efficient market, the blockchain price should track the traditional benchmark, and persistent deviations from parity therefore indicate frictions. To identify the main drivers of these inefficiencies, we regress Δ_t on blockchain-specific costs, stablecoin peg deviations, and intermediary constraints at the daily frequency,

$$\Delta_t = \beta_0 + \beta_1 \text{gasfee}_t + \beta_2 \sigma^{IV}_{ETH,t} + \beta_3 R_{ETH,t} + \beta_4 \left| p_{\text{USDC/USD},t} - 1 \right| + \beta_5 \left| p_{\text{EURC/EUR},t} - 1 \right| + \beta_6 \text{VLOOP}_t + \beta_7 \text{ICRF}_t + \varepsilon_t. \quad (8)$$

The results in Table 5 show that blockchain-native frictions are the dominant determinants of price deviations. Gas fees and Ether volatility have the most consistent explanatory power, with a one dollar increase in gas fees raising Δ_t by about 1.3 basis points, and a one basis point increase in implied Ether volatility adding 0.13 to 0.16 basis points. These findings indicate that execution costs and on-chain congestion constrain arbitrage efficiency (Barbon and Ranaldo, 2024; Foley et al., 2023).

Volatility matters not because of dealer balance sheet constraints but because it raises uncertainty for traders holding portfolios in cryptoassets, reducing their willingness to deploy capital for arbitrage. Higher volatility therefore allows deviations from parity to persist even in the absence of funding constraints.

Stablecoin fundamentals also play an important role. Deviations in the USDC/USD peg are highly significant and robust, with a one basis point increase associated with a 0.75 basis point rise in EURC/USDC deviations. In contrast, the EURC/EUR peg deviation is positive but statistically insignificant, consistent with tighter control of issuance and redemption for EURC.

By comparison, traditional measures of intermediary constraints such as VLOOP and the intermediary capital risk factor (ICRF) are statistically insignificant. This suggests that balance sheet frictions binding intermediaries in dealer-based FX markets do not carry over to decentralized venues. Instead, inefficiencies in blockchain FX markets arise mainly from blockchain-specific limits to arbitrage, including transaction costs, volatility, and stablecoin issuance frictions.

[INSERT TABLE 5 ABOUT HERE]

3.3 Trading Costs

Fact #3: *Transaction Costs Vary Widely by Trader Type; Gas Fees Dominate for Most, while Slippage Matter for Large Traders*

We now examine the transaction costs faced by individual traders in the EURC/USDC market. These costs are constructed at the transaction level and expressed in basis points relative to trade size. Each trade incurs four components—gas fees, LP fees, slippage, and private fees. Gas fees are payments to validators for processing transactions on the public blockchain and are largely fixed per trade, declining proportionally with trade size. LP fees are fixed percentage charges specified by the liquidity pool. Slippage reflects price impact and increases with trade size as trades consume liquidity along the bonding curve. Private fees represent payments to validators for authenticating privately routed transactions and are approximated using transfers to validator addresses, as detailed in Appendix C.2.4. Across trader types, median total costs range between 20 and 50 basis points.

Total transaction cost is calculated as the sum of (i) gas fees paid in ETH and converted to USD at execution time, (ii) the fixed LP fee of 5 basis points applied to the Uniswap V3 EURC/USDC pool, (iii) slippage, defined as the difference between the execution price and the prevailing pool price immediately before the trade, and (iv) private fees estimated from validator transfers for privately routed transactions.

Figure 7 presents a disaggregated breakdown of these costs by trader group. Panel (a) shows the interquartile range of total costs across seven account types, including Top 10 wallets, primary dealers, LPs, and a residual group of less active or retail-like traders. The distribution reveals substantial heterogeneity. Sophisticated wallets such as Top 10 and PMs tend to incur lower and more compressed transaction costs, while others face greater dispersion.

Panel (b) decomposes the median transaction cost by component. LP fees remain constant across trades, while gas fees account for the largest share of total costs for most participants. Private fees, which apply only to privately routed transactions, add around 16 basis points for Top 10 traders and 5.6 basis points for Top 10∩LP addresses. Slippage becomes more important for large traders, reaching 21 basis points for Top 10∩LPs, as bigger orders move prices further along the convex bonding curve even as these traders benefit from infrastructure that lowers their per-unit costs. Median total transaction costs range from about 14 basis points for LP∩PM addresses to just above 50 basis points for Top 10 wallets.

[INSERT FIGURE 7 ABOUT HERE]

Trading cost estimates across groups have two main implications. First, they allow us to test the role of arbitrage bounds and assess price efficiency once all sources of transaction costs—including gas, private, liquidity provider, and slippage fees—are considered. Appendix D constructs triangular arbitrage metrics to benchmark efficiency against these cost bounds. After incorporating DeFi transaction costs, the share of arbitrage bound violations is around 16–19%. When centralized exchange fees and OTC bid-ask spreads are included, this falls to 3–5%. These results indicate that the market is best described as constrained price efficient rather than fully efficient.

Second, trading costs shed light on whether decentralized FX markets can serve as a viable alternative to traditional OTC markets. Compared with inter-dealer markets, the contrast remains sharp: EUR/USD spreads averaged only 0.55 basis points in 2023 (Filippou et al., 2024), far below typical on-chain costs. Yet not all OTC participants access such favorable pricing. Hau et al. (2021) show that clients at the 90th percentile of FX derivatives trading face spreads of up to 50 basis points, comparable to median DEX costs. This suggests that decentralized platforms may offer value to participants excluded from preferred dealer pricing tiers. However, limited liquidity and depth still constrain scalability, with slippage particularly high for large trades. In sum, while scalability remains a limitation, DEXs may improve access and fairness for smaller users, offering an alternative venue to traditional OTC FX markets.

4 Informational Efficiency

Building on the stylized facts documented in Section 3, we now turn to the informational efficiency of blockchain currency markets. We consider two channels through which information is incorporated into prices. The first is a feedback trading channel, where traders respond to benchmark prices and public information. The second is an asymmetric information channel, where some traders act on private signals about fundamentals. This distinction clarifies how observed trading patterns reflect the extent to which blockchain prices incorporate both public and private information.

We first consider the feedback trading channel, where blockchain order flow responds to publicly observed price benchmarks and macroeconomic news. Traders monitor EUR/USD benchmark prices in the traditional market and adjust when the DEX reference rate (EURC/USDC) deviates from them. For instance, if the DEX trades at a premium, arbitrageurs sell EURC and buy USDC, bringing the on-chain price closer to parity. Traders also react to macroeconomic news, such as monetary policy announcements, which shift the benchmark in real time. These dynamics reflect the incorporation of public information into blockchain prices. Sophisticated traders, who can scale arbitrage trades more effectively and face lower per-unit transaction costs, are expected to rely more on feedback trading than primary dealers and LPs.

Hypothesis H1 (Feedback Trading)**.** *Blockchain order flow on DEX responds to deviations between the DEX price and the EUR/USD benchmark, as well as to macroeconomic news that shifts the benchmark. The response is stronger for sophisticated traders, who can scale larger arbitrage trades at lower per-unit costs, than for primary dealers and LPs.*

A second channel is asymmetric information. This theory posits that prices adjust in response to order flow because some traders possess private information about exchange rate fundamentals. Public information, such as scheduled macroeconomic announcements, tends to be incorporated quickly, whereas private signals are reflected more gradually through persistent order flow imbalances. In OTC FX markets, inter-dealer order flow transmits customer information and helps dealers manage inventory risk (Evans and Lyons, 2002). That two-tier structure is absent on chain. Instead, informed trading resembles the framework of Kyle (1985), with participants acting on private signals about fundamentals. Such trades may not generate price discovery in the sense of directly moving traditional EUR/USD rates, but they can still contain information that helps predict subsequent changes in benchmark prices.

We therefore hypothesize that the degree of asymmetric information varies systematically across participant groups. Primary dealers, connected to the interbank market through their access to EUR and USD deposits, are more likely to trade on fundamentals rather than on short-lived arbitrage opportunities. Sophisticated traders, operating at greater scale and facing lower effective frictions, are well positioned to transmit private information across venues. In contrast, LPs primarily manage inventory and are subject to adverse selection when prices move on fundamentals (Capponi and Jia, 2021). For example, when fundamentals imply an appreciation of the EUR/USD rate, arbitrageurs buy EURC and sell USDC, leaving LPs exposed to losses. These differences generate systematic heterogeneity in how groups process private information and in the persistence of their price impact.

Hypothesis H2 (Informational Heterogeneity)**.** *Sophisticated traders and primary dealers possess informational advantages that allow them to incorporate private signals about fundamentals into blockchain prices, while LPs remain largely uninformed and focus on inventory management, leaving them exposed to adverse selection.*

We examine these hypotheses by analyzing the determinants of feedback trading, trader behavior during stress episodes such as the USDC de-pegging, and the permanent price impact of order flow.

4.1 Blockchain Order Flow and Feedback Trading

To test Hypothesis H1, we examine whether DEX traders adjust their strategies in response to price differences between the DEX reference rate and the CLS benchmark rate. The analysis is conducted at the hourly frequency to capture short-term arbitrage dynamics and intra-day trading behavior. Specifically, we estimate equation (9), regressing blockchain order flow on the lagged price difference between DEX and traditional markets, with controls that include the lagged EURC/USDC return.

$$OF_{i,t} = \alpha + \beta_1 (p_{\text{EURC/USDC},t-1} - p_{\text{EUR/USD},t-1}) + \text{controls}_t + \epsilon_t \tag{9}$$

[INSERT TABLE 6 ABOUT HERE]

The results presented in Table 6 provide evidence of feedback trading behavior among sophisticated traders. In column (1), a one-unit increase in the lagged hourly price deviation between Uniswap and CLS rates is associated with a reduction in aggregate blockchain order flow of 0.15 million EURC. Column (4) shows a similar magnitude for traders who are both sophisticated and primary dealers, with order flow decreasing by 0.14 million EURC. In contrast, the effects for standalone primary dealers and LPs, reported in columns (2) and (3), are statistically insignificant.

These findings suggest that sophisticated traders and those with dual roles as primary dealers are more likely to engage in arbitrage between decentralized and traditional FX markets. This behavior is likely facilitated by their lower effective trading costs, such as gas fees, which represent a smaller share of transaction size. By contrast, primary dealers and LPs, who typically execute smaller trades, are less responsive to short-term price discrepancies because of higher relative transaction costs.

Primary dealers exhibit limited sensitivity to EURC/USDC price deviations because their smaller transaction sizes make arbitrage economically unattractive once gas fees and slippage are considered. As a result, they are less able to engage in feedback trading compared to sophisticated traders. Instead, their trading activity is more likely motivated by the processing of fundamental information such as macroeconomic news.

Feedback trading also provides a mechanism through which blockchain markets respond to public information. As an illustrative example, we examine market reactions to monetary policy announcements by the Federal Reserve. Appendix D.3 presents an event study comparing FOMC announcement days with a matched control group of non-announcement days. The results show that absolute price differentials between EURC/USDC on Uniswap and EUR/USD on CLS remain small and statistically indistinguishable from those on non-announcement days after the release. At the same time, trading volumes on both markets display a sharp and synchronized spike around announcement times, indicating heightened arbitrage and information-processing activity. These findings demonstrate that blockchain traders quickly incorporate public macroeconomic news into prices and actively arbitrage deviations across venues. Together, they highlight the importance of feedback trading in ensuring that macroeconomic information is efficiently reflected in blockchain prices, consistent with the mechanism proposed in Hypothesis H1.

4.2 Blockchain Order Flow and Fundamental Information

4.2.1 USDC De-Pegging Event

The USDC de-pegging event on March 11, 2023, provides a unique setting to analyze how blockchain market participants respond to market stress under asymmetric information. This event occurred when SVB, which held $3.3 billion of USDC reserves, declared bankruptcy, raising concerns about the backing of USDC and causing its price to drop to 87 cents. Confidence was restored on March 13 after the Federal Deposit Insurance Corporation (FDIC) guaranteed all SVB deposits.[19]

[INSERT FIGURE 8 ABOUT HERE]

[19] Further details on USDC's reserve composition and Circle's response to the de-pegging event are available at https://www.circle.com/blog/an-update-on-usdc-and-silicon-valley-bank.

We use this event to study resilience in the EURC/USDC market and analyze behavior across trader types. Figure 8 illustrates EURC/USDC price deviations from the EUR/USD market and blockchain order flow by trader groups.

Sophisticated Traders. Sophisticated traders showed positive order flow leading up to the event, suggesting informational advantages consistent with Hypothesis H2. This behavior is consistent with evidence presented in Liu et al. (2023), where informed investors responded similarly during the Terra Luna collapse. Transaction-level evidence in Appendix E indicates that sophisticated traders engaged in cross-exchange arbitrage during the de-pegging event. For instance, wallet '1c37' exhibited significant USDC selling pressure, executing large and frequent trades across pairs such as EURC/USDC, USDC-GYEN, and USDC-PRIME on Uniswap and SushiSwap.[20] The pattern suggests cross-venue arbitrage, likely involving the purchase of discounted USDC on centralized exchanges and its resale on decentralized venues.

While this behavior is consistent with informational advantages, another explanation is that sophisticated traders were better positioned to exploit mispricing and arbitrage opportunities because of their scale and routing advantages.[21] Appendix E.3 shows that during the de-pegging period, Top 10 wallets paid a median of only 12.4 USDC in gas fees per 10,000 EURC transacted, which is less than half the corresponding figure for non-Top 10 wallets (32.8 USDC). Gas costs on Ethereum are largely fixed per transaction, so larger investors who trade at higher notional sizes face a much lower effective cost per unit transacted. Sophisticated traders also benefit from routing

[20] Wallet '1c37' (full address: 0xd64137f743432392538a8f84e8e571fa09f21c37) conducted high-volume transactions during the de-pegging event, including major trades in USDC-PRIME, SYN-USDC, EURC-USDC, and USDC-GYEN pairs. Transaction logs show repeated inflows of USDC from Coinbase followed by USDC sales on Uniswap and SushiSwap pools. Detailed transaction logs are provided in Appendix E.

[21] Other channels may also contribute to this advantage. Sophisticated investors typically trade a larger variety of tokens, maintain higher transaction rates, and are able to batch or bundle transactions to optimize gas use, which together enhance execution efficiency relative to smaller participants.

large transactions through private mempools, paying a mostly fixed validator fee while avoiding repeated per-trade gas costs, frontrunning, or Maximum Extractable Value (MEV) extraction.[22] In contrast, retail and smaller institutional wallets rely on public routing, face higher proportional gas costs, and are more exposed to slippage and MEV. Together, these features give sophisticated traders a persistent cost advantage and faster, more secure execution, allowing them to arbitrage price discrepancies more effectively during stress events.

These two factors, scale and routing, help explain why smaller or retail wallets, including some primary dealers, participated less actively in the EURC/USDC market during the de-pegging episode, effectively pricing them out of profitable arbitrage. Combined with the evidence on timing and order flow, this suggests that the dominance of sophisticated traders during the crisis reflects a combination of informational advantages and lower relative trading costs. This interpretation complements the analysis of wallet `1c37` and highlights that access to scale, speed, and low-cost execution is central to understanding which participants engage in real-time arbitrage during de-pegging events.

Liquidity Providers (LPs). In contrast, LPs showed minimal strategic repositioning during this period, consistent with a passive role in which they absorb order flow rather than actively manage exposure.[23] As swap traders purchased EURC during the event, LPs' EURC reserves were drawn down, leaving a liquidity distribution that became increasingly one-sided. At the peak of the de-pegging, most remaining liquidity was concentrated in USDC on the bid side of the pool, reflecting

[22]MEV refers to the profit that validators or searchers can extract by reordering, inserting, or censoring transactions within a block. It arises because validators determine the order of transactions during block construction, which allows them to capture arbitrage or liquidation opportunities before public transactions are finalized.

[23]For example, the only LP withdrawal observed during the event occurred at 05:59 UTC on March 11, involving the removal of EURC and USDC at a mid-range price. This inactivity aligns with LPs focusing on inventory maintenance rather than informed trading.

that swap traders had absorbed nearly all available EURC liquidity. This pattern supports the view that LPs behave passively, with their liquidity positions reflecting mechanical inventory depletion rather than information-driven adjustments. Liquidity became highly asymmetric at the peak of the de-pegging on March 11, 2023 and returned to a more balanced state only after the USDC peg was restored. For further evidence on the evolution of tick-level liquidity during this episode, see Appendix E.2.

4.2.2 Permanent Price Impact

To assess whether blockchain-based trades convey private information about exchange rate fundamentals, we examine the permanent price impact of order flow. Permanent impact refers to the component of order flow that results in lasting changes in prices, distinguishing it from transitory fluctuations due to liquidity provision or noise trading. In the microstructure literature, order flow is often viewed as a proxy for the gradual incorporation of private information into prices, though this has typically been studied in the context of dealer-based OTC markets (Evans and Lyons, 2002).

This mechanism is particularly relevant in blockchain-based FX markets, where low liquidity and participant heterogeneity may amplify the informational role of trades. The availability of wallet-level data enables a more granular analysis of how different trader types, such as sophisticated investors, primary dealers, and LPs, have private information on underlying fundamentals that shifts prices over time.

We estimate a SVAR at the hourly frequency to capture the dynamic relationship between order flow and exchange rate changes. To isolate the informational content of blockchain order flow beyond what is already impounded through institutional trading activity, we control for traditional

FX order flow obtained from CLS data. This results in a block SVAR specification where variables are grouped into three components: traditional OTC order flow ($\mathbf{OF}_t^{OTC}$), blockchain order flow ($\mathbf{OF}_t^{DEX}$), and exchange rate changes (Δp_t):

$$\begin{bmatrix} \mathbf{OF}_t^{OTC} \\ \mathbf{OF}_t^{DEX} \\ \Delta p_t \end{bmatrix} = \alpha + \sum_{k=1}^{L} \mathbf{A}_k \begin{bmatrix} \mathbf{OF}_{t-k}^{OTC} \\ \mathbf{OF}_{t-k}^{DEX} \\ \Delta p_{t-k} \end{bmatrix} + \epsilon_t. \tag{10}$$

The traditional OTC order flow vector $\mathbf{OF}_t^{OTC}$ consists of buy-minus-sell imbalances disaggregated into non-bank financials, corporates, funds, and interbank dealers using CLS data. The blockchain order flow vector $\mathbf{OF}_t^{DEX}$ includes EURC/USDC transaction flows on Uniswap, grouped by wallet type: LPs, residual wallets ($\notin \{\text{Top10}, \text{PM}, \text{LP}\}$), Top10∩LP, PM, Top10, and Top10∩PM, as defined in Section 2. The dependent variable Δp_t denotes the log change in the EUR/USD exchange rate, either from the DEX mid-price or the CLS benchmark rate.

We identify the system using a recursive Cholesky decomposition, assuming the ordering $\mathbf{OF}_t^{OTC} \to \mathbf{OF}_t^{DEX} \to \Delta p_t$. This implies that traditional OTC order flow has a contemporaneous impact on blockchain order flow and prices, while DEX flows do not contemporaneously affect OTC flows. Prices respond immediately to all order flows. This identification scheme reflects the hierarchical structure of FX trading venues and is consistent with information transmission patterns documented in the microstructure literature. Full details of the identification assumptions and block matrix decomposition are provided in Appendix F.

Impulse response functions (IRFs) derived from this SVAR are shown in Figure 9, with Panel (a) reporting responses of EURC/USDC (DEX) returns and Panel (b) showing CLS benchmark EUR/USD returns. The IRFs trace the effect of shocks to DEX order flow components—including sophisticated traders, primary dealers, LPs, and intersecting groups—on returns over a 24-hour horizon.

The results indicate that order flows from sophisticated traders and primary dealers have the most persistent and substantial permanent impact on CLS benchmark returns. Specifically, a 1 million EURC shock to primary dealer order flow generates a permanent return impact of approximately 4.0 percent over 24 hours, while a 1 million EURC shock to sophisticated trader order flow generates a permanent impact of about 2.2 percent. The combined group of Top 10 wallets that are also primary dealers generates a permanent impact of 3.1 percent. In contrast, flows from LPs and residual wallets have negligible permanent effects, with estimates of approximately −0.1 and 0.2 percent respectively, consistent with uninformed trading or liquidity provision.

These findings support Hypothesis H2, revealing informational heterogeneity across blockchain market participants. When scaled to a one-standard-deviation shock in daily order flow, the implied permanent impact on CLS benchmark returns is approximately 1.6 basis points for sophisticated trader flows and 2.9 basis points for primary dealer flows,[24] consistent with earlier studies. Importantly, the informational effects we identify do not arise from inter-dealer price discovery mechanisms, as in Evans and Lyons (2002), but instead reflect private information about exchange rate fundamentals, in the spirit of informed trading models such as Kyle (1985).

[INSERT FIGURE 9 ABOUT HERE]

[24]The permanent price impact estimates reflect IRF responses expressed in percent returns per 1 million EURC order flow shock. Given that a one-standard-deviation shock in daily order flow is approximately 7,230 EURC, the implied return impact scales proportionally. These estimates align with earlier studies of the price impact of traditional FX order flow, including the 50 basis points per USD 1 billion reported by Evans and Lyons (2002) and Berger et al. (2008). The sample average daily trading volume is EURC 0.423 million (standard deviation EURC 0.674 million).

4.3 Private Information and Permanent Price Impact

The evidence presented above supports both microstructural hypotheses. Blockchain order flow responds to public benchmarks and macroeconomic announcements, consistent with feedback trading, and some trader flows generate permanent price impacts that are consistent with informational content. This raises a natural identification question regarding the extent to which the estimated permanent price impacts reflect private information rather than mechanical arbitrage. We address this question in two complementary ways. First, we exploit variation in routing by separating private transactions, which bypass the public mempool, from public transactions, and estimate the SVAR jointly to compare their permanent price effects. Second, we remove the feedback component from order flow by projecting it on lagged DEX–CLS price differences and other controls, and then test whether the residual component, which we interpret as informational order flow, generates permanent price impact. All figures and tables referenced in this subsection are reported in Appendix G.

Private Transactions. Private routing can affect prices through two main mechanisms. One is strategic routing to avoid MEV extraction or frontrunning when executing large arbitrage trades. The other is routing to conceal fundamental information during execution. Both mechanisms can generate price impact, but their motivations differ. MEV trades typically aim to exploit price discrepancies across venues and are closely tied to arbitrage and market efficiency, whereas information-concealing trades are motivated by private signals about fundamentals and are expected to have more persistent effects on benchmark prices. Private transactions are identified using Blocknative mempool archives, as detailed in Appendix C.2.4, and are concentrated among Top10 and Top10∩LP addresses, whereas primary dealers and LPs transact almost exclusively through the public mempool.

We re-estimate the SVAR model, splitting blockchain order flow into private and public components and restricting attention to Top10, primary dealers, and LPs, the groups with meaningful private activity. In this specification, each DEX order flow group is partitioned into a public and a private component, with the private component placed immediately after its corresponding public component in the ordering.[25] This ordering allows private transactions to adjust contemporaneously to public order flow shocks, consistent with the timing of information arrival.

Figure A10 shows that privately routed trades generate larger and more persistent price effects than public transactions, with the strongest responses for sophisticated traders (Top 10) and Top 10∩LP addresses. For benchmark CLS returns, the estimated permanent price impact at a 24-hour horizon of a one million EURC shock from Top 10 wallets is 2.87 bps (95% CI: 2.25, 3.54) for *private* transactions, compared to 1.15 bps (95% CI: 0.29, 1.82) for *public* transactions. For Top 10∩LP flows, the corresponding estimates are 1.12 bps (95% CI: −0.10, 1.82) and −0.58 bps (95% CI: −1.17, −0.13), respectively. These differences indicate that private flows are more informationally rich and associated with more persistent movements in benchmark FX prices.

To further understand these differences, Tables A13 and A14 examine whether public and private order flows are sensitive to DEX–CLS price differentials, which proxy for feedback trading and arbitrage incentives. Public transactions display strong and systematic sensitivity to these price gaps, consistent with arbitrage activity that mechanically aligns prices across venues. In contrast,

[25] Formally, if the baseline SVAR ordering of DEX order flow groups is $[OF_1, OF_2, OF_3, \ldots]$, the revised ordering is $[OF_{1,\text{public}}, OF_{1,\text{private}}, OF_{2,\text{public}}, OF_{2,\text{private}}, OF_{3,\text{public}}, OF_{3,\text{private}}, \ldots]$. This ordering reflects the information structure of blockchain execution: public transactions are broadcast through the mempool, making their information contemporaneously observable to private agents, who can condition on these signals before execution. The remainder of the SVAR structure follows Equation (10), with traditional OTC order flow ordered first, DEX order flow (public and private) ordered second, and exchange rate changes last.

private transactions—particularly those initiated by Top 10 traders—exhibit weak or no relationship with price differentials, suggesting that they are less driven by contemporaneous arbitrage incentives and more by private information on fundamentals. The main exception is the Top 10∩LP group, where private transactions display weakly significant feedback effects. This pattern is consistent with LPs using private routing primarily to execute large repositioning trades and to mitigate frontrunning or MEV extraction, rather than to exploit private information.

In sum, private transactions reflect the concealment of information on fundamentals, suggesting that the permanent price impact of sophisticated investors is driven more by private information than by pure arbitrage. This motivates our second test, which directly separates feedback-driven trading from informational order flow.

Feedback Trading vs Informational Order Flow. To distinguish informational effects from mechanical arbitrage more directly, we decompose blockchain order flow into two components—a feedback-driven component predicted by the lagged price difference between the EURC/USDC price on DEX and the EUR/USD benchmark, and a residual component that we interpret as informational order flow. This is implemented by estimating equation (9), regressing net EURC buyer flow on the lagged DEX–CLS price difference. The fitted values represent the feedback component, while the residual captures order flow orthogonal to these price discrepancies.

Figure A11 shows the impulse responses of EUR/USD CLS benchmark returns to these components. Panel (a) demonstrates that the residual (informational) component generates persistent and statistically significant price impacts, particularly for primary dealers and sophisticated traders. Panel (b) shows that the feedback component has no significant impact on benchmark FX prices. This pattern indicates that the permanent price effects we document are primarily driven by the incorporation of fundamental information rather than mechanical arbitrage between markets.

4.4 Robustness Tests

To validate the robustness of our findings, we conduct several additional tests, which are summarized below and detailed in Appendix H.

SVAR Price Impact. We assess the robustness of our price impact estimates by controlling for three potential sources of variation: liquidity provision, trading frequency and trade size. All corresponding SVAR impulse response figures are reported in Appendix H.1.

First, our price impact estimates could be an outcome of changes in liquidity provision rather than the informational content of order flow. Informed LPs may adjust their supply of currencies at different price levels in response to expected return movements, generating a contemporaneous correlation between order flow and prices. To address this, we re-estimate the baseline SVAR specification (Equation (10)) while controlling for net liquidity provision both near the best price (within $\pm 1\%$) and away from it (beyond $\pm 1\%$). The estimated coefficients remain quantitatively similar, suggesting that liquidity adjustments do not drive our main findings.

Second, we test whether informational content varies at higher frequencies. We re-estimate the SVAR using 5-minute intervals to capture finer timing between blockchain and benchmark market adjustments. At this frequency, the price impact of top wallets and primary dealers remains economically and statistically significant, particularly on the DEX leg. In contrast, other wallets and LPs show weaker and less persistent effects, consistent with a lower informational content of their trades.

Third, we examine whether price impact differs by trade size. Larger and more sophisticated traders tend to have higher turnover and may convey information or engage in meaningful arbitrage. We group traders into quintiles based on average transaction volume and re-estimate the SVAR within each group. Larger traders, in the top quintile, display the strongest and most persistent price impacts on both DEX and CLS returns, while smaller traders exhibit weaker effects.

Intra-day Variation. We next analyze the intra-day pattern of DEX order flow for each trading group (Appendix H.2). The price impacts of sophisticated traders and primary dealers are strongest between 13:00 and 15:00 UTC, overlapping with active European and U.S. trading hours when macroeconomic information is most abundant. LPs show no significant impacts during these periods, indicating that their activity is driven more by inventory management than by information processing.

Just-in-time Liquidity. A related concern is that some LPs may engage in high-frequency strategies such as just-in-time (JIT) liquidity provision, where providers add and remove liquidity strategically to capture fees while minimizing exposure to adverse selection (Capponi et al., 2024a). Our analysis in Appendix H.3 identifies one wallet, ending in 'ae13', that consistently exhibits JIT behavior in the EURC/USDC pool.[26] Although this behavior demonstrates a sophisticated liquidity strategy, it is rare and does not materially affect aggregate price dynamics.

Overall, these tests confirm that our results are robust to alternative specifications, trader composition, and liquidity behavior, reinforcing the interpretation that informed order flow drives blockchain price discovery.

[26]For example, on 23 August 2023, wallet 'ae13' deposited 50,249 EURC and 311,077 USDC within a price range of 1.0898 to 1.0909. Shortly after a large trade by another user, wallet '2cc4', it withdrew its liquidity, burning 32,048 EURC and 330,931 USDC. This pattern reflects strategic fee capture with minimal exposure to adverse selection.

5 Conclusion

DeFi platforms introduce new trading structures based on smart contracts and automated liquidity provision, offering a market design that differs from traditional dealer-intermediated FX markets. This paper examines the efficiency of blockchain-based currency markets and their connection to traditional FX markets, focusing on the EURC/USDC pair on the Uniswap V3 decentralized exchange. EURC/USDC prices track traditional FX benchmarks closely, with small deviations driven by blockchain frictions. Trading activity is concentrated during market openings, and while transaction costs exceed inter-dealer spreads, they remain comparable to those faced by less privileged OTC clients.

We analyze the mechanisms linking blockchain and traditional FX markets through two channels: feedback trading, which exploits price differences between venues, and asymmetric information, where traders act on private signals about exchange rate fundamentals. Using a SVAR framework, we show that order flow from sophisticated traders and primary dealers contains information about future benchmark movements, while LP trades have limited informational content. The evidence supports the asymmetric information view of FX markets, indicating that some blockchain traders possess private information about the fundamental exchange rate rather than merely reacting to observed price changes.

Although still small, blockchain currency markets have important implications for policy and market structure. Arbitrage and the processing of fundamental information enhance their efficiency and linkage with traditional benchmarks. As technology advances and blockchain frictions diminish, they are likely to develop as a complementary venue for cross-border trading and settlement.

References

Adams, Austin, Mary-Catherine Lader, Gordon Liao, David Puth, and Xin Wan, "On-Chain Foreign Exchange and Cross-Border Payments," *Available at SSRN 4328948*, 2023.

Adams, Hayden, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson, "Uniswap v3 whitepaper," *Tech. rep., Uniswap, https://uniswap.org/whitepaper-v3.pdf*, 2021.

Aldasoro, Iñaki, Rashad Ahmed, and Chanelle Duley, "Par for the course: Public information and stablecoin runs," *Available at SSRN*, 2023.

Barbon, Andrea and Angelo Ranaldo, "On The Quality Of Cryptocurrency Markets: Centralized Versus Decentralized Exchanges," *Available at SSRN 3984897*, 2024.

Berger, David W, Alain P Chaboud, Sergey V Chernenko, Edward Howorka, and Jonathan H Wright, "Order flow and exchange rate dynamics in electronic brokerage system data," *Journal of international Economics*, 2008, *75* (1), 93–109.

Bjønnes, Geir Høidal and Dagfinn Rime, "Dealer behavior and trading systems in foreign exchange markets," *Journal of Financial Economics*, 2005, *75* (3), 571–605.

Capponi, Agostino and Ruizhe Jia, "The Adoption of Blockchain-based Decentralized Exchanges: A Market Microstructure Analysis of the Automated Market Maker," *Available at SSRN 3805095*, 2021.

_ , _ , and Brian Zhu, "The paradox of just-in-time liquidity in decentralized exchanges: More providers can sometimes mean less liquidity," *Available at SSRN*, 2024.

— , — , and Shihao Yu, "Price discovery on decentralized exchanges," *Available at SSRN 4236993*, 2024.

Chaboud, Alain, Dagfinn Rime, and Vladyslav Sushko, "The foreign exchange market," in "Research Handbook of Financial Markets," Edward Elgar Publishing, 2023, pp. 253–275.

Eichengreen, Barry, My T Nguyen, and Ganesh Viswanath-Natraj, "Stablecoin Devaluation Risk," *WBS Finance Group Research Paper*, 2023.

Evans, Martin DD and Richard K Lyons, "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, 2002, *110* (1), 170–180.

Filippou, Ilias, Thomas A Maurer, Luca Pezzo, and Mark P Taylor, "Importance of transaction costs for asset allocation in foreign exchange markets," *Journal of Financial Economics*, 2024, *159*, 103886.

Foley, Sean, Peter O'Neill, and Tālis J Putniņš, "A Better Market Design? Applying 'Automated Market Makers' to Traditional Financial Markets," *Applying 'Automated Market Makers' to Traditional Financial Markets (May 26, 2023)*, 2023.

Gorton, Gary B, Elizabeth C Klee, Chase P Ross, Sharon Y Ross, and Alexandros P Vardoulakis, "Leverage and Stablecoin Pegs," Technical Report, National Bureau of Economic Research 2022.

Gromb, Denis and Dimitri Vayanos, "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of financial Economics*, 2002, *66* (2-3), 361–407.

Hagströmer, Björn and Albert J Menkveld, "Information revelation in decentralized markets," *The Journal of Finance*, 2019, *74* (6), 2751–2787.

Hasbrouck, Joel, "Measuring the information content of stock trades," *The Journal of Finance*, 1991, *46* (1), 179–207.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory pricing of over-the-counter derivatives," *Management Science*, 2021, *67* (11), 6660–6677.

He, Zhiguo, Bryan Kelly, and Asaf Manela, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126* (1), 1–35.

Hortaçsu, Ali and Samita Sareen, "Order flow and the formation of dealer bids: information flows and strategic behavior in the Government of Canada securities auctions," 2005.

Huang, Wenqian, Angelo Ranaldo, Andreas Schrimpf, and Fabricius Somogyi, "Constrained Dealers and Market Efficiency," *Journal of Financial Economics, Forthcoming*, 2025.

_ , Peter O'Neill, Angelo Ranaldo, and Shihao Yu, "HFTs and Dealer Banks: Liquidity and Price Discovery in FX Trading," *FCA Occasional Paper*, 2023, *63* (2023), 23–48.

King, Michael R, Carol Osler, and Dagfinn Rime, "Foreign exchange market structure, players, and evolution," *Handbook of exchange rates*, 2012, pp. 1–44.

Klein, Olga, Roman Kozhan, Ganesh Viswanath-Natraj, and Junxuan Wang, "Price Discovery in Cryptocurrencies: Trades versus Liquidity Provision," *Available at SSRN 4642411*, 2024.

Kloks, Peteris, Edouard Mattille, and Angelo Ranaldo, "Foreign Exchange Swap Liquidity," *Swiss Finance Institute Research Paper*, 2023, (23-22).

Krohn, Ingomar, Philippe Mueller, and Paul Whelan, "Foreign exchange fixings and returns around the clock," *The Journal of Finance*, 2024, *79* (1), 541–578.

Kyle, Albert S, "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, 1985, pp. 1315–1335.

Lehar, Alfred and Christine Parlour, "Decentralized exchange: The uniswap automated market maker," *The Journal of Finance*, 2025, *80* (1), 321–374.

Liu, Jiageng, Igor Makarov, and Antoinette Schoar, "Anatomy of a run: The terra luna crash," Technical Report, National Bureau of Economic Research 2023.

Lyons, Richard K, "Tests of microstructural hypotheses in the foreign exchange market," *Journal of Financial Economics*, 1995, *39* (2-3), 321–351.

_ **and Ganesh Viswanath-Natraj**, "What keeps stablecoins stable?," *Journal of International Money and Finance*, 2023, *131*, 102777.

Ma, Yiming, Yao Zeng, and Anthony Lee Zhang, "Stablecoin runs and the centralization of arbitrage," Technical Report, National Bureau of Economic Research 2025.

Malinova, Katya and Andreas Park, "Learning from DEFI: Would automated market makers improve equity trading?," *Available at SSRN*, 2024, *4531670.*

Ranaldo, Angelo and Fabricius Somogyi, "Asymmetric information risk in FX markets," *Journal of Financial Economics*, 2021, *140* (2), 391–411.

Figure 1: Structure of Traditional and Blockchain Market

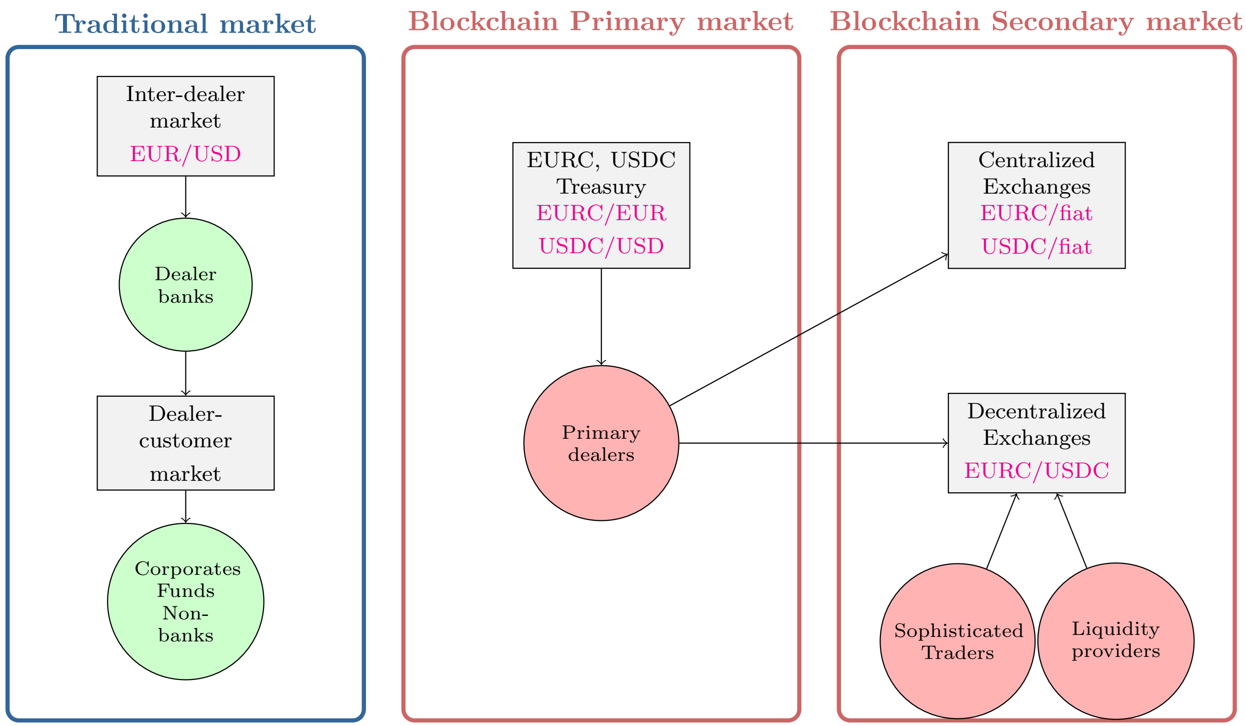

Note: This figure presents a schematic of both traditional and blockchain markets. Traditional markets have an inter-dealer market intermediated by dealer banks, that provide liquidity in the dealer-customer market, trading with corporates, funds and non-bank financial companies. The blockchain market has both a primary and secondary market. In the primary market, the Treasury, managed and operated by Circle, mint EURC tokens and USDC tokens, which are then distributed to "primary dealers", that distribute EURC and USDC tokens in the secondary market. Secondary market trading consists of trading in centralized exchanges that deal in limit order books, or alternatively on decentralized exchanges like Uniswap that trade on EURC/USDC. Other trading types on decentralized exchanges include LPs and sophisticated traders.

Figure 2: EURC/USDC Bonding Curves and Slippage

Panel (a): EURC/USDC Bonding Curves: Swap and Liquidity Trades

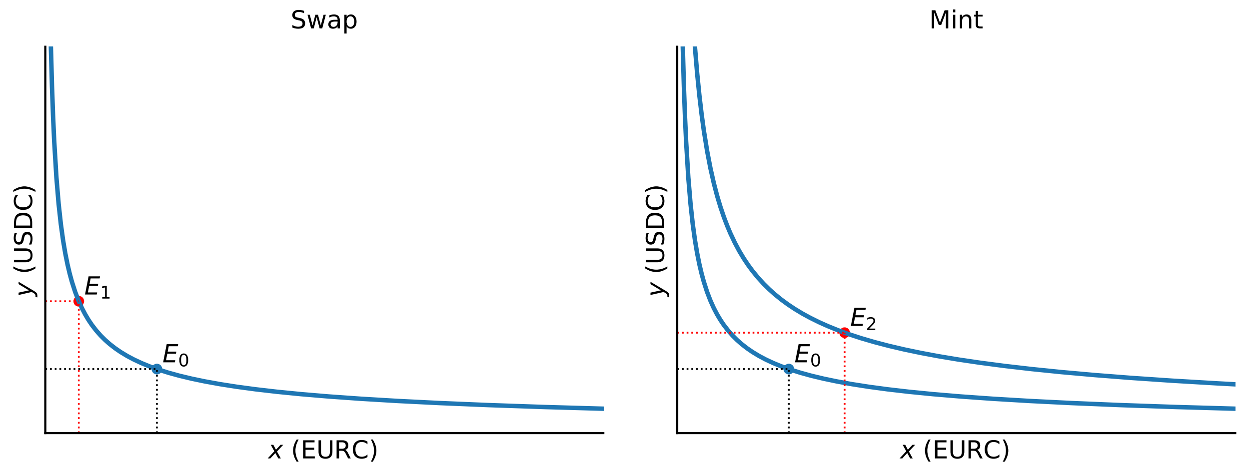

Panel (b): Bonding Curve and Slippage: Numerical Example

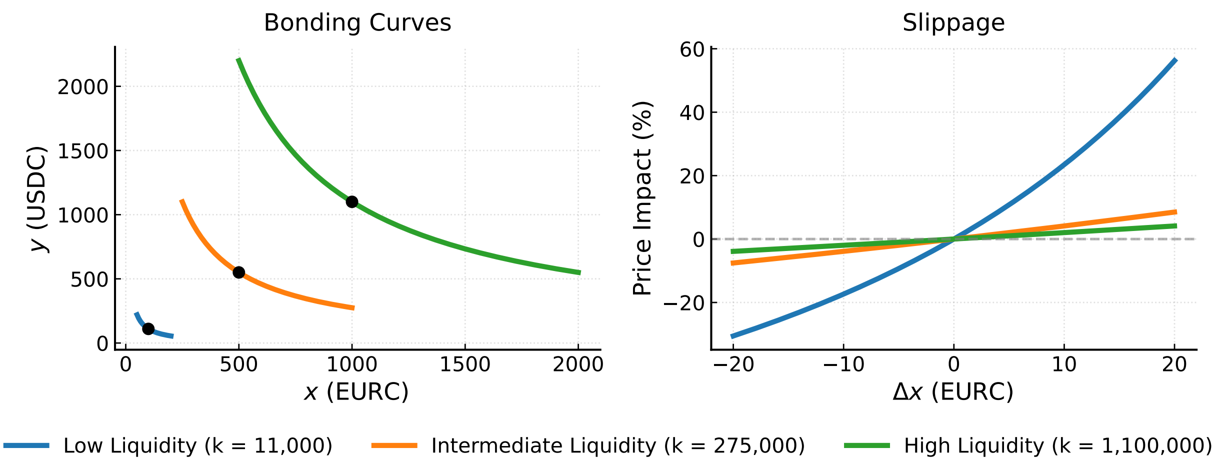

Note: Panel (a) illustrates the principles of a bonding curve and liquidity provision in Uniswap. The aggregate supply of liquidity is initially at point E_0, with a swap trade of purchasing EURC moving the equilibrium from E_0 to E_1, and an LP adding liquidity at the current price from E_0 to E_2. Panel (b) illustrates the relationship between pool size and price impact in a constant product market maker for the EURC/USDC pair. The left subpanel shows bonding curves for three liquidity levels. The right subpanel shows percentage price impact from trades of different sizes (in EURC). Price impact is lower in more liquid pools, reflecting reduced slippage.

Figure 3: **Snapshot of EURC/USDC Liquidity.**

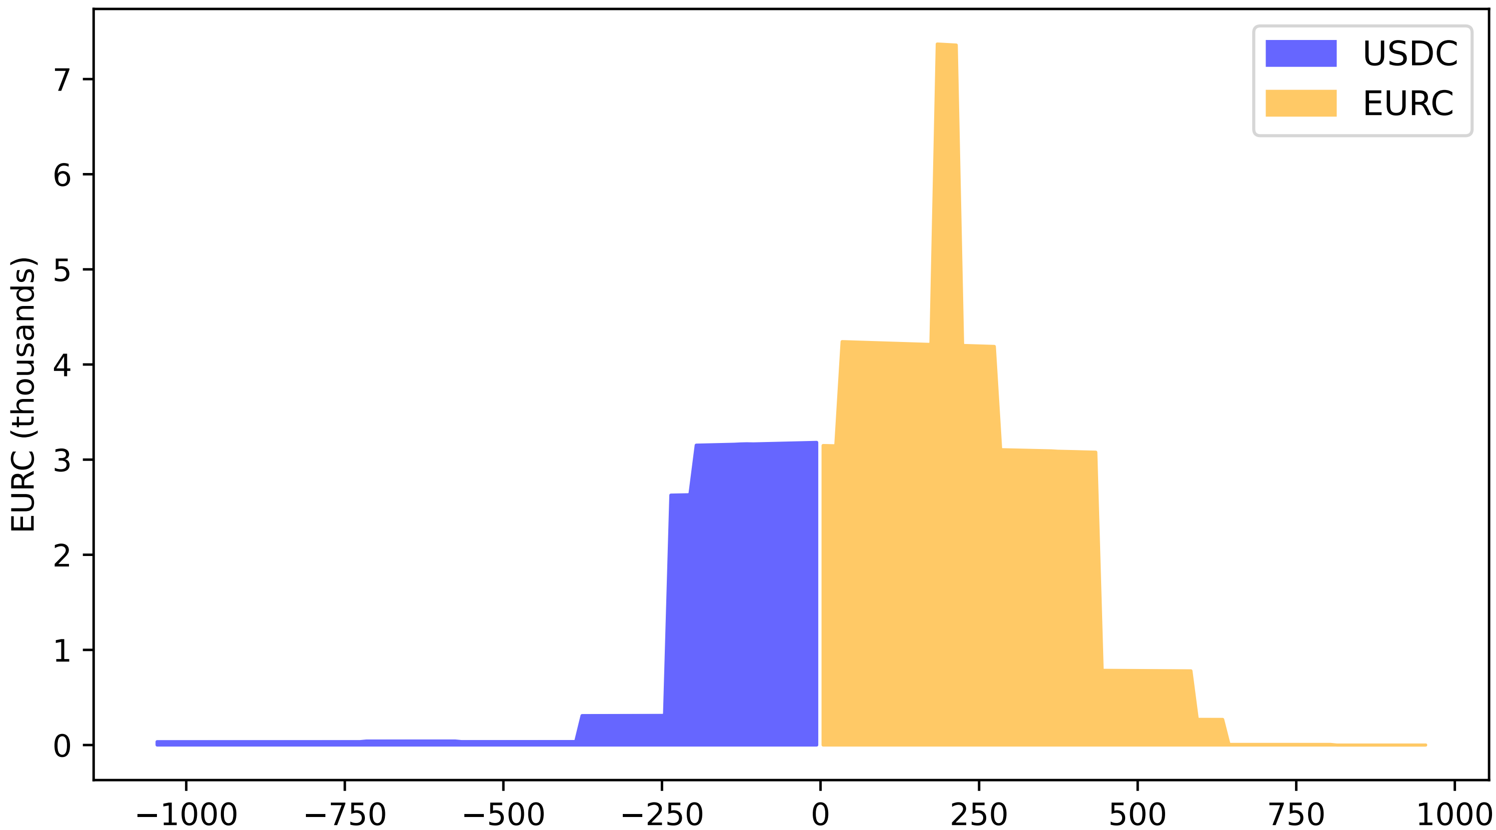

This figure displays the tick-level distribution of liquidity around the prevailing pool price for the EURC/USDC 0.05% Uniswap V3 pool, observed at block 19,771,559 on April 30, 2024. The horizontal axis measures tick distance from the current market price (tick 0), where each tick represents a discrete price interval in log base $\sqrt{1.0001}$ units. The pool has a fixed tick spacing of 10, implying price intervals of approximately 0.1% (10 basis points), which determine the granularity of allowable liquidity placements. Ticks to the left of zero represent liquidity *below* the current market price and correspond to *buy limit orders for EURC* (i.e., LPs are willing to purchase EURC by selling USDC). Ticks to the right of zero represent liquidity *above* the current market price and correspond to *sell limit orders for EURC* (i.e., LPs are willing to sell EURC in exchange for USDC). Liquidity in each tick is expressed in thousands of EURC on both sides of the book.

Figure 4: EURC/USDC Prices

Panel (a): EURC/USDC Price (Uniswap) and EUR/USD Price (CLS)

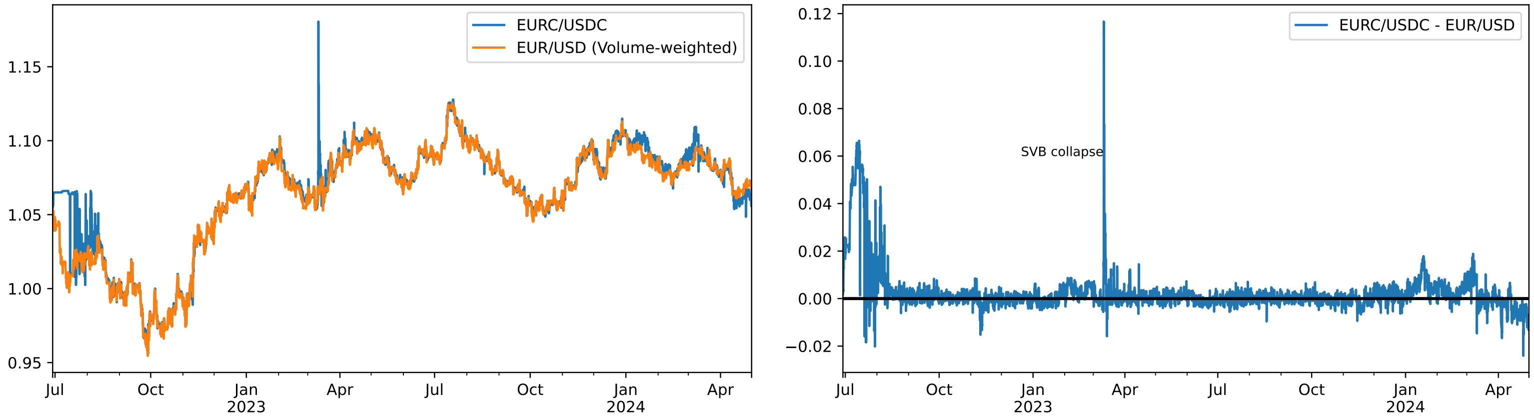

Panel (b): EURC/USDC Price and Cumulative Blockchain Order Flow

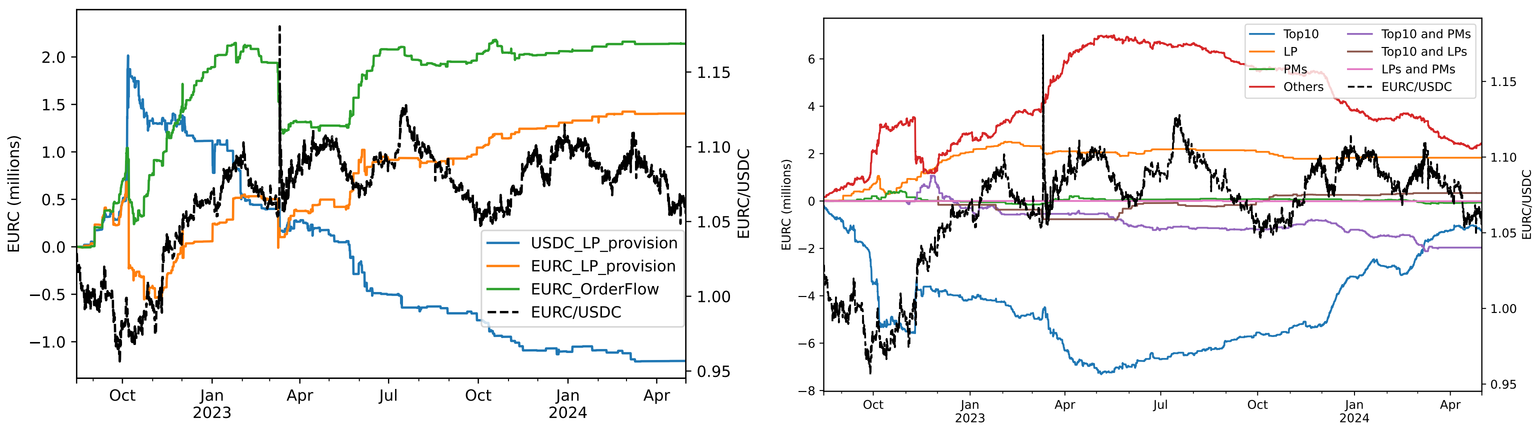

Note: This figure plots EURC/USDC and EUR/USD prices. EURC/USDC prices are sourced from Uniswap V3, and EUR/USD prices are sourced from CLS. Panel (a) shows the EURC/USDC price, the traditional (CLS) EUR/USD price, and their price difference across markets. Panel (b) reports cumulative order flow and price in the EURC/USDC market. The left panel of (b) includes aggregate order flow, price, and liquidity provision in both EURC and USDC. The right panel disaggregates cumulative blockchain order flow into the following trading groups: sophisticated traders (Top 10 wallets), primary dealers (PM), and liquidity providers (LPs). It also reports activity for the intersections of sophisticated traders with primary dealers and LPs, as well as all other traders not belonging to these groups. The sample period for Panel (a) spans 28 June 2022 to 30 April 2024, and for Panel (b) spans 15 August 2022 to 30 April 2024.

Figure 5: Hourly FX Trading Volume

Panel (a): DEX Trading Volume

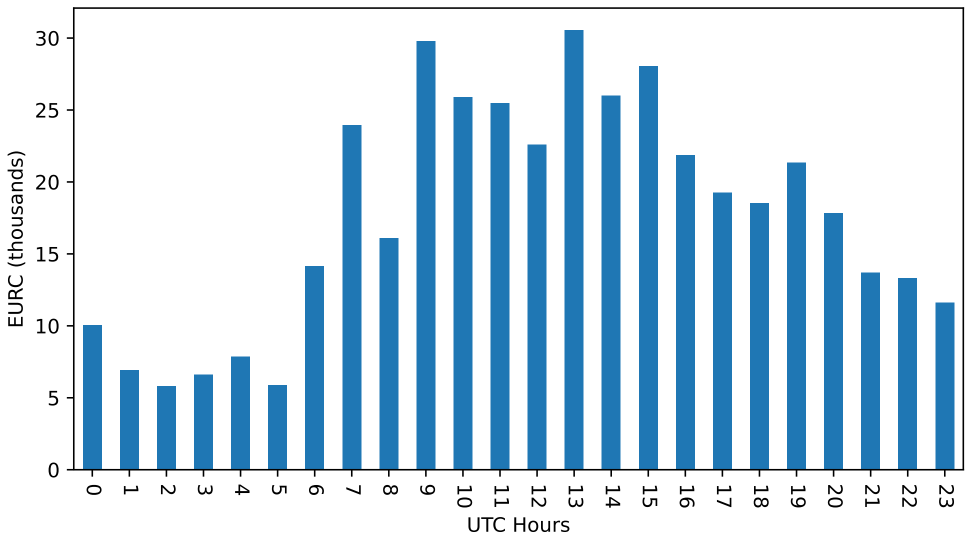

Panel (b): CLS Trading Volume

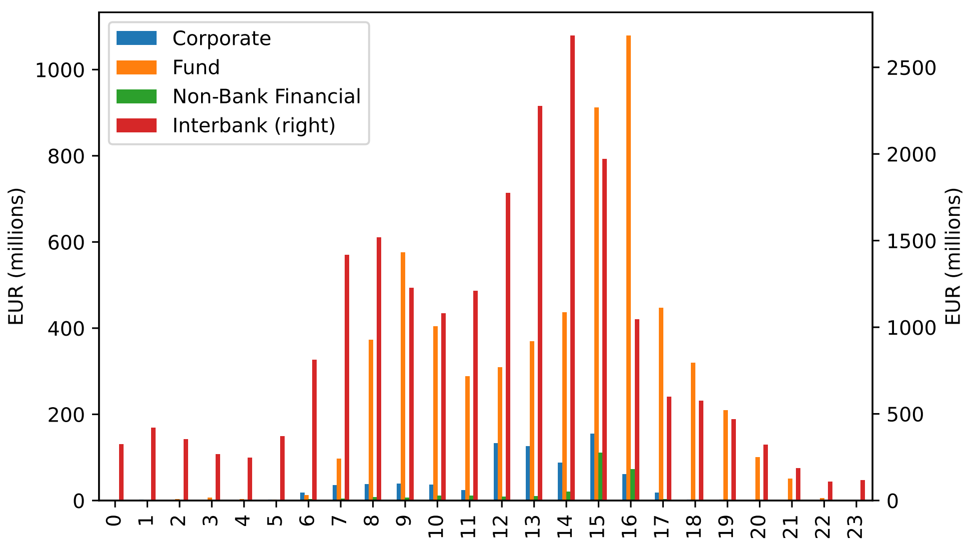

Note: This figure plots hourly trading volume. In Panel (a), we report trading on Uniswap V3 in the EURC/USDC Market in EURC Millions. In Panel (b), we report trading volume on CLS for the EUR/USD market, disaggregated by sectors: Bank-Bank, Bank-Fund, Bank-Corporate, and Non-Bank Financial-Bank. CLS Volume is in EUR Million. The total sample period starts on 15 August 2022, and ends on 30 April 2024.

Figure 6: Weekend and Weekday Volume by Trader type

Panel (a): Weekday trading: traditional hours versus close

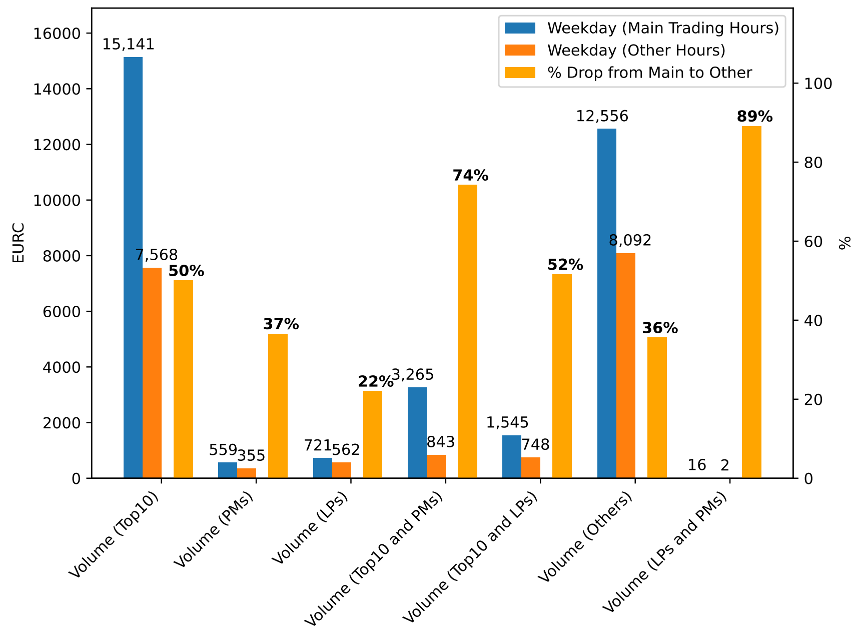

Panel (b): Weekday vs Weekend trading

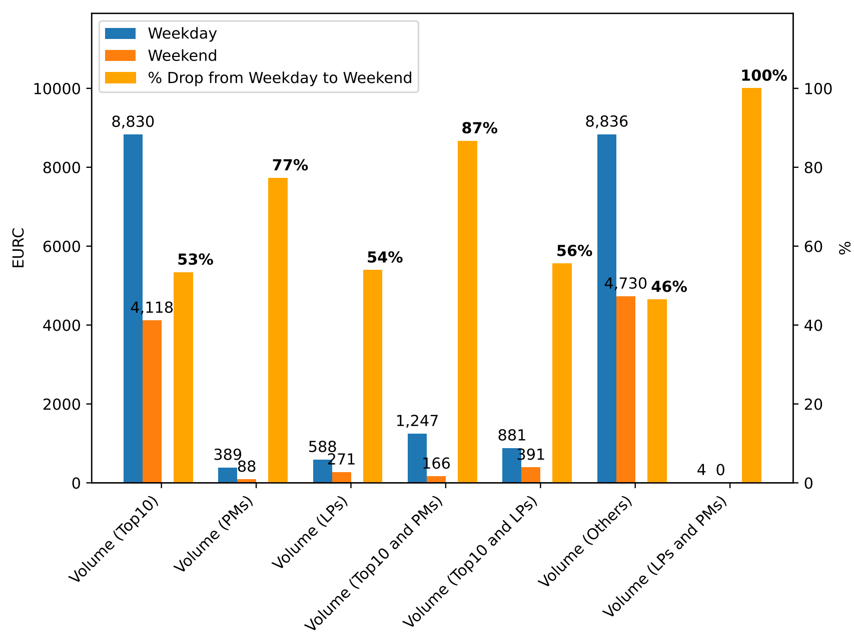

Note: The figure plots average hourly trading volume, distinguishing between weekday and weekend trading for each group. In Panel (a), we compare trading volume for each group during traditional primary opening hours (13–16 UTC) versus other hours on weekdays. Panel (b) presents average trading volume for each group over weekdays and weekends. All volumes are expressed in EURC. Blockchain volume is categorized into seven sub-groups: sophisticated traders (top 10 wallets), primary dealers, LPs, and combinations of these groups. The sample period spans from 15 August 2022 to 30 April 2024.

Figure 7: Trading Costs

Panel (a): Trading Costs: Inter-Quartile Range

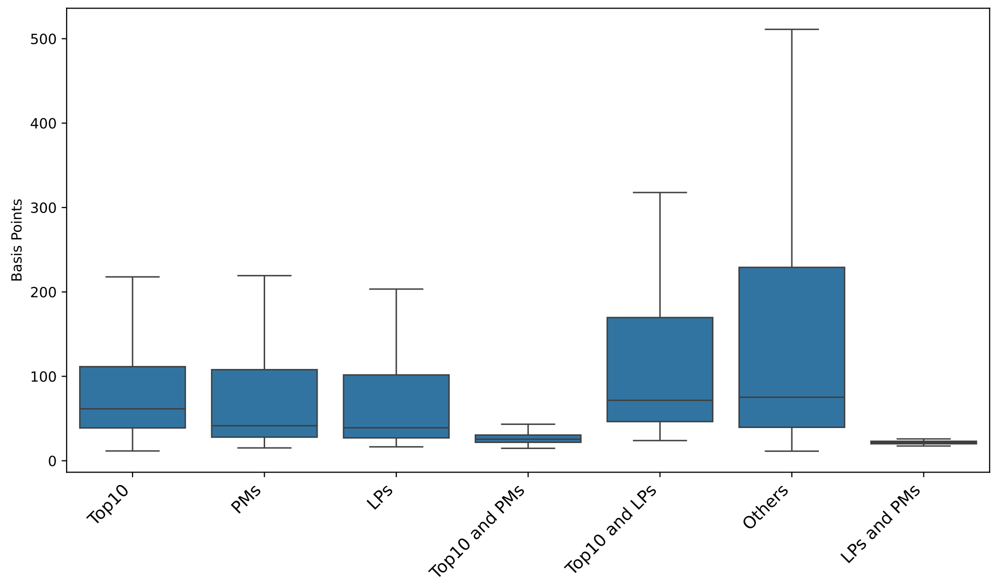

Panel (b): Trading Costs: Median Decomposition

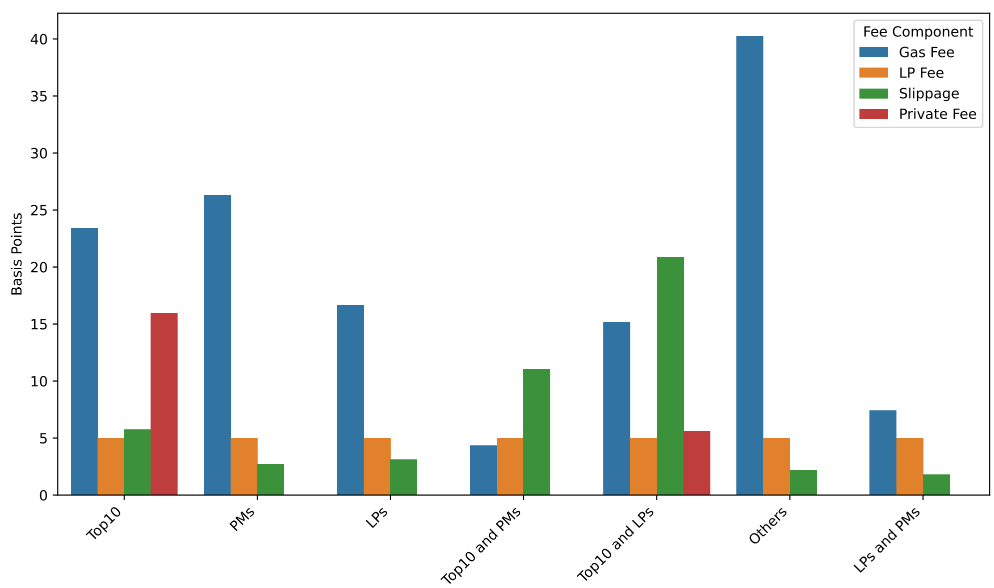

Note: This figure presents trading cost metrics for the EURC/USDC market. Panel (a) plots the inter-quartile range of total transaction costs across account types. The measure combines gas fees (based on ETH transaction fees converted to USD), private fees (transfers to validators for authenticating privately routed transactions), LP fees (5 basis points for the EURC–USDC pool), and slippage, all measured in basis points. Panel (b) decomposes the median transaction cost into its component parts: gas fees, private fees, LP fees, and slippage. Fee components are expressed in basis points and grouped by account type, with medians shown. Account types are categorized into seven sub-groups: sophisticated traders (Top 10 wallets), primary dealers, LPs, and combinations of these groups. Sample period is from 1 March 2023 to 30 April 2024.

Figure 8: USDC De-Pegging event: blockchain order flow of different trading groups

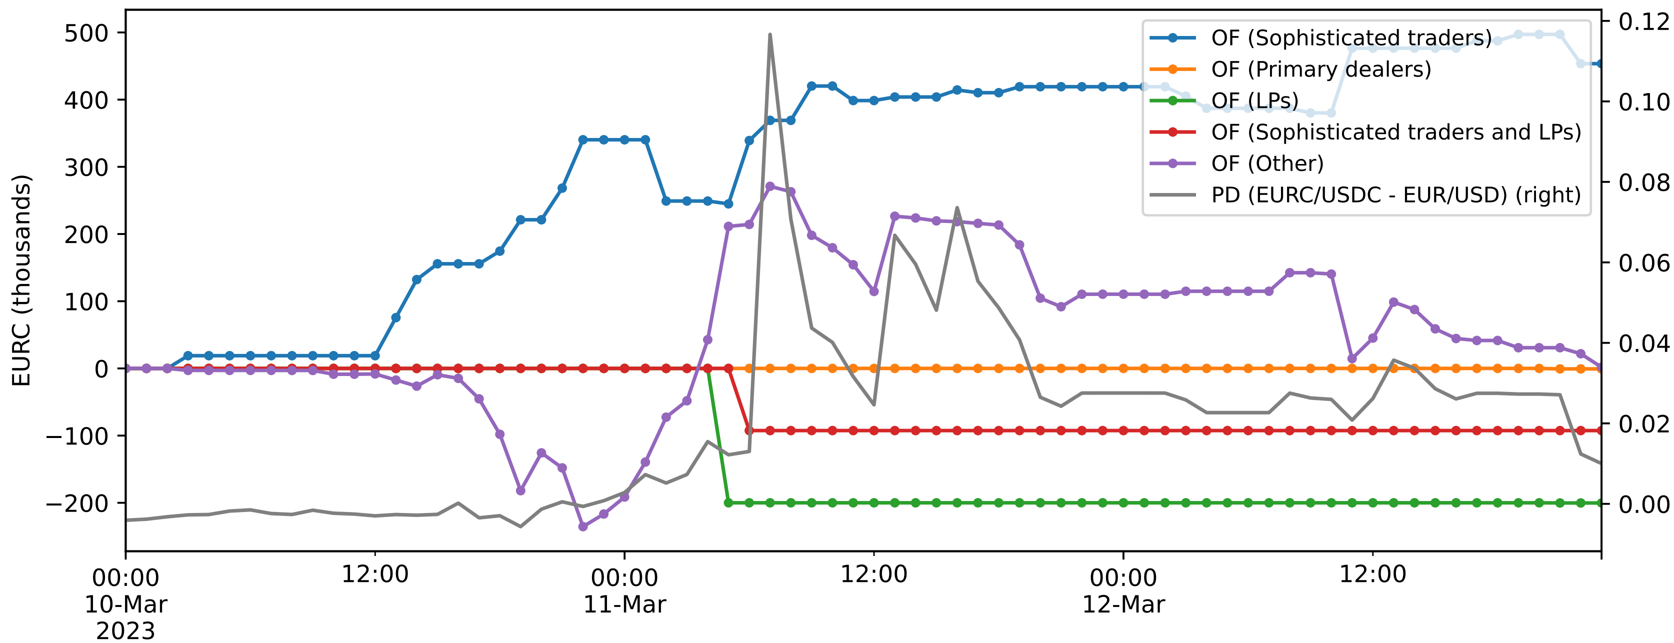

Note: This figure plots the response of blockchain order flow to the de-pegging event of USDC. PD is the difference between EURC/USDC and EUR/USD prices, sourced from Uniswap V3 and CLS respectively. OF is measuring the net buyer transactions of purchasing EURC, and is sourced from Uniswap V3 trade data. Cumulative blockchain order flow is divided into the following sub-categories: sophisticated traders (top 10 wallets), primary dealers, and LPs, denoted by OF_{top10}, OF_{PM} and OF_{LP} respectively. Additionally, we include blockchain order flow of the intersection of sophisticated traders and primary dealers, $OF_{top10 \cap PM}$, and the intersection of sophisticated traders and LPs, $OF_{top10 \cap LP}$, and blockchain order flow of traders that do not belong to the three groups, $OF_{\notin top10, PM, LP}$. Total sample period is from 10 March 2023 to 12 March 2023.

Figure 9: Price impact of blockchain order flow by trader groups

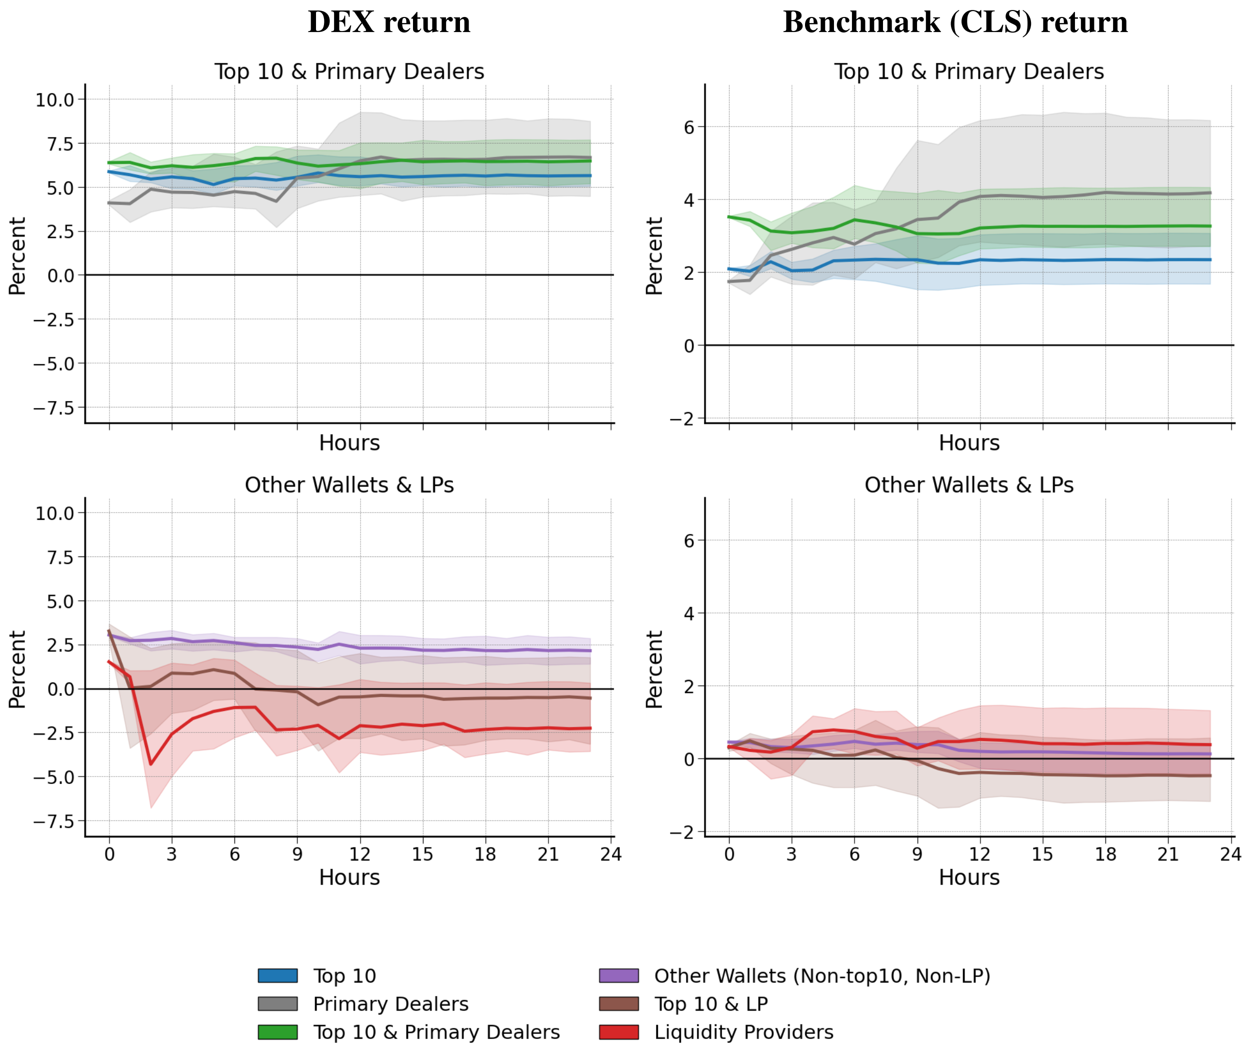

Note: This figure plots impulse responses of returns to shocks in blockchain order flow using hourly data. The top row corresponds to transactions by Top 10 wallets and primary dealers, while the bottom row corresponds to other wallets and LPs. The left column shows EURC/USDC returns from Uniswap V3, and the right column shows EUR/USD returns from CLS. Responses are estimated using a structural VAR with 1,000 bootstrap replications. The sample covers 15 August 2022 to 30 April 2024.

Table 1: Institutional Differences between Traditional and Blockchain-Based Currency Markets

Dimension	Traditional FX Market	Blockchain-Based Market (Uniswap)
Market structure	Over-the-counter dealer market organized around inter-dealer and dealer–client segments.	Decentralized exchange structure based on algorithmic smart contracts and liquidity pools.
Price formation	Dealer quotes reflect order flow, inventory, and competition across platforms.	Prices determined mechanically by the AMM invariant (e.g., $x \times y = k$ or its V3 generalization).
Liquidity provision	Dealers supply liquidity subject to balance sheet and risk constraints.	Liquidity providers supply token pairs across specified price ranges; liquidity is continuous and rule-based.
Transparency	Trade and quote data are proprietary and often delayed.	All transactions and liquidity positions are publicly observable on-chain at the wallet level.
Settlement	Typically T+2 via CLS or correspondent banking networks.	Atomic settlement within each blockchain block.
Arbitrage mechanism	Cross-venue arbitrage constrained by capital, latency, and credit limits.	On-chain arbitrage is permissionless but subject to gas fees and transaction latency.
Primary–secondary linkage	No redemption commitment; exchange rates determined in OTC trading.	Stablecoin issuers guarantee par convertibility (e.g., 1 USDC = 1 USD), linking on- and off-chain prices.
Data granularity	Aggregated by institution or platform; trader-level data rarely observed.	Wallet-level data reveal participant behavior and enable classification by trading role.

Note: This table compares the institutional structure of traditional foreign exchange (FX) markets with that of blockchain-based markets such as Uniswap. AMM refers to automated market maker. CLS denotes Continuous Linked Settlement, the global FX settlement system. Traditional FX markets rely on dealer intermediation, while blockchain-based markets enable peer-to-peer trading through algorithmic liquidity pools.

Table 2: Trader classification

Panel (a): Number of transactions

Group	top10	PrimaryDealer	LP	$N_{addresses}$	Tx	Tx/$N_{addresses}$
Top10	✓	×	×	76	4439	58.41
PM	×	✓	×	68	363	5.34
LP	×	×	✓	90	446	4.96
Top10 ∩ PM	✓	✓	×	6	534	89.00
Top10 ∩ LP	✓	×	✓	7	249	35.57
PM ∩ LP	×	✓	✓	3	6	2.00
$\notin Top10, PM, LP$	×	×	×	2342	9118	3.89

Panel (b): Volume per transaction (EURC)

Group	mean	std	min	25%	50%	75%	max
Top10	25,301	48,886	1	7,845	13,715	27,545	1,040,295
PM	12,528	18,558	3	991	8,000	18,596	183,500
LP	16,752	25,887	1	1,149	8,079	24,260	289,800
Top10 ∩ PM	26,373	10,664	100	20,000	25,000	30,000	95,990
Top10 ∩ LP	44,665	62,339	100	4,290	31,212	50,000	343,333
PM ∩ LP	7,537	9,931	352	2,394	4,556	6,262	27,256
$\notin Top10, PM, LP$	12,611	21,334	0	1,061	5,069	15,169	557,076

Note: Panel (a) presents summary statistics for the number of transactions (Tx) of different trading groups, and the transactions per unique address (Tx/$N_{address}$). Panel (b) presents summary statistics for the volume per transaction in EURC for different trading groups. We characterize wallets in the following trading groups: sophisticated traders (top 10 wallets), primary dealers, and are LPs, denoted by Top10, PM and LP respectively. Additionally, we include sub-categories of traders that are the intersection of sophisticated traders and have primary dealers, Top10 ∩ PM, the intersection of sophisticated traders and LPs, Top10 ∩ LP, and traders that do not belong to the three groups, $\notin \{Top10, PM, LP\}$. Sample period is from 15 August 2022 to 30 April 2024.

Table 3: Summary statistics: Prices, Volume and Blockchain Variables

	count	mean	std	min	25%	50%	75%	max
Panel (a): Trading Volume (CLS) - EUR Billion								
Volume-Corporate-Bank	625	0.777	1.255	0.000	0.000	0.450	0.924	11.018
Volume-Fund-Bank	625	6.003	6.062	0.000	0.000	6.111	8.552	44.678
Volume-Non-Bank Financial-Bank	625	0.275	1.106	0.000	0.000	0.030	0.106	10.331
Volume-Interbank	625	21.366	15.671	0.000	0.354	25.560	31.197	82.861
Volume-Aggregate	625	28.421	20.657	0.000	0.354	34.114	42.077	94.397
Panel (b): Trading Volume (Uniswap)- EURC Million								
Volume (Aggregate)	625	0.423	0.674	0.0001	0.103	0.232	0.490	8.545
Volume (top10)	625	0.180	0.341	0.0	0.015	0.067	0.199	3.453
Volume (PM)	625	0.007	0.020	0.0	0.000	0.000	0.002	0.184
Volume (LP)	625	0.012	0.036	0.0	0.000	0.000	0.002	0.464
Volume (top10 $\cap$ PM)	625	0.023	0.047	0.0	0.000	0.000	0.030	0.343
Volume (top10 $\cap$ LP)	625	0.018	0.084	0.0	0.000	0.000	0.000	1.381
Volume ($\notin \{Top10, PM, LP\}$)	625	0.184	0.360	0.0	0.042	0.097	0.193	5.259
Volume (PM $\cap$ LP)	625	0.0001	0.0013	0.0	0.000	0.000	0.000	0.027
Panel (c): Additional Variables								
$\mathrm{p}_{\mathrm{EURC/USDC}}$	625	1.067	0.035	0.962	1.058	1.078	1.091	1.128
$\mathrm{p}_{\mathrm{EUR/USD}}$	625	1.066	0.035	0.960	1.058	1.077	1.089	1.124
$\lvert \mathrm{p}_{\mathrm{EUR/USD}} - \mathrm{p}_{\mathrm{EURC/USDC}} \rvert$	625	0.002	0.003	0.000	0.001	0.002	0.003	0.028
σ_{ETH}	625	0.007	0.002	0.003	0.005	0.006	0.008	0.013
GasFee	625	0.006	0.001	0.004	0.005	0.006	0.007	0.009
R_{ETH}	624	0.001	0.031	-0.189	-0.012	0.000	0.015	0.160

Note: Panel (a) presents summary statistics of trading volume for EUR/USD pair from CLS. CLS volume is measured in EUR Billions, and is aggregated as well as in the following sub-categories: BuySide Bank-SellSide, Corporate-Bank, Fund-Bank and Non-Bank Financial-Bank volume. Panel (b) presents summary statistics of trading volume for the EURC/USDC pair from Uniswap. DEX volume is divided into different trading groups based on whether they are sophisticated traders (top10), primary dealers (PM), or are LPs. See classification in Table 2 for more details. Panel (c) presents summary statistics of a series of price, blockchain and traditional FX market statistics. Blockchain characteristics include the returns and volatility of Coinbase ETH/USD, and an index of gas fees. Sample period is from 15 August 2022 to 30 April 2024.

Table 4: DEX and CLS Volume correlations

	V_{top10}	V_{PM}	V_{LP}	$V_{top10 \cap PM}$	$V_{top10 \cap LP}$	$V_{LP \cap PM}$	$V_{\notin top10,PM,LP}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Interbank	4.3478***	0.1984***	0.3286**	0.8337***	0.4106*	-0.0001	3.2545***
	(0.7699)	(0.0413)	(0.1309)	(0.1027)	(0.2480)	(0.0006)	(0.7452)
Corporate-Bank	1.5545	-0.0026	0.3532	0.5860*	-0.4185**	-0.0018	2.2923
	(1.5778)	(0.1899)	(0.3098)	(0.3324)	(0.1945)	(0.0013)	(1.9787)
Fund-Bank	1.1120***	0.0353	0.0166	0.2303***	0.0369	0.0017	0.9016***
	(0.3931)	(0.0308)	(0.0419)	(0.0651)	(0.0844)	(0.0017)	(0.3087)
Non-Bank Financial-Bank	2.3239	0.3554	-0.0312	0.7064	0.0518	-0.0002	6.8670
	(3.7332)	(0.3012)	(0.1768)	(0.6986)	(0.1026)	(0.0002)	(7.7577)
constant	3261.9288***	113.7215***	190.3928**	111.9940	379.6192***	2.7742	4390.3679***
	(529.2313)	(35.7043)	(92.8885)	(68.7749)	(136.5073)	(2.3496)	(600.3514)
R-squared	0.017	0.005	0.005	0.028	0.001	0.000	0.018
No. observations	14,999	14,999	14,999	14,999	14,999	14,999	14,999

Note: This table presents the results of regressing CLS volume on DEX volume. DEX volume is measuring the aggregate buy and sell transactions in EURC, and is sourced from Uniswap V3 trade data. DEX volume is divided into sub-categories: sophisticated traders (top 10 wallets), primary dealers, and LPs, denoted by V_{top10}, V_{PM} and V_{LP} respectively. Additionally, we include DEX trading volume of the intersection of sophisticated traders and primary dealers, $V_{top10 \cap PM}$, and the intersection of sophisticated traders and LPs, $V_{top10 \cap LP}$, and traders that do not belong to the three groups, $V_{\notin top10,PM,LP}$. CLS volume is measured in EUR Millions, and is aggregated as well as in the following sub-categories: BuySide Bank-SellSide, Corporate-Bank, Fund-Bank and Non-Bank Financial-Bank volume. Total sample period is from 15 August 2022 to 30 April 2024. Standard errors are Newey-West (HAC) and reported in parentheses. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Table 5: Determinants of EURC-USDC Peg Deviations

	|EURC/USDC – EUR/USD| Peg Deviations				
	(1)	(2)	(3)	(4)	(5)
gasfee	1.3283***				1.1787**
	(0.5036)				(0.4718)
σ^{IV}_{ETH}		0.1349*			0.1582**
		(0.0803)			(0.0723)
R_{ETH}		0.0050			0.0011
		(0.0042)			(0.0031)
$\|p_{\text{USDC/USD}} - 1\|$			0.6438***		0.7532***
			(0.0931)		(0.0272)
$\|p_{\text{EURC/EUR}} - 1\|$			0.1797		
			(0.1137)		
VLOOP				-1.0602	-0.3698
				(0.8011)	(0.7163)
ICRF				32.9633	49.2961
				(111.2642)	(98.4555)
constant	15.9478***	15.0894***	20.6290***	23.9748***	5.1853
	(2.9195)	(5.3608)	(2.1497)	(1.8659)	(5.9025)
R-squared	0.0552	0.0139	0.2805	0.0019	0.2294
No. observations	625	624	429	625	624

Note: This table reports OLS regressions of daily absolute peg deviations $|p_{\text{EURC/USDC}} - p_{\text{EUR/USD}}|$. Gas fees measure average transaction costs on the Ethereum network, expressed in USD per transaction. σ^{IV}_{ETH} denotes the 30-day implied volatility for Ether from the EthVol index. R_{ETH} is the daily return on ETH-USDC based on closing prices. VLOOP is the standardized first principal component of FX arbitrage violations (triangular no-arbitrage deviations) across nine G10 currency pairs. ICRF is the intermediary capital risk factor, from He et al. (2017), reflects shocks to U.S. dealer capital constraints based on equity-to-asset ratios. $|p_{\text{USDC/USD}} - 1|$ and $|p_{\text{EURC/EUR}} - 1|$ are peg deviations of the USDC and EURC stablecoins from their respective fiat reference values. All data is measured at the daily frequency, and peg-price deviations, returns, and volatility measures are expressed in basis points. Standard errors are Newey-West (HAC) and reported in parentheses. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Table 6: Determinants of EURC/USDC Order Flow

	$OF_{top10,t}$	$OF_{PM,t}$	$OF_{LP,t}$	$OF_{top10\cap PM,t}$	$OF_{top10\cap LP,t}$	$OF_{LP\cap PM,t}$	$OF_{\notin top10,PM,LP,t}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$P_{DEX,t-1} - P_{CLS,t-1}$	-0.1454**	-0.0097	-0.0207*	-0.1374***	-0.0032	-0.0003	-0.2247***
	(0.0609)	(0.0072)	(0.0121)	(0.0297)	(0.0085)	(0.0002)	(0.0354)
DEXReturn_{t-1}	-0.0077**	-0.0002	0.0003	-0.0012	0.0002	-0.0000	-0.0008
	(0.0035)	(0.0002)	(0.0006)	(0.0010)	(0.0002)	(0.0000)	(0.0026)
$OF_{top10,t-1}$	0.1995***						
	(0.0687)						
$OF_{PM,t-1}$		0.0257**					
		(0.0128)					
$OF_{LP,t-1}$			0.0153				
			(0.0138)				
$OF_{top10\cap PM,t-1}$				0.0654**			
				(0.0261)			
$OF_{top10\cap LP,t-1}$					-0.0888		
					(0.1617)		
$OF_{LP\cap PM,t-1}$						0.0000	
						(0.0001)	
$OF_{\notin top10,PM,LP,t-1}$							0.1332
							(0.0863)
constant	0.0001	0.0000	0.0001**	-0.0000	0.0000	-0.0000	0.0003*
	(0.0002)	(0.0000)	(0.0001)	(0.0001)	(0.0000)	(0.0000)	(0.0002)
R-squared	0.042	0.001	0.000	0.012	0.008	0.000	0.020
No. observations	14,998	14,998	14,998	14,998	14,998	14,998	14,998

Note: This table presents the results of regressing order flow on the price difference between the DEX and CLS exchange rates. OF measures net buyer transactions of EURC, sourced from Uniswap V3 data. $P_{DEX} - P_{CLS}$ measures the price difference between DEX and CLS exchange rates. Order flow is divided into sub-categories such as top 10 wallets, access to primary markets, and LPs. The sample period is from 15 August 2022 to 30 April 2024. Standard errors are Newey-West (HAC) and reported in parentheses. *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

Internet Appendix to
"Blockchain Currency Markets"

(Not for publication)

We provide a roadmap of each section of our Appendix.

1. Appendix A examines primary market issuance of USDC and EURC, describing Treasury transactions and their impact on stablecoin supply and circulation.

2. Appendix B describes the mechanics of liquidity provision in Uniswap V3, including liquidity aggregation across ticks, price setting through virtual reserves, and the construction of on-chain measures of net new liquidity from mint and burn flows aggregated to hourly frequency.

3. Appendix C reports detailed statistics on trader activity and liquidity provision. It includes transaction volumes, trading frequencies, participant classifications, distribution of liquidity across tick ranges, intraday liquidity adjustments, and JIT provision patterns.

4. Appendix D provides supplementary information on market efficiency. It details the construction of arbitrage bounds between EURC/USDC and EUR/USD markets, incorporating transaction costs, gas fees, and slippage estimates, and analyzes the reaction of exchange rates to Federal Reserve monetary announcements using an event study methodology.

5. Appendix E provides transaction-level evidence on sophisticated investors and LPs during the USDC de-pegging event (March 10–12, 2023).

6. Appendix F details the SVAR identification strategy, including recursive Cholesky assumptions and matrix construction for traditional OTC and blockchain order flows.

7. Appendix G presents analysis on the role of private information versus feedback trading in explaining permanent price impact.

8. Appendix H presents robustness tests, including intraday price impact analysis, controls for liquidity provision, and documentation of JIT liquidity provision by sophisticated LPs.

Appendix A: Primary Market Issuance

We obtain data on the primary market issuance from the Ethereum blockchain API. The primary market issuance uses a Circle Treasury address of the EURC and USDC Treasury. This dataset provides an entire history of Treasury transactions, with details on the size, timestamp, and the type of transaction. USDC tokens are created through a "grant" when new USDC tokens are minted. USDC tokens are destroyed through a "revoke" when USDC tokens are redeemed. Transactions between the Treasury and secondary market recipients are recorded based on whether counter parties are listed on the "send" and "receive" sides of the transaction.[27] The supply of USDC and EURC is shown in Figure A1. In addition to documenting the aggregate supply of USDC and EURC, we net out the amount of Circle tokens held by the Treasury that is not circulating in private wallets. This is indicated by the labels "USDC Total Circulation" and "EURC Total Circulation". The USDC primary market started issuance in early 2019, and reached a peak of nearly 60 USDC Billion in 2022. In contrast, the EURC Issuance started in June 2022 and reached a peak of 75 EURC Million.[28]

An important function of the USDC and EURC Treasury is guaranteeing a primary market rate, which is the rate at which the Treasury is willing to exchange USDC for dollars. The primary market rate is 1 USDC:USD for the Circle USDC Treasury, and 1 EURC:EUR for the Circle EURC Treasury. Trading of USDC/USD and EURC/EUR are on select centralized exchanges, that we can use to construct measures of market efficiency in the following subsection. Stability of the USDC and EURC pegs are based on a decentralized arbitrage mechanism (Lyons and Viswanath-Natraj, 2023; Ma et al., 2025). If the secondary market price of USDC (EURC) trades above one dollar, an investor can buy USDC (EURC) from the Treasury at a one-for-one rate, and sell USDC (EURC) at the prevailing market rate to profit, resulting in a flow of USDC (EURC) from the Treasury to the secondary market.

[27] The USDC Treasury address we use to retrieve the transaction history is "0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48". The EURC Treasury address is "0x1abaea1f7c830bd89acc67ec4af516284b1bc33c"

[28] One caveat regarding the primary market issuance data is that we can only download activities related to the transfer of ERC-20 tokens. As a result, we might miss certain transaction activities, such as internal transactions. However, our data is representative and valid for understanding the overall trend in primary market issuance.

Figure A1: Primary Market Issuance

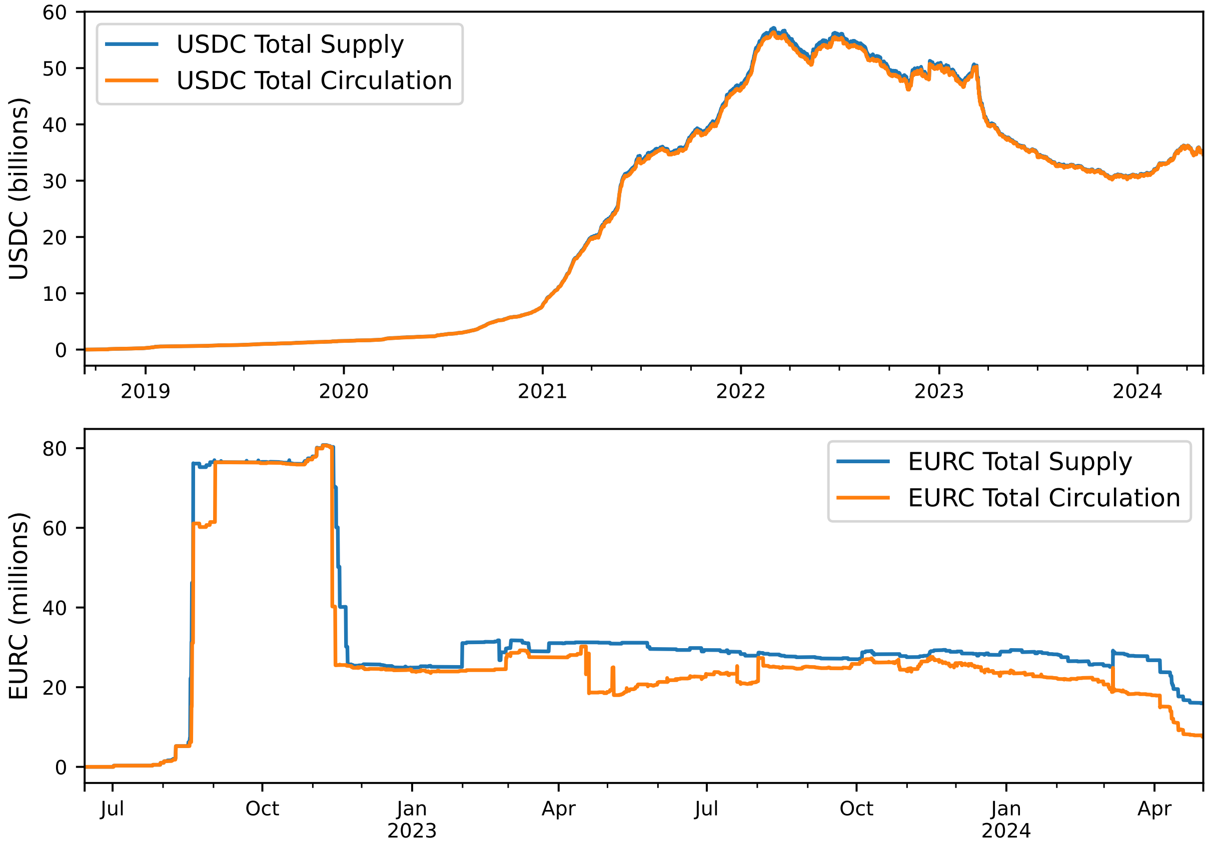

Note: This figure plots the total supply of USDC and EURC, as well as the total in circulation (net of Treasury). The top panel reports the total supply of USDC, and the bottom panel reports the total supply of EURC. The total sample period for the top two figures is from 28 June 2022, to 30 April 2024. For the bottom two figures, the sample period goes back to the early issuance dates of USDC and EURC. We use data starting from 10 September 2018, for USDC and from 23 June 2020, for EURC.

Appendix B: Liquidity Provision: Supplementary Details

B.1 Liquidity Aggregation and Execution

Uniswap V3 aggregates liquidity across price intervals by summing the liquidity parameters L_i from all LP positions that overlap the current price. Each LP contributes a liquidity parameter L_i based on the token amounts deposited and the selected price range $[P_a^{(i)}, P_b^{(i)}]$. When the market price P lies within this range, the position is active and contributes its full L_i to market liquidity. The total active liquidity at price P is therefore

$$L_{\text{active}}(P) = \sum_{i:\, P_a^{(i)} \leq P \leq P_b^{(i)}} L_i.$$

This aggregate parameter governs pricing within the current tick. If a trade moves the price beyond a tick boundary, the protocol updates L_{active} by adding or removing positions that start or end at that boundary and continues pricing in the next tick.

The relationship between each LP's liquidity parameter L_i and the token amounts deposited within a tick follows Adams et al. (2021). For a tick range $[i, i + l]$,

$$L_{EURC,i} = \frac{L_i}{\sqrt{z_i}} - \frac{L_i}{\sqrt{p_{i+l}}}, \tag{11}$$

$$L_{USDC,i} = L_i(\sqrt{z_i} - \sqrt{p_i}), \tag{12}$$

where $L_{EURC,i}$ and $L_{USDC,i}$ are the quantities of EURC and USDC locked within the range. The intermediate price z_i depends on the current market price p_M and is defined as

$$z_i = \begin{cases} p_i & \text{if } p_M \leq p_i, \\ p_M & \text{if } p_i < p_M < p_{i+l}, \\ p_{i+l} & \text{if } p_{i+l} \leq p_M. \end{cases}$$

These expressions are linear in L_i, so the token amounts within a tick and the overall active liquidity both sum directly across LPs. The aggregate $L_{\text{active}}(P)$ therefore enters the same constant-product invariant that determines prices at the pool level.

This aggregation mechanism produces a liquidity distribution that is dense near the prevailing market price and thinner in the tails, in contrast to the uniform profile in Uniswap V2. The result resembles a continuous limit order book, where liquidity above the current price represents offers to sell EURC for USDC and liquidity below represents bids to buy EURC with USDC. Because LPs can choose arbitrary tick intervals, the distribution reflects their expectations about where future trading is most likely to occur. LPs can also supply liquidity asymmetrically. If the chosen range lies entirely above the current market price, only EURC is deposited, similar to a sell limit order. If the range lies entirely below, only USDC is deposited, similar to a buy limit order. When the range straddles the current price, both tokens are deposited and the position provides liquidity on both sides of the market.

B.2 Price Setting in Uniswap V3

The aggregate liquidity parameter $L_{\text{active}}(P)$ determines the slope of the bonding curve that governs trading within the active tick. Prices are computed using *virtual reserves* rather than actual token reserves. Let the price of EURC in USDC be

$$P_{x/y} \equiv P = \frac{L_{USDC,v}}{L_{EURC,v}}, \tag{13}$$

where $L_{EURC,v}$ and $L_{USDC,v}$ denote the virtual reserves of EURC and USDC. These reserves are determined by the liquidity parameter L as

$$L_{EURC,v} = \frac{L}{\sqrt{P}}, \qquad L_{USDC,v} = L\sqrt{P}, \qquad L_{EURC,v}L_{USDC,v} = L^2. \tag{14}$$

For an LP that provides liquidity over a price interval $[P_a, P_b]$, the liquidity parameter relates to the actual reserves through

$$L = \frac{x + y}{\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right) + \left(\sqrt{P} - \sqrt{P_a}\right)}. \tag{15}$$

A buy of Δx EURC, which removes EURC and adds USDC, satisfies

$$\left(L_{EURC,v} - \Delta x\right)\left(L_{USDC,v} + \Delta y\right) = L^2, \tag{16}$$

and the post-trade price can be written as

$$P' = \frac{L_{USDC,v} + \Delta y}{L_{EURC,v} - \Delta x} = \frac{L^2}{(L_{EURC,v} - \Delta x)^2}. \tag{17}$$

For small trades, slippage is approximately proportional to $\Delta x / L_{EURC,v}$, so a larger L implies smaller price impact near the current price.

Virtual and Actual Reserves. At price P, the actual reserves associated with a liquidity interval $[P_a, P_b]$ are

$$x = L\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right), \qquad y = L\left(\sqrt{P} - \sqrt{P_a}\right). \tag{18}$$

By contrast, the virtual reserves $(L_{EURC,v}, L_{USDC,v})$ depend only on (L, P) and not on (P_a, P_b). They represent the notional reserves that reproduce the constant-product pricing curve at P. The virtual and actual reserves coincide only in the limiting case $[P_a, P_b] = [0, \infty)$, corresponding to Uniswap V2. For any finite interval, the virtual reserves exceed the actual reserves, providing greater effective depth for the same token balances.

Numerical example. Consider two LP positions with the same token budget $x+y = 210$, matching the small-pool example in V2. The initial price is $P = 1.10$. We examine the price impact of a buy order of $\Delta x = 5$ EURC.

Narrow range $[1.0, 1.2]$. At $P = 1.10$, we have $\sqrt{P} = 1.049$ and $1/\sqrt{P} = 0.953$. With $\sqrt{P_a} = 1$ and $1/\sqrt{P_b} = 1/\sqrt{1.2} = 0.913$, the liquidity parameter is

$$L_{\text{narrow}} = \frac{210}{\left(\frac{1}{\sqrt{1.10}} - \frac{1}{\sqrt{1.2}}\right) + \left(\sqrt{1.10} - \sqrt{1.0}\right)} = \frac{210}{(0.953 - 0.913) + (1.049 - 1.000)} = \frac{210}{0.089} \approx 2{,}349. \tag{19}$$

The virtual reserves at the initial price are

$$L_{EURC,v} = \frac{2{,}349}{1.049} \approx 2{,}240, \qquad L_{USDC,v} = 2{,}349 \times 1.049 \approx 2{,}464.$$

Using the price update formula,

$$P'_{\text{narrow}} = \frac{2{,}349^2}{(2{,}240 - 5)^2} \approx 1.105,$$

which implies slippage of $(1.105 - 1.10)/1.10 \approx 0.45\%$.

Wide range $[0.6, 1.6]$. Applying the same steps gives

$$L_{\text{wide}} = \frac{210}{\left(\frac{1}{\sqrt{1.10}} - \frac{1}{\sqrt{1.6}}\right) + \left(\sqrt{1.10} - \sqrt{0.6}\right)} = \frac{210}{(0.953 - 0.791) + (1.049 - 0.775)} = \frac{210}{0.436} \approx 480.4,$$

so that $L_{EURC,v} \approx 458.1$ and $L_{USDC,v} \approx 503.9$. The same trade yields

$$P'_{\text{wide}} = \frac{480.4^2}{(458.1 - 5)^2} \approx 1.124, \qquad \text{slippage} \approx 2.22\%.$$

For comparison, in the V2 pool with $(x, y) = (100, 110)$, virtual and actual reserves coincide. A trade of 5 EURC gives $P'_{\text{V2}} \approx 1.219$, with slippage of 10.8%.

This example shows that concentrating liquidity in a narrow range increases L and the virtual reserves relative to actual balances, thereby reducing slippage for a given liquidity budget.

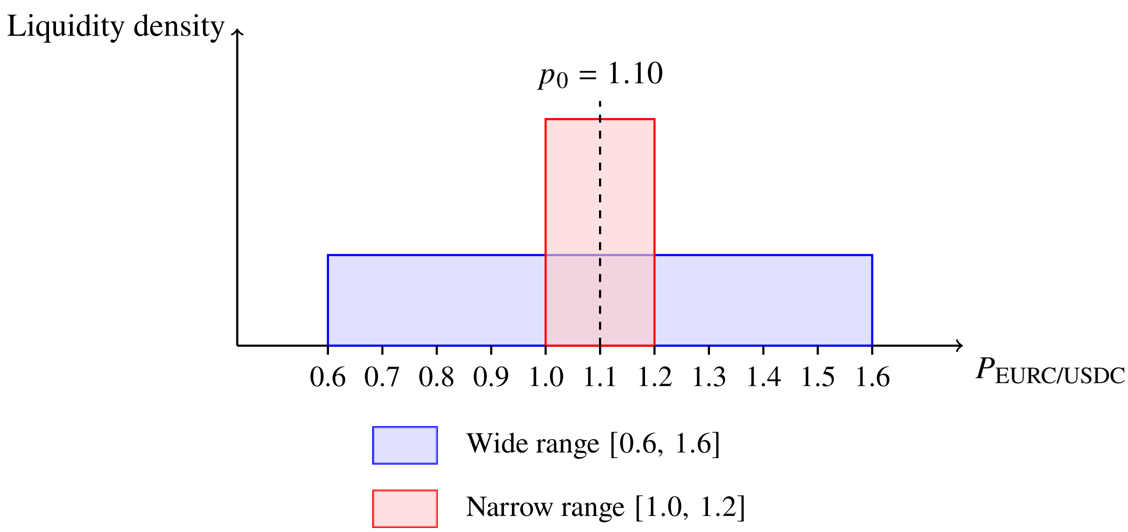

Figure A2: Liquidity distribution along the price axis. The narrow range concentrates liquidity around $p_0 = 1.10$, increasing L and lowering price impact for the same trade size.

B.3 Construction of Net Liquidity Variables

We follow Klein et al. (2024) in constructing on-chain measures of net liquidity from Uniswap V3 “mint” and “burn” events, which record additions and withdrawals of liquidity over price intervals $[P_a, P_b]$ at the block level. Because LPs can concentrate liquidity within selected price ranges, these measures can take positive or negative values. A positive value indicates a disproportionate addition or withdrawal of liquidity on the ask side, analogous to an excess of sell-limit orders in a traditional order book.

Net Mints and Burns. For each block k, we define the net mints of liquidity on the ask and bid sides as

$$mint_{(k)} = mint_{(k)}^{ask} - mint_{(k)}^{bid}, \tag{20}$$

where

$$mint_{(k)}^{ask} = \Delta x \cdot P_{x/y}, \tag{21}$$

$$mint_{(k)}^{bid} = \Delta y. \tag{22}$$

Here, Δx and Δy denote the token amounts added to the pool during the mint event, representing changes in the aggregate reserves of each token. $P_{x/y}$ is the contemporaneous market price of EURC in USDC. When multiple events occur within the same block, values are summed before computing $mint_{(k)}$.

Analogously, the net burns are defined as

$$burn_{(k)} = burn_{(k)}^{ask} - burn_{(k)}^{bid}, \tag{23}$$

with

$$burn_{(k)}^{ask} = \Delta x \cdot P_{x/y}, \tag{24}$$

$$burn_{(k)}^{bid} = \Delta y, \tag{25}$$

where Δx and Δy represent the token amounts withdrawn from the pool. A positive $mint_{(k)}$ or $burn_{(k)}$ therefore indicates that more ask-side (EURC-side) liquidity was added or withdrawn relative to bid-side (USDC-side) liquidity.

Distance from the Current Price. To account for proximity to the market price, we classify each liquidity event according to its distance from the current tick i_M. For every mint or burn defined over a tick range $[i_a, i_b]$, we compute the absolute distances between the lower and upper ticks (i_a, i_b) and the current tick i_M. Positions within 100 basis points of the current price are classified as *best*, while those beyond this threshold are labeled *away*. The notation follows: the superscript

b refers to best, a to away, and the subscripts ask and bid indicate whether the position lies above or below the current market price, respectively.

Some liquidity positions span both best and away regions, so we disaggregate the total liquidity accordingly. For each mint event over a price range $[P_a, P_b]$, we first compute the total position liquidity for each provider i using their reserves (x_i, y_i) and the corresponding price interval:

$$L_i = \frac{x_i + y_i}{\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right) + \left(\sqrt{P} - \sqrt{P_a}\right)}. \tag{26}$$

This parameter L_i measures the effective liquidity contributed by LP i across their full price range. When the interval $[P_a, P_b]$ overlaps both best and away regions, L_i is decomposed proportionally using the Uniswap V3 price–tick mapping, so that part of each LP's position is attributed to each region.[29]

After this decomposition, we aggregate across all LPs active in block k to obtain total additions and withdrawals within each range:

$$mint^{b,ask}_{(k)} = \sum_i \Delta x_i^{(b)} P_{x/y}, \qquad mint^{a,ask}_{(k)} = \sum_i \Delta x_i^{(a)} P_{x/y}, \tag{27}$$

and similarly for the bid side and for burn events. These aggregated measures of best and away liquidity form the building blocks for the block-level net liquidity changes and, subsequently, the hourly liquidity variable $Liquidity^{net}_{t,h}$ used in the empirical analysis.

Net Liquidity Flow. Combining mints and burns, the block-level net liquidity flow is defined as

$$Liquidity^{net}_{(k)} = mint_{(k)} - burn_{(k)}. \tag{28}$$

A positive $Liquidity^{net}_{(k)}$ indicates that, on balance, more EURC-side liquidity has been added than withdrawn in block k.

[29] For example, if the current price is $P_M = 1.10$ and the best range is defined as ±100 basis points $[1.089, 1.111]$, an LP providing liquidity over $[1.08, 1.12]$ contributes partially to both regions. The decomposition follows from the relationship between liquidity and token quantities:

$$\Delta x_i^{(b)} = L_i\left(\frac{1}{\sqrt{P_M}} - \frac{1}{\sqrt{1.111}}\right), \qquad \Delta x_i^{(a)} = L_i\left(\frac{1}{\sqrt{1.111}} - \frac{1}{\sqrt{1.12}}\right),$$

and analogously for $\Delta y_i^{(b)}$ and $\Delta y_i^{(a)}$. By construction, $\Delta x_i^{(b)} + \Delta x_i^{(a)} = \Delta x_i$ and $\Delta y_i^{(b)} + \Delta y_i^{(a)} = \Delta y_i$. This ensures that each LP's liquidity is consistently partitioned between best and away ranges before aggregation across LPs.

For empirical analysis, we aggregate these block-level flows to the hourly frequency by summing across all blocks within each hour:

$$Liquidity^{net}_{t,h} = \sum_{k \in h} \big(mint_{(k)} - burn_{(k)} \big), \tag{29}$$

where h indexes the hourly interval within trading day t. The resulting series $Liquidity^{net}_{t,h}$ captures the net change in active liquidity per hour and serves as our main proxy for new liquidity provision in the empirical analysis.

Appendix C: Blockchain Data and Trader/LP Statistics

This section summarizes the blockchain data used in the analysis and outlines the appendices that provide detailed information on transaction variables, trading activity, and liquidity provision.

Glossary and Example Transactions (Appendix C.1). We begin with a glossary of key variables, illustrated through an example swap and liquidity (mint) transaction from the EURC/USDC pool on Uniswap V3. The examples show how execution prices, gas costs, and liquidity ranges are recorded on-chain and how pool state variables such as `sqrtPriceX96` and tick indices are used to infer market prices. They clarify the definitions of execution price, pool price, and slippage and demonstrate how blockchain data allow precise reconstruction of trade and liquidity events.

Trading Statistics, Persistence, and Private Transactions (Appendix C.2). This section reports statistics on trading activity and wallet characteristics. Around 200 addresses trade each month, compared with roughly 5–10 LPs active in minting or burning tokens. Monthly trading volume peaked at 39 million EURC in November 2022, while liquidity provision reached 13 million EURC in October 2022. Sophisticated traders account for 50–60 percent of aggregate volume, and the top five LPs provide over 90 percent of liquidity. Sophisticated traders have older addresses, higher transaction frequencies, and greater token diversity, while LPs and residual wallets are more passive. Private transactions that bypass the public mempool are concentrated among Top 10 and Top 10∩LP wallets and are typically larger and confirmed earlier within blocks, indicating execution priority and potential informational advantages.

Liquidity Provision and Inventory Management (Appendix C.3). The final section focuses on liquidity provision, reporting the number of LPs, aggregate liquidity, and its concentration across providers, and examining intra-day patterns in mint and burn activity. Liquidity peaks during core FX trading hours, consistent with overlap between on-chain and traditional market activity, but shows no systematic timing in net provision. The appendix also documents inventory management by large LPs: when order flow reduces EURC holdings, they withdraw USDC, purchase EURC through swaps, and re-mint liquidity at updated price ranges. These adjustments restore balanced inventories and indicate that LP activity reflects passive liquidity management rather than informed trading.

C.1 Glossary: Example Swap and Liquidity Transaction

C.1.1 Swap Transaction

Table A1: Example Swap Transaction and Variable Glossary: EURC/USDC on Uniswap V3

Variable	Example Value	Definition
`UTC_time`	28/06/2022 16:06	Human-readable timestamp of the transaction in UTC.
`transaction.id`	0x28f1554a0ad5974e6d252545440e0092d503be974457fcaaf5b1c17fc6e7531f	Unique transaction hash on Ethereum.
`timestamp`	1656432368	Unix timestamp (seconds since 1970-01-01).
`sender`	0x68b3465833fb72a70ecdf485e0e4c7bd8665fc45	Address initiating the transaction (taker).
`recipient`	0x73a5dba52df247a66798575f4e2bb3747f8c16d3	Address receiving the output tokens.
`amount0` (EURC)	–20	Negative means trader bought EURC from the pool; positive means sold to the pool.
`amount1` (USDC)	21.078359	Negative means trader bought USDC; positive means sold to the pool.
`amountUSD`	21.078359	USD notional value of the trade.
`sqrtPriceX96`	8.13652×10^{28}	Square root of price in Q96 format at execution.
`tick`	532	Tick index corresponding to pool price.
`Price`	1.054674	Derived from `sqrtPriceX96`.
`pool.id`	0x95dbb3c7546f22bce375900abfdd64a4e5bd73d6	Uniswap V3 pool address.
`pool.feeTier`	500	Pool fee tier in basis points (500 bps = 0.05%).
`token0`, `token1`	EURC (6 decimals), USDC (6 decimals)	ERC-20 token pair in the pool.
`transaction.blockNumber`	15040496	Ethereum block number.
`transaction.gasUsed`	323,938	Gas units used for execution.
`transaction.gasPrice`	64,993,467,673	Gas price in wei.

Note: `amount0` and `amount1` are reported from the pool's perspective. A negative EURC value means the trader buys EURC and removes it from the pool, increasing the USDC balance; a positive value means the trader sells EURC to the pool. Price is derived from `sqrtPriceX96` using $P = \left(\frac{\texttt{sqrtPriceX96}}{2^{96}}\right)^2$ (for equal decimals), which matches the tick-based price $P = 1.0001^{\text{tick}}$.

C.1.2 Liquidity Provision Transaction

Table A2: Example Mint Transaction and Variable Glossary: EURC/USDC on Uniswap V3

Variable	Example Value	Definition
`UTC_time`	28/06/2022 14:15	Human-readable timestamp of the transaction in UTC.
`transaction.hash`	0x912e3d8411e5d21f53503caed74e6922f18bfa1cae00a1f88f5dc3f203b01583	Unique transaction hash on Ethereum.
`block_number`	15040055	Ethereum block in which the transaction was confirmed.
`pool_address`	0x95dbb3c7546f22bce375900abfdd64a4e5bd73d6	Uniswap V3 pool address where liquidity was posted.
`user_address`	0x48b516f12d44fd0f8ce6f634813f078514ee8b6a	Liquidity provider wallet address.
`price`	1.052110	Spot price at the time of minting, implied by `sqrtPriceX96` and tick.
`symbol1`, `symbol2`	EURC, USDC	ERC-20 token symbols for token0 and token1.
`amount1` (EURC)	100	Token0 amount deposited into the pool.
`amount2` (USDC)	110.368749	Token1 amount deposited into the pool.
`lower_ticker`	380	Lower tick bound of the liquidity position.
`upper_ticker`	630	Upper tick bound of the liquidity position.

Note: Mint transactions add liquidity to the Uniswap V3 pool within a specified price interval, defined by the lower and upper ticks. The LP deposits both tokens (EURC and USDC) according to the bonding curve at the prevailing price. The price is derived from `sqrtPriceX96` via $P = \left(\frac{\texttt{sqrtPriceX96}}{2^{96}}\right)^2$, or equivalently from the tick as $P = 1.0001^{\text{tick}}$ for equal decimals. The relative composition of EURC and USDC depends on the position of the spot price within the tick range.

C.2 Trading Statistics

C.2.1 Trading Volume

Figure A3: Summary statistics of trading volume

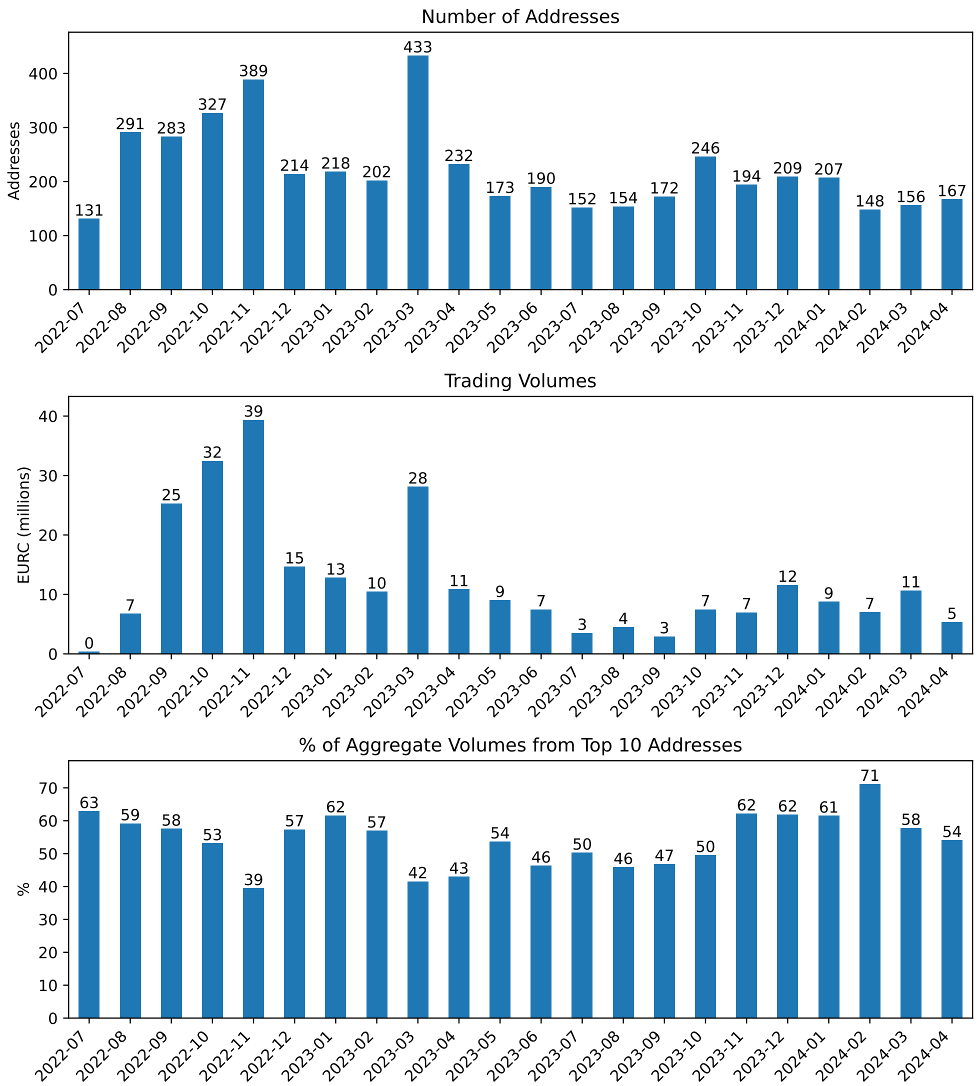

Note: This figure plots monthly summary statistics of the distribution of trading volume. It shows the number of addresses, the trading volume, and the percentage of trading volume from sophisticated traders (top 10 wallets). The total sample period is from 1 July 2022 to 30 April 2024.

C.2.2 Blockchain characteristics

In this section, we provide supplementary information on blockchain-level characteristics of trader wallets, specifically wallet age, the number of tokens transferred, and the average frequency of transactions per day. We examine these characteristics across different trader types: sophisticated traders, primary dealers (PMs), and LPs. Table A3 presents summary statistics across seven mutually exclusive groups.

Panel (a) reports statistics for sophisticated traders. These wallets are relatively older, with a median age of 742 days, and show the highest transaction frequency, with a median of 0.68 transactions per day. They also tend to transfer a larger number of tokens, with a median of 54.

Panel (b) shows that primary dealers have somewhat younger wallets (median age 613 days) and lower token transfer activity (median of 20), but still transact relatively frequently, with a median frequency of 0.58 transactions per day.

Panel (c) presents LPs, who have the oldest median wallet age (813 days), moderate token transfer activity (median of 30), and a lower transaction frequency (median of 0.32 per day), suggesting more infrequent but potentially larger or passive transactions.

Panels (d) to (f) report statistics for wallets belonging to multiple categories. These subgroups show a range of behaviors, but generally have higher transaction frequency and token transfers than their single-category counterparts. For instance, sophisticated traders who are also LPs exhibit a notably high median number of tokens transferred (88) and higher transaction frequency (1.13 per day).

Finally, Panel (g) covers wallets that do not fall into any of the above categories. These wallets have the lowest medians across most metrics, including a transaction frequency of 0.24 per day and only 14 tokens transferred, highlighting the less active behavior of the residual group.

Overall, the summary statistics show differences in activity patterns across trader types. Sophisticated traders and PMs tend to be more active and older, while LPs operate with lower frequency. However, the correlation between wallet-level blockchain activity and trader classification remains moderate, suggesting blockchain activity alone does not fully explain observed trading roles or strategies.

Table A3: Blockchain characteristics by address type

	count	mean	std	min	0.25	0.50	0.75	max
Panel (a): Sophisticated traders								
Age (days)	75	805.01	465.37	154.00	535.50	742.00	990.00	2624.00
Number of Tokens Transferred	75	100.83	105.77	5.00	15.50	54.00	184.00	383.00
Frequency (transactions per day)	75	10.29	47.62	0.01	0.07	0.68	2.20	384.94
Panel (b): Primary dealers								
Age (days)	68	750.31	495.65	15.00	412.50	612.50	942.25	2389.00
Number of Tokens Transferred	68	57.16	108.64	1.00	5.00	19.50	49.50	643.00
Frequency (transactions per day)	68	1.75	3.46	0.02	0.13	0.58	1.82	23.99
Panel (c): LPs								
Age (days)	90	911.97	428.14	194.00	601.25	813.00	1101.25	2301.00
Number of Tokens Transferred	90	44.28	46.04	2.00	14.25	29.50	55.00	258.00
Frequency (transactions per day)	90	0.56	0.96	0.02	0.16	0.32	0.56	8.00
Panel (d): Sophisticated traders and primary dealers								
Age (days)	6	507.50	152.10	376.00	394.50	475.00	546.50	781.00
Number of Tokens Transferred	6	21.00	7.92	11.00	15.25	21.00	26.00	32.00
Frequency (transactions per day)	6	3.94	3.40	0.63	1.80	2.38	6.90	8.23
Panel (e): Sophisticated traders and LPs								
Age (days)	7	630.86	370.45	341.00	373.00	421.00	781.50	1345.00
Number of Tokens Transferred	7	357.00	527.28	11.00	53.50	88.00	449.00	1395.00
Frequency (transactions per day)	7	4.51	7.67	0.12	0.36	1.13	4.16	21.32
Panel (f): LPs and primary dealers								
Age (days)	3	1337.00	1005.61	696.00	757.50	819.00	1657.50	2496.00
Number of Tokens Transferred	3	105.67	78.68	36.00	63.00	90.00	140.50	191.00
Frequency (transactions per day)	3	1.21	0.47	0.79	0.96	1.13	1.42	1.71

Panel (g): Not sophisticated traders, primary dealers and LPs

	count	mean	std	min	0.25	0.50	0.75	max
Age (days)	2316	707.10	492.28	1.00	406.75	585.50	944.50	2834.00
Number of Tokens Transferred	2316	64.64	251.83	1.00	4.00	14.00	46.00	7631.00
Frequency (transactions per day)	2316	2.28	16.88	0.00	0.07	0.24	0.92	558.01

Note: This table presents summary statistics of blockchain characteristics, based on age (days since wallet started trading), number of tokens transferred by the wallet, and frequency (measured in transactions per day). We compute summary statistics for 7 trading groups, including sophisticated traders, primary dealers, LPs, the intersection of sophisticated traders and primary dealers, the intersection of sophisticated traders and LPs, LPs and primary dealers, and traders that do not belong to the three groups. Total sample period is from 15 August 2022 to 30 April 2024.

C.2.3 Persistence of Wallet Classifications

Table A4 reports how often addresses remain in their assigned group across the sample, using thresholds for the share of months traded. LPs are highly stable: $86/90$ appear in every month of their trading history (96%), and all 90 trade in at least half of the months. Primary dealers (PMs) are similarly persistent, with $62/68$ present throughout (91%). By contrast, the Top10 group shows more entry and exit: only $17/76$ (22%) are present in all months, while $24/76$ (32%) appear in at most 30% of months. Intersection wallets remain sticky: all 3 PM∩LP addresses are observed continuously, and $3/7$ Top10∩LP addresses trade in every month. Finally, among addresses outside these groups, $2283/2342$ (98%) are present in all months. Weighted by group size, roughly 95% of addresses remain active throughout, suggesting that the classifications are generally stable, with most turnover concentrated in the Top10 (see Figure A4).

Table A4: Persistence of wallet classifications

Type	N	$N_{=100\%}$	$N_{>70\%}$	$N_{>50\%}$	$N_{>30\%}$	$N_{\leq 30\%}$
Top10	76	17	20	29	52	24
PM	68	62	62	63	68	0
LP	90	86	88	90	90	0
Top10 ∩ PM	6	0	0	3	6	0
Top10 ∩ LP	7	3	3	3	5	2
PM ∩ LP	3	3	3	3	3	0
$\notin \{Top10, PM, LP\}$	2342	2283	2307	2320	2339	3

Figure A4: Persistence of Wallet Classifications

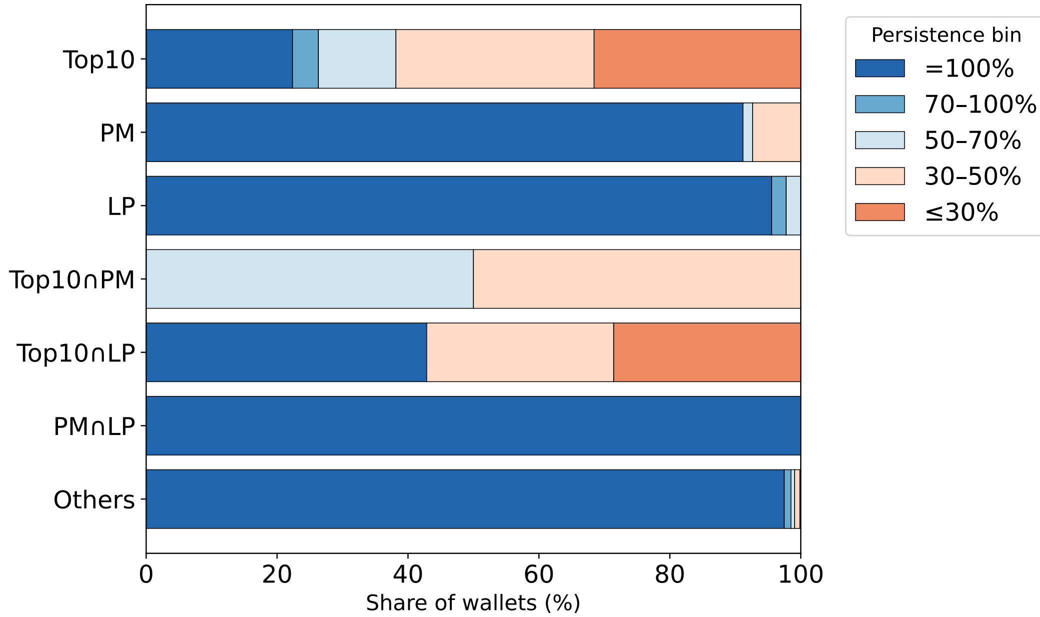

Note: This figure shows the persistence of addresses in each wallet category over the sample period. Bars depict the percentage of wallets falling into disjoint persistence bins: present in every month (= 100%), in 70–100% of months, 50–70%, 30–50%, and in at most 30% of months. Each bar sums to 100% of wallets within that group, allowing a comparison of turnover across Top10 wallets, primary dealers (PM), liquidity providers (LP), intersection groups, and addresses outside these categories.

C.2.4 Private versus public transactions

Private transactions are defined as those that do not appear in the public mempool prior to inclusion on-chain. We classify these transactions using Blocknative mempool archives,[30] based on two conditions. First, the transaction must be confirmed on-chain. Second, the `timePending` value, which measures the time in milliseconds that the transaction waited before inclusion, must be equal to zero, indicating an absence of public mempool presence. We also identify 17 transactions with no mempool record, which we classify as private. Private activity is concentrated among Top10 and Top10∩LP addresses, while PMs and LPs transact almost exclusively through the public mempool.

Table A5 reports descriptive statistics comparing private and public transactions across account types, block positions, and trade sizes.

Panel (a) shows that private activity is concentrated in a few groups. Among the Top10 wallets,

[30] See Blocknative documentation at https://docs.blocknative.com/data-archive/mempool-archive.

3,270 of 4,439 transactions (74%) are routed privately. The reliance on private submission is even higher for Top10∩LP addresses, where 220 of 249 transactions (88%) bypass the public mempool. By contrast, primary dealers (PMs) and standalone LPs rely almost exclusively on public submission, with only 3% of their trades routed privately. Addresses outside the Top10, PM, or LP categories split their activity more evenly, with 40% of trades private.

Panel (b) combines statistics on block positions and transaction volumes. Private trades enter blocks much earlier, with a median position of 7 compared to 82 for public transactions, consistent with direct relay to validators. Transaction sizes are also larger for private activity: the median private trade is about 11,100 USDC versus 6,400 USDC for public trades, and the upper tail is thicker, with maximum sizes exceeding 1 million USDC. These patterns suggest that private submission is primarily used by large Top10 wallets and Top10∩LP addresses to obtain execution priority and to handle block-sized trades, while smaller traders and PMs rely on the public mempool.

Table A5: Summary statistics of private and public transactions

Panel (a): Number of transactions by account type

	N_{Private}	N_{Public}	Total	% Private
Top10	3,270	1,169	4,439	74%
PM	10	353	363	3%
LP	14	432	446	3%
Top10 $\cap$ PM	1	533	534	0%
Top10 $\cap$ LP	220	29	249	88%
PM $\cap$ LP	0	6	6	0%
$\notin$ {Top10, PM, LP}	3,642	5,474	9,116	40%

Panel (b): Block positions and volumes per transaction (USDC equivalent)

	$\text{BlockPos}_{\text{Private}}$	$\text{BlockPos}_{\text{Public}}$	$\text{Vol/TX}_{\text{Private}}$	$\text{Vol/TX}_{\text{Public}}$
count	7,157	7,996	7,157	7,996
mean	26	88	21,270	14,039
std	47	69	42,321	22,112
min	0	0	4	0
25%	3	37	5,031	1,000
50%	7	82	11,079	6,382
75%	21	120	22,887	20,000
max	658	650	1,040,295	557,076

Note: This table reports descriptive statistics of private and public transactions. Panel (a) shows the number of transactions by account type and the share of private flows. Panel (b) combines statistics on block positions (ordering within blocks) and transaction volumes (in USDC equivalent) per trade. For each metric, separate columns are reported for private and public transactions. Sample period as in Figure X.

C.3 Liquidity Provision Statistics

C.3.1 Summary Statistics

Figure A5: Summary statistics of liquidity provision

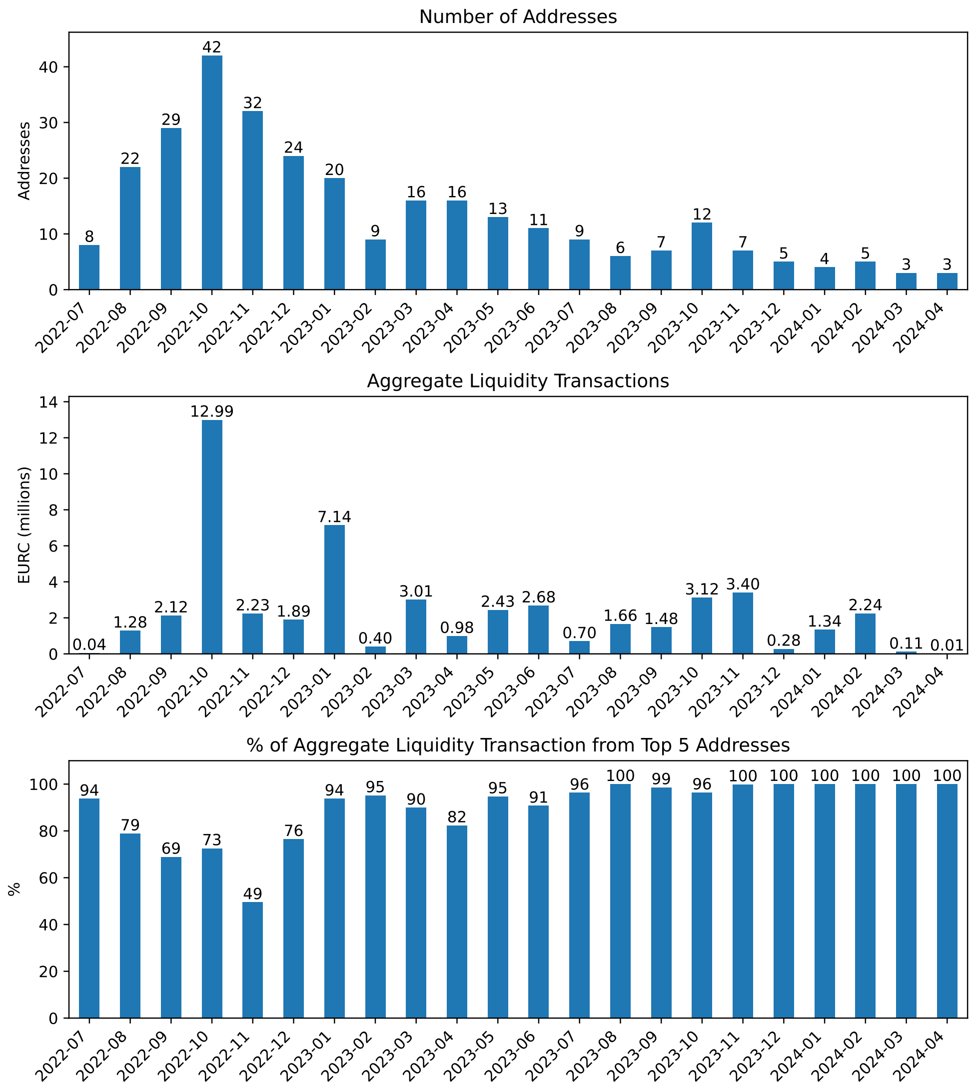

Note: This figure plots monthly summary statistics of the distribution of liquidity provision. It shows the number of addresses, the aggregate liquidity provision, and the percentage of liquidity provided by the top 5 LPs. The total sample period is from 1 July 2022 to 30 April 2024.

C.3.2 Intra-day patterns

Figure A6: Intra-day LP Mints and Burns

Panel (a): Number of transactions

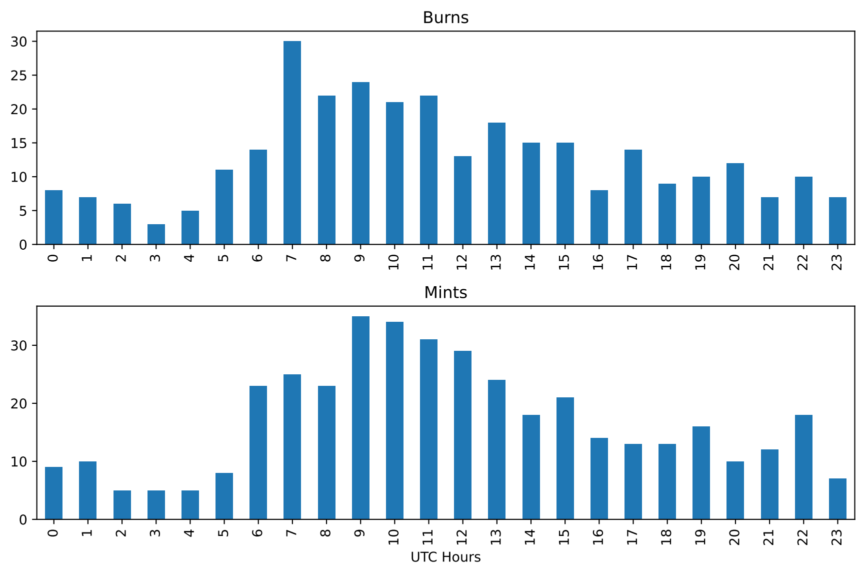

Panel (b): Volume

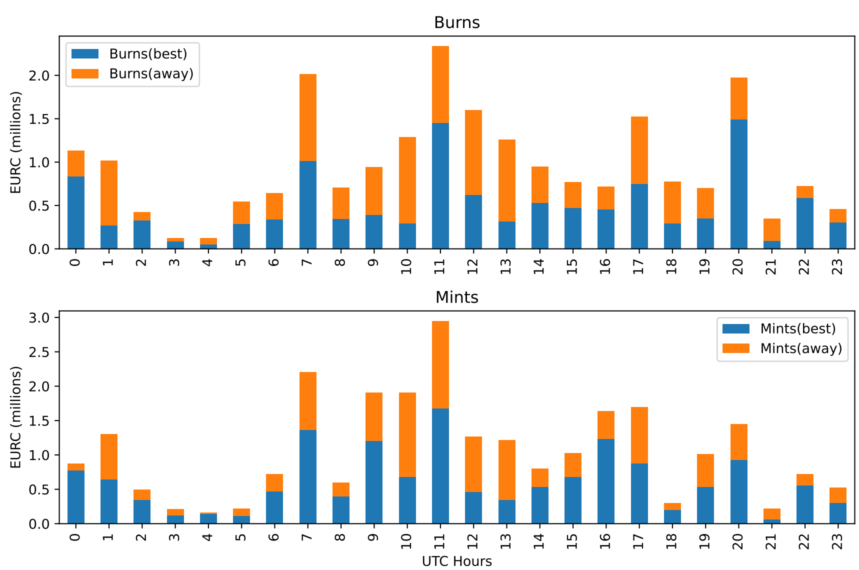

Note: Figure plots hourly liquidity provision, classified into mints (additions of liquidity) and burns (withdrawals of liquidity). Panel (a) reports LPs' transaction counts of mints and burns. Panel (b) reports LPs' mint and burn volumes, disaggregated into 'best' and 'away' regions. Liquidity is classified as 'best' when provided within ±1% of the current market price, corresponding to active liquidity at the best bid and ask. Liquidity placed outside this range is classified as 'away' and represents inactive liquidity that becomes active only after a significant price move. The sample period spans 15 August 2022 to 30 April 2024.

C.3.3 Liquidity Provision as Inventory Management

LPs often adjust their positions to stabilise inventory when order flow reduces their holdings of one token. When swap traders buy EURC from the pool, the LP's share of EURC declines. To maintain target exposure, LPs may act as swap traders themselves, by buying EURC against USDC and then redepositing it as liquidity. Table A6 illustrates these dynamics through selected transactions.

We find evidence of inventory management behaviour by large LPs during Q4 2022, a period in which the EUR/USD exchange rate appreciated from 0.98 USD to 1.07 USD. The three most active LPs, measured by swap amounts exceeding 10 000 EURC, are wallets *75e6*, *2488* and *2200*. Their transactions reveal a consistent pattern in which order flow depletes their EURC holdings, prompting them to withdraw USDC, buy EURC through a swap, and re–deposit liquidity at a refreshed price range. This reallocation restores their inventory balance and redistributes liquidity symmetrically around the market price, consistent with an inventory management motive.

Illustrative example of wallet *75e6* on 22 Nov 2022 Negative amounts indicate a decline of that currency in the pool. A negative swap amount in EURC indicates the trader is buying EURC. Mints are positive amounts of EURC and USDC into the pool, and burns are negative amounts of EURC and USDC respectively.

- Swap of −19,638 EURC at price 1.028.
- Mint of 19,638 EURC and 11,708 USDC with bounds 1.018–1.044.
- Burn of −40,701 USDC from a stale range 0.976–1.021.
- Swap of −25,765 EURC at price 1.028.
- Mint of 25,749 EURC and 17,870 USDC again at 1.018–1.044.

These actions show how *75e6* withdrew surplus USDC, purchased EURC and re–minted liquidity at updated bounds to rebalance its position. Comparable sequences for wallets *2488* and *2200*—also displayed in Table A6—follow the same logic, though with smaller trade sizes.

Table A6: Transaction Details for Selected Wallets (75e6, 2488, 2200)

Date (UTC)	Blk Num	Type	User	EURC	USDC	Lower Price	Upper Price	Price
2022-10-06 15:05		swap	75e6	-5550.000				0.984
2022-10-06 19:15		swap	75e6	-509.130				0.982
2022-10-06 19:17		mint	75e6	6059.130	3773.114	0.969	1.001	
2022-10-28 12:29		swap	75e6	-19956.539				1.005
2022-10-28 12:32		swap	75e6	500.000				0.999
2022-10-28 12:38		mint	75e6	19493.650	21007.628	0.976	1.021	
2022-11-10 10:27		burn	75e6	-1990.440	-7793.174	0.969	1.001	
2022-11-10 10:31		swap	75e6	2070.240				0.994
2022-11-22 12:14		swap	75e6	-19638.000				1.028
2022-11-22 12:15		mint	75e6	19638.000	11707.874	1.018	1.044	
2022-11-22 12:19		burn	75e6	0.000	-40700.656	0.976	1.021	
2022-11-22 12:23		swap	75e6	-25765.296				1.028
2022-11-22 12:23		mint	75e6	25749.030	17869.645	1.018	1.044	
2022-12-03 15:41		burn	75e6	0.000	-76597.072	1.018	1.044	
2022-12-03 15:44		swap	75e6	139.847				1.053
2022-11-10 19:05		swap	2488	-990.764				1.009
2022-11-10 19:39		mint	2488	14630.780	15986.258	1.009	1.011	
2022-11-11 19:09		burn	2488	0.000	-30771.748	1.009	1.011	
2022-11-11 19:14		swap	2488	-18552.669				1.036

Continued on next page

Table A6: Transaction Details for Selected Wallets (cont.)

Date (UTC)	**Blk Num**	**Type**	**User**	**EURC**	**USDC**	**Lower Price**	**Upper Price**	**Price**
2022-11-11 19:16		mint	2488	14424.180	11632.420	1.029	1.031	
2022-11-11 19:18		mint	2488	4181.240	0.000	1.029	1.031	
2022-11-14 21:26		burn	2488	0.000	-30810.961	1.029	1.031	
2022-11-14 21:47		swap	2488	94.100				1.035
2022-10-24 22:25		mint	2200	0.000	12903.813	0.952	0.986	
2022-11-05 11:20		burn	2200	0.000	-12903.813	0.952	0.986	
2022-11-05 11:54		swap	2200	-6290.000				0.996
2022-11-05 11:56		mint	2200	6295.920	6692.779	0.968	1.023	
2022-11-12 10:50		burn	2200	0.000	-13048.534	0.968	1.023	
2022-11-12 10:56		swap	2200	-4560.000				1.032
2022-11-12 10:58		mint	2200	4570.370	8404.033	0.950	1.080	
2023-01-22 10:16		burn	2200	0.000	-13228.607	0.950	1.080	
2023-01-22 10:34		swap	2200	-1361.336				1.087
2023-01-22 10:36		mint	2200	1345.640	12083.623	1.000	1.097	
2023-05-02 09:23		burn	2200	0.000	-13553.126	1.000	1.097	

Note: This table reports transactions for wallets ending *75e6* (0x33e0. . . 75e6), *2488* (0xe75b. . . 2488), and *2200* (0x3231. . . 2200) in the EURC–USDC pool. For *swap* rows, a negative amount indicates the trader has purchased EURC and is removing EURC from the pool, and positive means the trader is selling EURC; "USDC" is left blank. "mint"/"burn" rows display net liquidity changes in each token; and the price range denoted by the lower and upper price. All numeric entries are formatted to three decimal places; block numbers were unavailable. Sample: Oct 2022–May 2023.

Appendix D: Market Efficiency: Supplementary Analysis

D.1 Volume Correlations

Table A7: Correlation Matrix of Trading Volume: Liquid Trading Hours vs Off Hours

	V (top10)	V (top10 ∩ PM)	V (PM)	V ($\notin \{Top10, PM, LP\}$)	V (top10 ∩ LP)	V (LP)	V (PM ∩ LP)	Interbank	Fund-Bank	Non-Bank Financial-Bank	Corporate-Bank
Panel (a): Trading Hours (Europe + New York, 07:00–22:00 UTC)											
V (top10)	1.00	0.07	0.09	0.35	0.03	0.10	-0.00	0.10	0.01	0.01	0.02
V (top10 ∩ PM)	0.07	1.00	-0.01	0.08	0.02	0.02	-0.00	0.13	0.02	0.03	0.04
V (PM)	0.09	-0.01	1.00	0.08	0.01	0.00	-0.00	0.06	0.00	0.02	0.02
V ($\notin \{Top10, PM, LP\}$)	0.35	0.08	0.08	1.00	0.06	0.14	-0.00	0.09	0.01	0.05	0.04
V (top10 ∩ LP)	0.03	0.02	0.01	0.06	1.00	0.03	-0.00	0.02	-0.00	-0.00	-0.00
V (LP)	0.10	0.02	0.00	0.14	0.03	1.00	-0.00	0.07	0.00	0.00	0.02
V (PM ∩ LP)	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	1.00	-0.01	0.00	-0.00	-0.00
Interbank	0.10	0.13	0.06	0.09	0.02	0.07	-0.01	1.00	-0.06	-0.04	0.11
Fund-Bank	0.01	0.02	0.00	0.01	-0.00	0.00	0.00	-0.06	1.00	0.04	0.08
Non-Bank Financial-Bank	0.01	0.03	0.02	0.05	-0.00	0.00	-0.00	-0.04	0.04	1.00	0.42
Corporate-Bank	0.02	0.04	0.02	0.04	-0.00	0.02	-0.00	0.11	0.08	0.42	1.00
Panel (b): Off Hours											
V (top10)	1.00	0.09	0.24	0.52	-0.01	0.13	–	0.08	0.01	-0.00	0.06
V (top10 ∩ PM)	0.09	1.00	-0.01	0.03	-0.00	0.05	–	0.02	-0.02	-0.00	0.00
V (PM)	0.24	-0.01	1.00	0.31	0.03	0.04	–	-0.01	0.00	-0.00	0.00
V ($\notin \{Top10, PM, LP\}$)	0.52	0.03	0.31	1.00	0.08	0.19	–	0.05	0.01	-0.00	0.07
V (top10 ∩ LP)	-0.01	-0.00	0.03	0.08	1.00	0.12	–	-0.00	0.02	-0.00	-0.01
V (LP)	0.13	0.05	0.04	0.19	0.12	1.00	–	0.03	0.02	0.01	0.02
V (PM ∩ LP)	–	–	–	–	–	–	–	–	–	–	–
Interbank	0.08	0.02	-0.01	0.05	-0.00	0.03	–	1.00	0.25	0.09	0.42
Fund-Bank	0.01	-0.02	0.00	0.01	0.02	0.02	–	0.25	1.00	0.04	0.21
Non-Bank Financial-Bank	-0.00	-0.00	-0.00	-0.00	-0.00	0.01	–	0.09	0.04	1.00	0.11
Corporate-Bank	0.06	0.00	0.00	0.07	-0.01	0.02	–	0.42	0.21	0.11	1.00

Note: This table reports pairwise Pearson correlation coefficients of trading volume across decentralized and traditional FX market segments. Panel (a) shows correlations during active trading hours (Europe + New York, 07:00–22:00 UTC), while Panel (b) reports correlations during off hours. On-chain volumes correspond to decentralized EURC/USDC trading on Uniswap V3, decomposed by participant groups: top 10 wallets (V(top10)), primary dealers (V(PM)), liquidity providers (V(LP)), their intersections, and residual addresses (V($\notin \{Top10, PM, LP\}$)). Traditional volumes correspond to CLS transaction volume by counterparty type (interbank, fund, non-bank financial, and corporate). Correlations are computed at the hourly frequency on weekdays. Gray cells indicate missing or economically negligible observations. Sample period is 15 August 2022 to 30 April 2024.

D.2 Arbitrage Bounds

Finite liquidity and transaction fees create frictions that limit arbitrage, allowing deviations from efficient prices to persist and reducing the informativeness of transaction prices (Barbon and Ranaldo, 2024). One approach to quantifying these inefficiencies is through triangular arbitrage, which identifies violations of the law of one price in a closed triplet of currency pairs $X \leftrightarrow Y$, $Y \leftrightarrow Z$, and $Z \leftrightarrow X$:

$$\Delta = |1 - P_{XY} P_{YZ} P_{ZX}|\,, \tag{30}$$

where P_{AB} represents the quoted price of currency A in units of currency B. A triangular trade is profitable only if the magnitude of Δ is sufficiently large to exceed transaction costs, including liquidity fees, slippage, and validator payments on private transactions. Under efficient arbitrage, such deviations should not persist and should align with these frictions.

Extending this framework, we construct alternative efficiency measures that apply triangular arbitrage bounds to combinations of EURC/USDC DEX prices and centralized exchange rates, as defined in Equation (31).[31] We define three measures corresponding to different trading paths. The first, Δ_1, follows a cycle converting 1 EURC to USDC on DEX, to USD on Kraken, and back to EURC via the EURC/USD pair on Coinbase. The second, Δ_2, converts 1 EURC to USDC on DEX, to EUR on Coinbase via USDC/EUR, and back to EURC via EURC/EUR. The third, Δ_3, traces a four-currency path: 1 EURC → USDC → USD → EUR → EURC.

$$\begin{aligned}
\Delta_1 &= \left|1 - \frac{p_{\text{EURC/USDC}} \cdot p_{\text{USDC/USD}}}{p_{\text{EURC/USD}}}\right| \\
\Delta_2 &= \left|1 - \frac{p_{\text{EURC/USDC}} \cdot p_{\text{USDC/EUR}}}{p_{\text{EURC/EUR}}}\right| \\
\Delta_3 &= \left|1 - \frac{p_{\text{EURC/USDC}} \cdot p_{\text{USDC/USD}}}{p_{\text{EUR/USD}} \cdot p_{\text{EURC/EUR}}}\right|
\end{aligned} \tag{31}$$

These metrics quantify inefficiencies between decentralized and centralized markets and capture frictions arising from liquidity, gas fees, validator payments, and execution costs.

Panel (a) of Table A8 reports the distribution of these triangular arbitrage deviations. Median deviations lie between 0.2 and 0.3 percent, with Δ_1 and Δ_2 slightly lower on average than Δ_3. The

[31] Centralized exchanges are the only platforms with access to USD- or EUR-denominated pairs. EURC/USD and EURC/EUR are listed on Coinbase, while USDC/USD is listed on Kraken, which offers the most liquid pair for USDC/USD.

largest observed deviation reaches 8 percent, highlighting short-lived but material inefficiencies.

Panels (b) and (c) show the incidence of arbitrage-bound violations under alternative cost assumptions. In Panel (b), we account for on-chain frictions: gas fees (in ETH converted to USD), a fixed 5 basis-point liquidity-provider fee, slippage, and private fees paid to validators for authenticated private transactions.[32] When incorporating these frictions, the share of violations is 16.1 percent for Δ_1, 18.6 percent for Δ_2, and 19.1 percent for Δ_3.

Panel (c) expands the bounds to include off-chain intermediation costs such as centralized-exchange (CEX) taker fees and OTC bid–ask spreads.[33] Including these costs lowers the incidence of violations to 3.4 percent on average (3.4 percent for Δ_1, 3.1 percent for Δ_2, and 4.6 percent for Δ_3).

Although most deviations fall within estimated bounds, arbitrage spreads remain large relative to traditional FX markets, where the average hourly VLOOP across EUR–USD–X triplets is only 1.6 basis points. We therefore interpret these results as evidence of *constrained efficiency*, reflecting the inherent costs and frictions of decentralized trading and settlement (Gromb and Vayanos, 2002).

[32] For execution prices, we use the ratio of tokens exchanged (e.g., if a swap adds 10 EURC and removes 11 USDC, the execution price is 1.10 EURC/USDC). The initial price is the last completed trade before execution. Private fees are approximated from validator transfers described in Appendix C.2.4.

[33] Coinbase charges 60 bps for trades under $10,000, falling to 5 bps for institutional volumes; Kraken follows a similar schedule. See Coinbase fee schedule and Kraken fee schedule. The EUR/USD bid–ask spread of 0.55 bps follows Filippou et al. (2024).

Table A8: Triangular arbitrage conditions and transaction costs: violations of the upper bound

	count	mean	std	min	25%	50%	75%	max
Panel (a): Triangular arbitrage metrics								
Δ_1	9049	0.003	0.006	0.000	0.001	0.002	0.003	0.080
Δ_2	9049	0.004	0.007	0.000	0.001	0.002	0.004	0.071
Δ_3	9049	0.004	0.008	0.000	0.001	0.002	0.004	0.079
Δ_{UB}	9049	4.683	431.213	0.001	0.003	0.006	0.015	41017.587
Panel (b): Transaction costs: gas fees + liquidity fees + slippage + private fees								
Δ_1 Arbitrage Bound Violation	9049	0.161	0.367	0.000	0.000	0.000	0.000	1.000
Δ_2 Arbitrage Bound Violation	9049	0.176	0.381	0.000	0.000	0.000	0.000	1.000
Δ_3 Arbitrage Bound Violation	9049	0.186	0.389	0.000	0.000	0.000	0.000	1.000
Panel (c): Transaction costs: gas fees + liquidity fees + slippage + private fees + CEX fees								
Δ_1 Arbitrage Bound Violation	9049	0.034	0.182	0.000	0.000	0.000	0.000	1.000
Δ_2 Arbitrage Bound Violation	9049	0.031	0.174	0.000	0.000	0.000	0.000	1.000
Δ_3 Arbitrage Bound Violation	9049	0.046	0.210	0.000	0.000	0.000	0.000	1.000

Note: This table reports summary statistics on violations of no-arbitrage conditions based on triangular arbitrage metrics (Δ_1, Δ_2, Δ_3) constructed from EUR/USD and EURC/USDC prices. Panel (a) summarizes absolute percentage deviations and the estimated upper bound (Δ_{UB}). Panel (b) reports the share of observations where the arbitrage metric exceeds estimated on-chain transaction costs, including gas fees (ETH $\rightarrow$ USD), a 0.05 % LP fee, slippage, and private validator fees. Panel (c) adds centralized-exchange (CEX) taker fees and bid–ask spreads for traditional FX trades. Binary indicators equal 1 when deviations exceed these bounds. Gas fees and costs are winsorized at the top 1 %. Sample: 1 March 2023 – 30 April 2024.

D.3 Monetary Announcements

An efficient market should incorporate public news into prices rapidly and without systematic delay. We test this hypothesis by examining how *blockchain market prices* (EURC/USDC) respond to scheduled public information releases by the Federal Reserve. Specifically, we use high-frequency timestamps of Federal Open Market Committee (FOMC) announcements, which are released at 2pm Eastern Time, to measure the intra-day response of both the on-chain EURC/USDC price on Uniswap V3 and the off-chain EUR/USD price on CLS.

To construct the event study, we center the time series on each announcement timestamp and compute average values of key variables over a ±6 hour window around the event. We then align the same intra-day time (2pm ET) for all non-announcement weekdays in the sample to form a *placebo group*. This approach allows us to isolate the effect of scheduled public news arrivals from normal intra-day variation. Figure A7 presents the aggregate response, while Figure A8 plots the reaction for each of the 13 individual announcements between August 2022 and March 2024.

Panel (a) of Figure A7 plots the absolute price difference between the EURC/USDC and EUR/USD exchange rates. The placebo series provides a benchmark for normal intra-day pricing differences, while the announcement series captures deviations during public news events. Panel (b) reports trading volume in both markets, with DEX trading volume scaled to thousands of EURC, and CLS Bank volume scaled to millions of EUR. Trading activity exhibits a clear spike at the time of announcements in both markets, consistent with liquidity surges and price discovery around scheduled macroeconomic news releases. Importantly, the blockchain-based EURC/USDC price responds to FOMC announcements, closely tracking the off-chain EUR/USD price response.

To formally assess whether price differences during announcements are statistically distinguishable from normal intra-day variation, we perform t-tests comparing the absolute price difference between the announcement and placebo groups for each hour in the ±6 hour window. Table A9 reports the results. Price differences are significantly lower during the pre-announcement period (-6 to -1 hours), reflecting reduced dispersion before scheduled public news releases, possibly due to anticipatory arbitrage or liquidity provision. Crucially, the difference is statistically insignificant during the post-announcement period (0 to +6 hours), indicating that on-chain and off-chain prices remain closely aligned after public news is incorporated into prices.

Importantly, the absence of a significant post-announcement price difference coincides with a marked increase in trading volume at the time of announcements, as shown in Panel (b) of Figure A7. This suggests that higher on-chain trading activity, likely reflecting arbitrage or informed trading, helps align prices across the two markets when public news arrives. In contrast, on placebo (non-announcement) days, trading volume typically declines at 2pm US ET—corresponding to the late European trading session—particularly in the blockchain market. This pattern indicates that elevated trading activity during FOMC announcements plays a key role in maintaining a tight price alignment between EURC/USDC and EUR/USD across market venues.

Figure A7: Event Study Around Federal Reserve Monetary Announcements

Panel (a): Absolute Price Difference (bps)

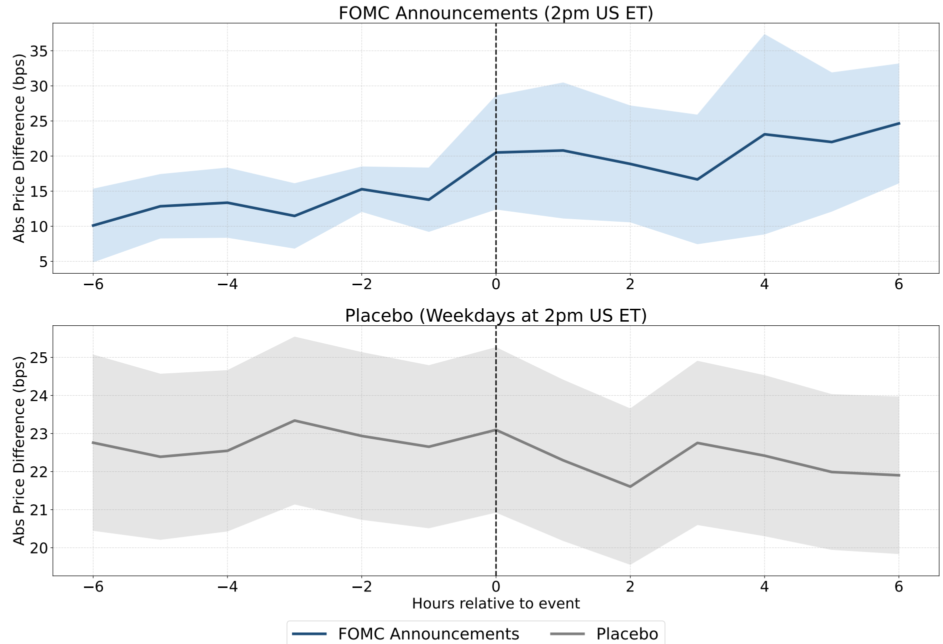

Panel (b): Trading Volume (DEX & CLS)

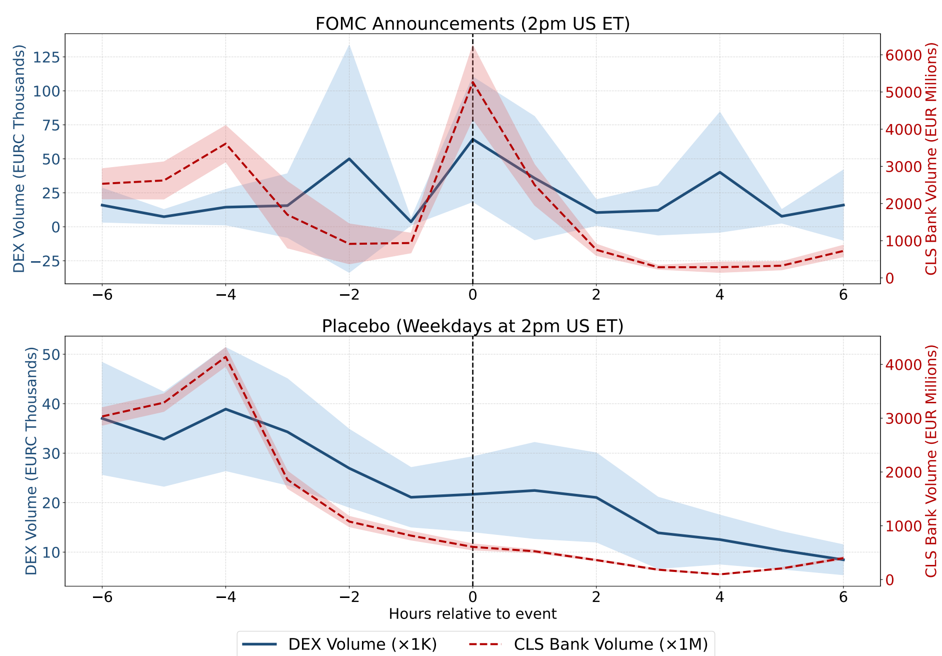

Note: This figure presents event studies around monetary announcements of the Federal Reserve (2pm US ET). Panel (a) plots the absolute price difference (in basis points) between EURC/USDC on Uniswap V3 and EUR/USD on CLS. Panel (b) plots trading volumes in both markets: DEX trading volume is reported in EURC thousands and CLS Bank volume is reported in EUR millions. The placebo group corresponds to weekdays at 2pm ET, excluding announcement days. The sample period covers 15 August 2022 to 30 April 2024.

Figure A8: Federal Reserve Monetary Announcements

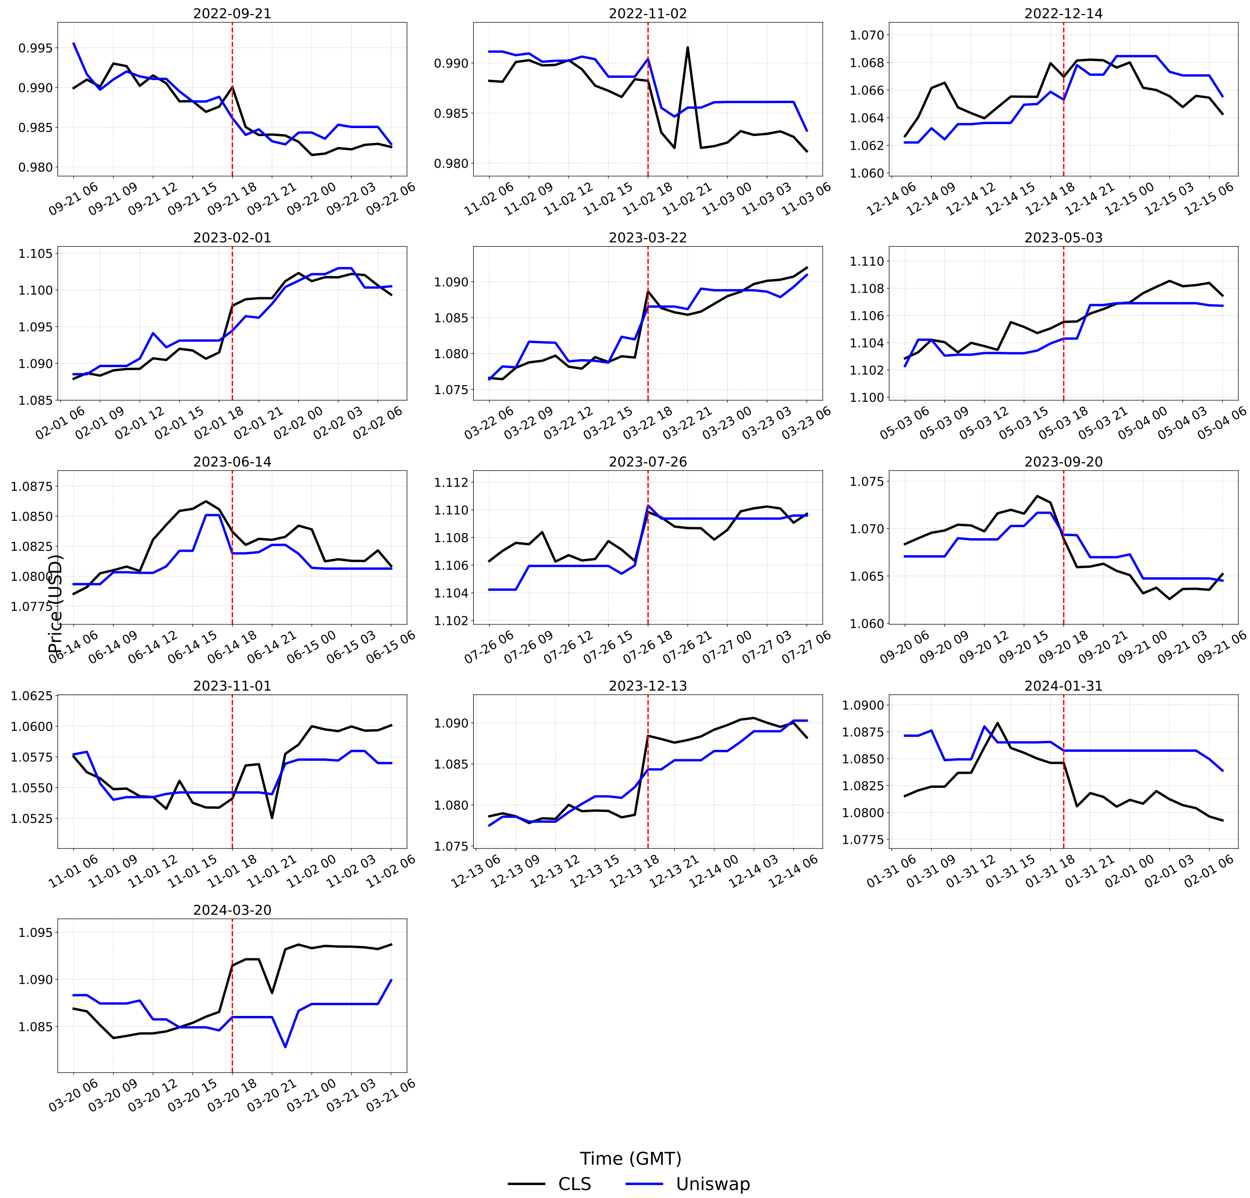

Note: This figure plots individual event studies of EURC/USDC and EUR/USD around each of the 13 monetary policy announcements by the Federal Reserve between August 2022 and March 2024. Each panel shows a 12-hour window around the 2pm ET announcement time.

Table A9: Welch T-test: Absolute Price Difference – FOMC vs Placebo

Hour (rel.)	**Mean FOMC. (bps)**	**Mean Placebo. (bps)**	**t-stat**	**p-value**
−6	10.11	22.76	−4.36	0.0004
−5	12.85	22.39	−3.71	0.0016
−4	13.36	22.55	−3.35	0.0039
−3	11.47	23.34	−4.56	0.0002
−2	15.28	22.94	−3.88	0.0006
−1	13.78	22.65	−3.47	0.0027
0	20.51	23.10	−0.61	0.5536
1	20.80	22.30	−0.30	0.7704
2	18.88	21.60	−0.63	0.5415
3	16.67	22.75	−1.26	0.2278
4	23.10	22.42	0.09	0.9269
5	22.00	21.99	0.00	0.9982
6	24.65	21.90	0.62	0.5487

Note: This table reports Welch t-test statistics comparing the absolute price difference (in basis points) between EURC/USDC and EUR/USD around FOMC announcements (treatment) and placebo weekdays at 2pm ET (control). The test is conducted for each hourly interval in a ±6 hour window around the event. Differences are statistically significant before the announcement, but not afterwards, consistent with tight price alignment across venues once public news is incorporated.

Appendix E: USDC De-Pegging Event

E.1 USDC De-Pegging Event: Sophisticated Investor Trades

Table A10: Transactions of Sophisticated Investor during USDC De-Pegging Event (2023-03-10 to 2023-03-12)

Date (UTC)	Hash	From	To	USDC	From Name	To Name
2023/03/10 00:14:47	ea98	1c37	02ce	500	trader	Uniswap V3: USDC-PRIME 2
2023/03/10 01:29:47	4daa	1c37	60ae	1500	trader	SushiSwap: SYN-USDC
2023/03/10 02:11:11	36a6	1c37	60ae	2000	trader	SushiSwap: SYN-USDC
2023/03/10 02:28:23	62e5	1c37	60ae	2000	trader	SushiSwap: SYN-USDC
2023/03/10 03:25:35	d18f	1c37	60ae	4000	trader	SushiSwap: SYN-USDC
2023/03/10 03:47:47	c30e	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 04:06:11	46f0	3e43	1c37	100000	Coinbase	trader
2023/03/10 09:49:47	65dc	1c37	1690	1000	trader	SushiSwap: DDX-USDC
2023/03/10 12:35:47	18de	1c37	1690	1500	trader	SushiSwap: DDX-USDC
2023/03/10 13:30:11	cfa9	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 13:34:59	3601	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 13:43:35	5de7	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 14:11:11	ae67	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 14:24:47	6aa6	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC

Continued...

Date (UTC)	Hash	From	To	USDC	From Name	To Name
2023/03/10 14:29:11	5102	3e43	1c37	100000	Coinbase	trader
2023/03/10 14:29:59	b043	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 14:36:59	ebaf	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 14:43:35	021d	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 15:03:47	3c82	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 15:11:23	103a	1c37	1690	1000	trader	SushiSwap: DDX-USDC
2023/03/10 15:39:35	2426	1c37	73d6	5000	trader	Uniswap V3: EURC-USDC
2023/03/10 15:55:11	e414	1c37	02ce	5000	trader	Uniswap V3: USDC-PRIME 2
2023/03/10 16:00:59	3e42	1c37	02ce	4000	trader	Uniswap V3: USDC-PRIME 2
2023/03/10 16:05:11	4a84	3e43	1c37	100000	Coinbase	trader
2023/03/10 16:05:59	2e85	1c37	02ce	4000	trader	Uniswap V3: USDC-PRIME 2
2023/03/10 18:31:11	69e9	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 19:55:47	c9e0	1c37	73d6	7500	trader	Uniswap V3: EURC-USDC
2023/03/10 21:30:35	5f29	1c37	73d6	10000	trader	Uniswap V3: EURC-USDC
2023/03/10 21:34:47	419e	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 21:40:47	8acd	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 22:26:11	54a8	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:26:23	f5f5	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:29:35	56e4	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:31:11	2f23	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN

Continued...

Date (UTC)	Hash	From	To	USDC	From Name	To Name
2023/03/10 22:31:11	3521	3e43	1c37	100000	Coinbase	trader
2023/03/10 22:33:35	0a02	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:38:47	707e	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:42:23	0393	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:43:35	2d24	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:46:23	410c	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:47:23	23ed	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:52:35	7987	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 22:55:47	e358	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 22:57:11	f239	3e43	1c37	100000	Coinbase	trader
2023/03/10 22:59:11	9bb9	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC
2023/03/10 23:09:47	f719	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 23:19:35	536f	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 23:19:47	5c76	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 23:21:35	cd37	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/10 23:24:23	a95a	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:13:23	9b13	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:13:35	1421	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:18:35	5a94	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:18:47	4725	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN

Continued...

Date (UTC)	Hash	From	To	USDC	From Name	To Name
2023/03/11 00:19:47	ab0c	1c37	2286	3000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:19:59	41e0	1c37	2286	5000	trader	Uniswap V3: USDC-GYEN
2023/03/11 00:42:11	5c32	3e43	1c37	100000	Coinbase	trader
2023/03/11 01:43:23	8e21	1c37	02ce	5000	trader	Uniswap V3: USDC-PRIME 2
2023/03/11 02:02:59	b24c	1c37	e180	4000	trader	Uniswap V3: BTRST-USDC
2023/03/11 02:40:11	aab8	1c37	b3e3	500	trader	Uniswap V3: FORT-USDC
2023/03/11 02:44:59	67b6	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 02:45:11	9ee4	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 02:45:23	b6c0	1c37	2286	10000	trader	Uniswap V3: USDC-GYEN
2023/03/11 02:52:59	e5a4	1c37	e180	3000	trader	Uniswap V3: BTRST-USDC
2023/03/11 03:08:23	465c	1c37	2286	5000	trader	Uniswap V3: USDC-GYEN
2023/03/11 03:27:35	b37a	1c37	2286	5000	trader	Uniswap V3: USDC-GYEN
2023/03/12 19:01:59	b0a1	1c37	1690	2000	trader	SushiSwap: DDX-USDC
2023/03/12 21:30:11	2d77	1c37	02ce	1500	trader	Uniswap V3: USDC-PRIME 2
2023/03/12 21:35:23	205d	1c37	02ce	1500	trader	Uniswap V3: USDC-PRIME 2
2023/03/12 23:00:35	2fc2	1c37	1690	1500	trader	SushiSwap: DDX-USDC

Note: This table presents swap transactions from the sophisticated investor with wallet ID ’0xd64137f743432392538a8f84e8e571fa09f21c37’, abbreviated to wallet ’1c37’, during the USDC de-pegging event on March 10-12, 2023. Transactions are sourced from Etherscan API. This wallet was the largest single source of USDC selling pressure during the de-pegging event. The ’From’ and ’To’ refer to transfers of USDC. Transactions typically show transfers of USDC from Coinbase to wallet ’1c37’. Wallet ’1c37’ then transfers USDC to decentralized exchange pools in Uniswap V3. The sample period is from 10 March 2023 to 12 March 2023.

E.2 USDC De-Pegging Event: Liquidity Provision

Table A11: Liquidity provision during the USDC de-pegging event

Date (UTC)	Blk Num	Type	User (last 4)	EURC	USDC	Lower Price	Upper Price	Price
2023-03-10 05:57		mint	4f35	48656.685	62725.785	1.013	1.094	1.057
2023-03-11 05:59		burn	4f35	-92233.623	-355866.065	1.013	1.094	1.076
2023-03-11 06:57		swap	4f35	-92509.174				1.071
2023-03-11 09:47		burn	4f8f	0.000	-312108.039	1.000	1.080	1.110
2023-03-11 09:51		mint	4f8f	0.000	312665.183	1.035	1.107	1.108
2023-03-12 21:32		swap	ebb3	-252.598				1.091
2023-03-12 21:34		mint	ebb3	0.000	506.468	1.005	1.075	1.091

Note: This table reports all transactions by LPs during the USDC de-pegging event (10–12 March 2023). The "Type" column identifies swaps, mints, and burns. For swaps, the amount in "EURC" equals the order flow from the pool's perspective; a negative value means the trader bought EURC from the pool. For mints and burns, EURC and USDC give the amounts added or withdrawn; negative values indicate withdrawals. "Price" denotes the swap price for swaps and the prevailing market price for mints/burns, while "Lower Price" and "Upper Price" refer to the tick range at which liquidity was posted. Wallet identifiers use the last four characters of the address: `4f35` corresponds to `0x767f840400070112ead7b6f64603897ce0144f35`, `4f8f` to `0xf550786c496bd9b99d2f91b3db6a01ce32704f8f`, and `ebb3` to `0x251691e49c2ea15882883c4ed3a4fdcd28abebb3`.

In addition to the transaction-level evidence, Figure A9 presents a series of tick-level liquidity snapshots for the EURC/USDC pool before, during, and after the de-pegging event. Because LPs are relatively passive, their liquidity distribution reflects the net effect of swap activity rather than active rebalancing. As swap traders began buying EURC, they absorbed the EURC reserves held by LPs, leading to a visible decline in the amount of EURC liquidity on the book. At the peak of the de-pegging, liquidity became highly asymmetric, with most provision concentrated in USDC on the bid side and very limited EURC remaining. This pattern indicates that LPs were left holding predominantly USDC reserves as swap traders depleted the EURC side of the pool. By March 15, when the USDC peg had stabilized, the liquidity distribution became more balanced.

Figure A9: Snapshots of EURC/USDC Liquidity Around the USDC De-Pegging Event.

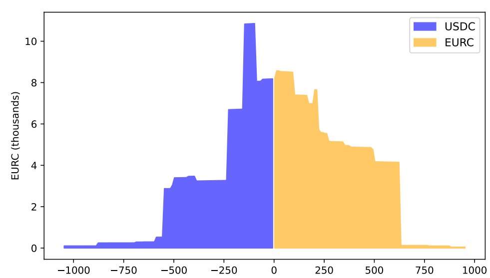

(a) 7 Mar 2023 — Pre-event

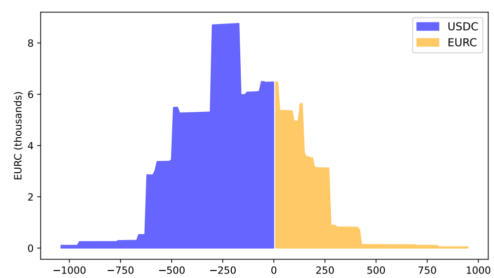

(b) 10 Mar 2023 — Onset of de-pegging

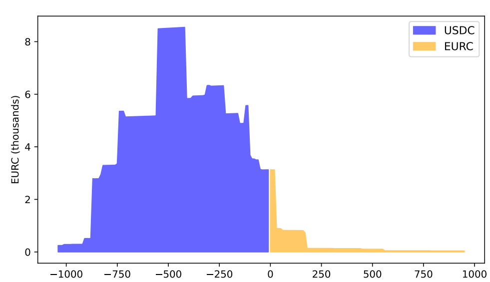

(c) 11 Mar 2023 — Peak dislocation

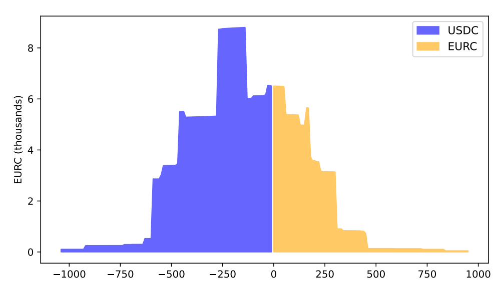

(d) 15 Mar 2023 — Post-stabilization

This figure shows the evolution of tick-level liquidity around the prevailing pool price for the EURC/USDC 0.05% Uniswap V3 pool at four points surrounding the USDC de-pegging episode. Each panel plots liquidity on both sides of the market relative to the current tick (tick 0). The horizontal axis measures tick distance in log base $\sqrt{1.0001}$ units, where the pool's fixed tick spacing of 10 corresponds to roughly 0.1% (10 basis points) price intervals. Liquidity to the left of zero represents positions below the current price (*buy-side* liquidity, LPs willing to buy EURC with USDC), while liquidity to the right represents positions above the current price (*sell-side* liquidity, LPs willing to sell EURC). Comparing panels (a)–(d) illustrates how liquidity became asymmetric during the de-pegging and rebalanced after stabilization.

E.3 Gas Fees across Trader Types: Full Sample and USDC De-Pegging Event

Table A12: Gas Fees in USDC per 10,000 EURC Transacted

Panel (a): Full Sample (2022-08-15 to 2024-04-30)

Group	Count	Mean	Std	Min	25%	50%	75%	Max
All	15,155	359,118.47	40,583,248.55	0.045	5.090	17.542	73.478	4,975,681,243.48
Top10	4,439	87.14	2,550.26	0.045	3.775	12.481	43.005	169,622.03
PM	363	325.92	2,056.57	0.300	4.933	19.560	75.545	31,146.34
LP	446	175.82	1,296.89	0.298	2.458	7.667	37.351	25,916.66
Top10 ∩ PM	534	5.10	12.42	0.388	1.594	2.981	5.317	251.50
Top10 ∩ LP	249	168.98	643.03	0.203	3.362	15.831	83.927	8,430.81
PM ∩ LP	6	42.83	66.52	4.321	6.983	11.090	40.961	173.97
$\notin \{Top10, PM, LP\}$	9,118	596,820.78	52,320,660.66	0.045	7.564	25.010	121.569	4,975,681,243.48

Panel (b): USDC De-Pegging Period (2023-03-10 to 2023-03-12)

Group	Count	Mean	Std	Min	25%	50%	75%	Max
All	299	162.47	752.35	0.611	11.119	24.928	78.663	8,665.39
Top10	54	42.87	125.03	1.355	4.826	12.439	24.949	890.29
PM	1	181.20	–	181.20	181.20	181.20	181.20	181.20
LP	3	72.95	119.29	3.948	4.076	4.203	107.45	210.70
Top10 ∩ LP	1	15.97	–	15.97	15.97	15.97	15.97	15.97
$\notin \{Top10, PM, LP\}$	240	191.04	835.42	0.611	12.547	32.820	102.360	8,665.39

Note: This table reports gas fees paid in USDC per 10,000 EURC transacted across trader types. Panel (a) reports full-sample results. Panel (b) isolates the period of the USDC de-pegging crisis from March 10–12, 2023. Groups are defined in line with Table 2, including intersections such as Top10 ∩ PM, and a residual category of traders not belonging to any primary group. Median gas fees were substantially lower for Top10 wallets relative to other groups, particularly during crisis periods.

Appendix F: SVAR Identification Assumptions

This appendix provides further detail on the identification strategy for the structural vector autoregression (SVAR) used to estimate the permanent price impact of blockchain order flow. Specifically, we describe the variable blocks, recursive ordering assumptions, and the structure of the Cholesky decomposition. We estimate a structural VAR to examine the contemporaneous and dynamic relationship between sector-level order flow and exchange rate changes. The structural form is given by

$$AY_t = A_0 + \sum_{j=1}^{L} A_j Y_{t-j} + \varepsilon_t, \tag{32}$$

where $Y_t = \left[\mathbf{OF}_t^{OTC},\ \mathbf{OF}_t^{DEX},\ \Delta p_t\right]^\top$, and ε_t is a vector of orthogonal structural shocks.

The reduced-form VAR is obtained by applying A^{-1} to both sides:

$$Y_t = C_0 + CY_{t-1} + B\varepsilon_t, \tag{33}$$

where $B = A^{-1}$, $C_0 = A^{-1}A_0$, and $C = A^{-1}A_1$. The matrix A is the Cholesky (impact) matrix and imposes the identifying restrictions in the SVAR.

We define the order flow vectors as

$$\mathbf{OF}_t^{OTC} = \left[\mathrm{OF}_{\text{non-bank}},\ \mathrm{OF}_{\text{corporate}},\ \mathrm{OF}_{\text{fund}},\ \mathrm{OF}_{\text{bank}}\right]^\top,$$

$$\mathbf{OF}_t^{DEX} = \left[\mathrm{OF}_{\text{LP}},\ \mathrm{OF}_{\notin\{\text{Top10,PM,LP}\}},\ \mathrm{OF}_{\text{Top10}\cap\text{LP}},\ \mathrm{OF}_{\text{PM}},\ \mathrm{OF}_{\text{Top10}},\ \mathrm{OF}_{\text{Top10}\cap\text{PM}}\right]^\top.$$

We impose a recursive causal structure on matrix A, which is lower triangular and decomposed into three blocks:

$$A = \begin{bmatrix} A^{OTC} \\ A^{DEX} \\ A^{\Delta p} \end{bmatrix}, \quad A^{OTC} \in \mathbb{R}^{4\times 11},\ A^{DEX} \in \mathbb{R}^{6\times 11},\ A^{\Delta p} \in \mathbb{R}^{1\times 11}.$$

1. OTC Block A^{OTC}.

$$A^{OTC} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ a_{41} & a_{42} & a_{43} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

This ordering follows Huang et al. (2023), which assumes that dealer-to-customer (D2C) flows (non-bank, corporate, and fund) affect dealer-to-dealer (D2D) flows (bank) contemporaneously, but not vice versa. This hierarchy generates a lower bound on the price impact and information share of inter-dealer trading and is consistent with inventory-based models of exchange rates in which inter-dealer markets learn from customer order flow.

2. DEX Block A^{DEX}.

$$A^{DEX} = \begin{bmatrix} a_{51} & a_{52} & a_{53} & a_{54} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} & 1 & 0 & 0 & 0 & 0 & 0 \\ a_{71} & a_{72} & a_{73} & a_{74} & a_{75} & a_{76} & 1 & 0 & 0 & 0 & 0 \\ a_{81} & a_{82} & a_{83} & a_{84} & a_{85} & a_{86} & a_{87} & 1 & 0 & 0 & 0 \\ a_{91} & a_{92} & a_{93} & a_{94} & a_{95} & a_{96} & a_{97} & a_{98} & 1 & 0 & 0 \\ a_{10,1} & a_{10,2} & a_{10,3} & a_{10,4} & a_{10,5} & a_{10,6} & a_{10,7} & a_{10,8} & a_{10,9} & 1 & 0 \end{bmatrix}$$

DEX order flows can contemporaneously respond to OTC flows. Within DEX, wallet types are ordered by size and sophistication, with LPs first and larger, more informed traders (Top10, PM) later. This ordering reflects the idea that informed traders react to aggregate flows, provides a lower bound on their information share, and is consistent with theoretical and empirical work on information transmission in DEX markets (Capponi et al., 2024b; Klein et al., 2024).

3. Price Equation $A^{\Delta p}$.

$$A^{\Delta p} = \begin{bmatrix} a_{11,1} & a_{11,2} & a_{11,3} & a_{11,4} & a_{11,5} & a_{11,6} & a_{11,7} & a_{11,8} & a_{11,9} & a_{11,10} & 1 \end{bmatrix}$$

The exchange rate is assumed to respond contemporaneously to all OTC and DEX order flows. This ordering reflects the hierarchical structure between markets and their informational structure.

Appendix G: Permanent Price Impact: the Role of Private Information

G.1 Public versus Private Transactions

Figure A10: Price impact of private versus public blockchain transactions

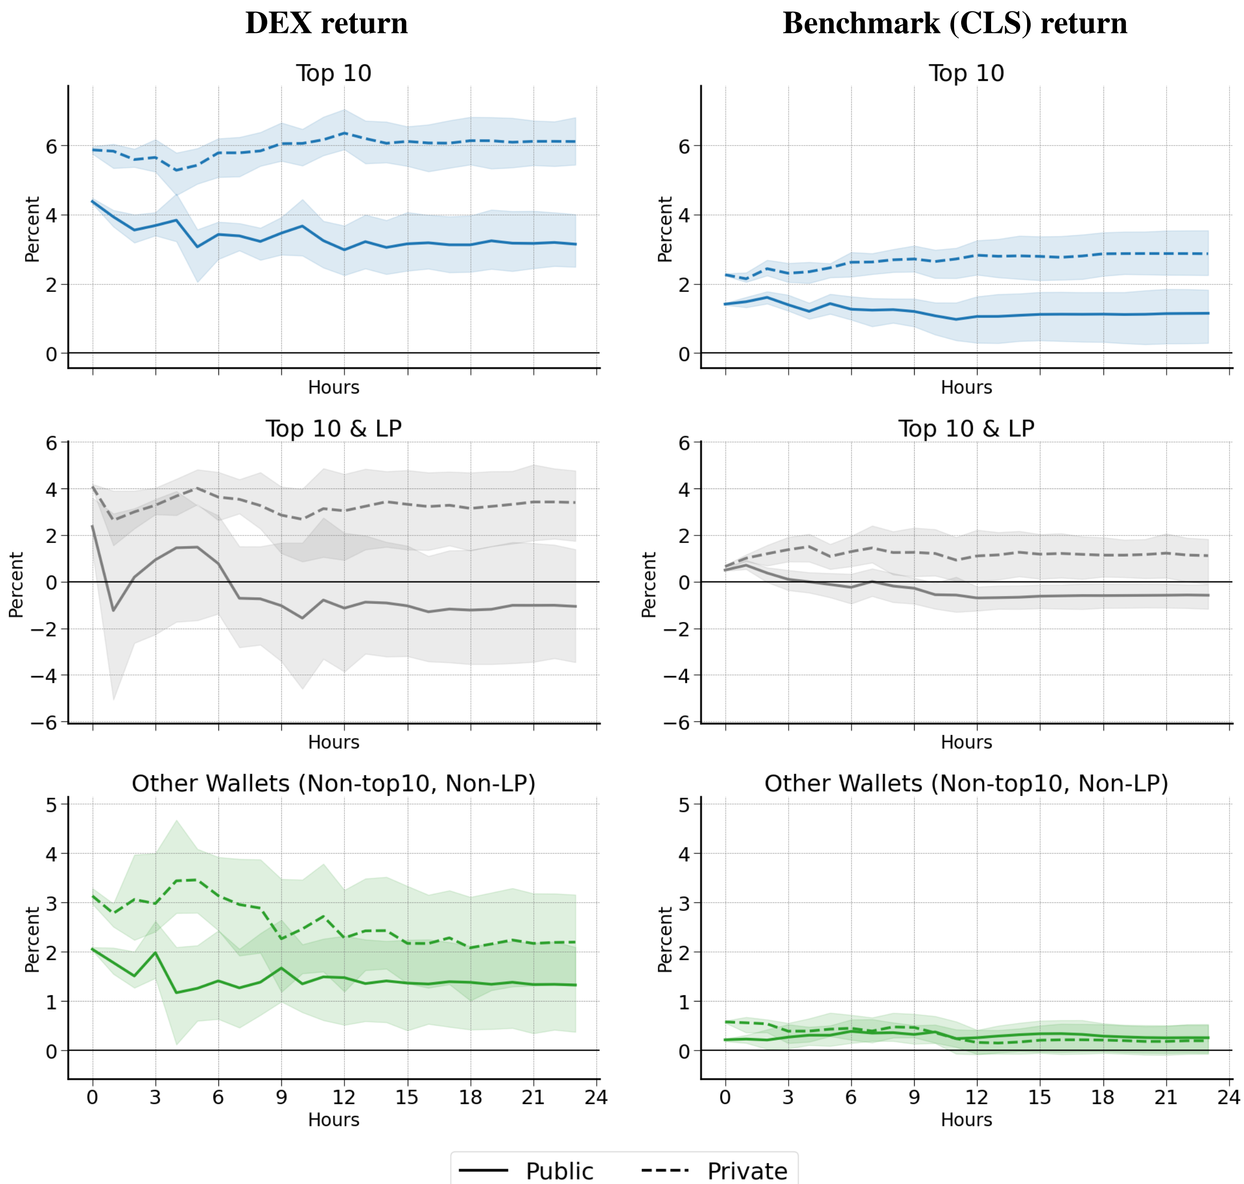

Note: This figure plots impulse responses of returns to a one million EURC shock in blockchain order flow, comparing public and private transactions across trading groups. The left column shows responses of EURC/USDC returns (Uniswap V3), and the right column shows responses of EUR/USD returns (CLS). Results are estimated using a structural VAR with 1,000 bootstrap replications. Private transactions are interpreted as those routed off-chain before execution, while public transactions occur directly on-chain. Groups include the top 10 wallets, top 10 with LPs, and all other wallets. The sample covers 15 August 2022 to 30 April 2024.

Table A13: Determinants of EURC/USDC Order Flow: Public Transactions

	OF_{top10}	OF_{PM}	OF_{LP}	$OF_{top10\cap PM}$	$OF_{top10\cap LP}$	$OF_{LP\cap PM}$	$OF_{\notin top10,PM,LP}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$P_{DEX,t-1} - P_{CLS,t-1}$	-0.1531***	-0.0145**	-0.0227*	-0.1371***	0.0034	-0.0003	-0.1936**
	(0.0334)	(0.0060)	(0.0120)	(0.0297)	(0.0079)	(0.0002)	(0.0979)
DEXReturn$_{t-1}$	-0.0009	-0.0002	0.0003	-0.0012	0.0000	-0.0000	-0.0009
	(0.0009)	(0.0002)	(0.0006)	(0.0010)	(0.0003)	(0.0000)	(0.0030)
$OF_{top10,t-1}$	0.0813***						
	(0.0233)						
$OF_{PM,t-1}$		0.0276**					
		(0.0138)					
$OF_{LP,t-1}$			0.0141				
			(0.0143)				
$OF_{top10\cap PM,t-1}$				0.0656**			
				(0.0261)			
$OF_{top10\cap LP,t-1}$					-0.1347		
					(0.2284)		
$OF_{LP\cap PM,t-1}$						0.0000	
						(0.0001)	
$OF_{\notin top10,PM,LP,t-1}$							0.2335*
							(0.1332)
constant	-0.0003***	0.0000	0.0002***	-0.0000	-0.0001**	-0.0000	0.0003
	(0.0001)	(0.0000)	(0.0001)	(0.0001)	(0.0000)	(0.0000)	(0.0002)
R-squared	0.010	0.001	0.000	0.012	0.018	0.000	0.055
No. observations	14,998	14,998	14,998	14,998	14,998	14,998	14,998

Note: This table presents the results of regressing *public* order flow on the price difference between the DEX and CLS exchange rates. OF measures net buyer transactions of EURC sourced from Uniswap V3 data. $P_{DEX} - P_{CLS}$ measures the price difference between DEX and CLS exchange rates. Order flow is divided into sub-categories such as top 10 wallets, access to primary markets, and LPs. Standard errors are Newey–West (HAC) and reported in parentheses. *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A14: Determinants of EURC/USDC Order Flow: Private Transactions

	OF_{top10}	OF_{PM}	OF_{LP}	$OF_{top10\cap PM}$	$OF_{top10\cap LP}$	$OF_{LP\cap PM}$	$OF_{\notin top10,PM,LP}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$P_{DEX,t-1} - P_{CLS,t-1}$	-0.0054	0.0049	0.0020*	-0.0004	-0.0063*		-0.0342
	(0.0432)	(0.0044)	(0.0012)	(0.0004)	(0.0038)		(0.1159)
DEXReturn$_{t-1}$	-0.0055**	-0.0000	-0.0000	0.0000	0.0002*		-0.0022
	(0.0023)	(0.0000)	(0.0000)	(0.0000)	(0.0001)		(0.0018)
$OF_{top10,t-1}$	0.2224**						
	(0.0921)						
$OF_{PM,t-1}$		-0.0004					
		(0.0004)					
$OF_{LP,t-1}$			-0.0002				
			(0.0002)				
$OF_{top10\cap PM,t-1}$				-0.0001			
				(0.0001)			
$OF_{top10\cap LP,t-1}$					-0.0010***		
					(0.0003)		
$OF_{\notin top10,PM,LP,t-1}$							0.1217**
							(0.0592)
constant	0.0003***	-0.0000	-0.0000	0.0000	0.0001***		0.0000
	(0.0001)	(0.0000)	(0.0000)	(0.0000)	(0.0000)		(0.0001)
R-squared	0.051	0.000	0.000	0.000	0.000		0.015
No. observations	14,998	14,998	14,998	14,998	14,998	14,998	14,998

Note: This table presents the results of regressing *private* order flow on the price difference between the DEX and CLS exchange rates. OF measures net buyer transactions of EURC sourced from Uniswap V3 data. $P_{DEX} - P_{CLS}$ measures the price difference between DEX and CLS exchange rates. Order flow is divided into sub-categories such as top 10 wallets, access to primary markets, and LPs. Standard errors are Newey–West (HAC) and reported in parentheses. *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

G.2 Decomposing Order Flow into a Feedback Trading vs Residual Component

Figure A11: Price impact of blockchain order flow: information versus feedback trading (EUR/USD CLS Return)

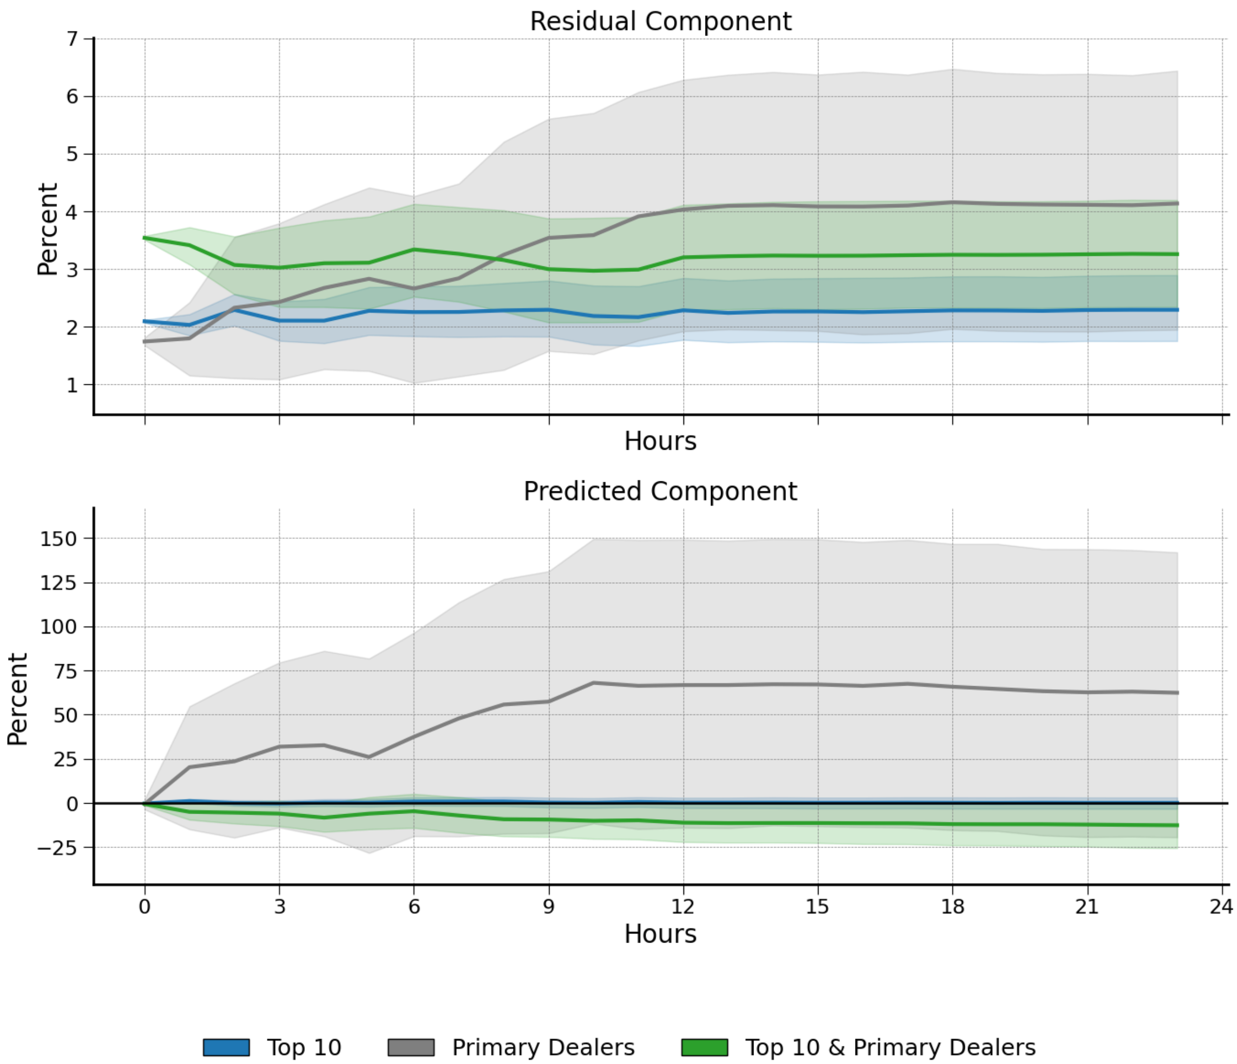

Note: This figure plots the impulse response of spot returns to a one million EURC shock in blockchain order flow using a structural VAR with 1,000 bootstrap replications. Blockchain order flow measures net EURC buyer transactions from Uniswap V3 trade data, while EUR/USD prices are sourced from CLS. To isolate informational content from feedback trading and arbitrage effects between DEX and traditional markets, we decompose order flow by regressing it on the lagged price difference between markets, separating it into a feedback component and a residual component. The top panel shows the response of EUR/USD returns to the residual (information) component, and the bottom panel shows the response to the predicted (feedback/arbitrage) component. Results are presented for blockchain order flow sub-categories: sophisticated traders (top 10 wallets), primary dealers, and their intersecting group. The sample covers 15 August 2022 to 30 April 2024.

Appendix H: Robustness Tests

This appendix presents robustness tests for the structural VAR (SVAR) analysis described in Section 4.4. Each test extends the baseline specification in Appendix F to address potential concerns related to liquidity provision, trading frequency, and trader heterogeneity. We also examine intraday variation in price impact and identify just-in-time liquidity (JIT) behavior among LPs.

The structural form of the model is given by

$$AY_t = A_0 + \sum_{j=1}^{L} A_j Y_{t-j} + \varepsilon_t, \tag{34}$$

where $Y_t = [\mathbf{OF}_t^{OTC}, \mathbf{OF}_t^{DEX}, \Delta p_t]^\top$, and ε_t is a vector of orthogonal structural shocks. The reduced-form VAR is obtained by premultiplying both sides by A^{-1}

$$Y_t = C_0 + CY_{t-1} + B\varepsilon_t, \tag{35}$$

where $B = A^{-1}$, $C_0 = A^{-1}A_0$, and $C = A^{-1}A_1$. The matrix A is lower triangular and imposes the identifying restrictions in the SVAR.

Order flow vectors in the baseline specification are defined as

$$\mathbf{OF}_t^{OTC} = [\mathrm{OF}_{\text{non-bank}}, \mathrm{OF}_{\text{corporate}}, \mathrm{OF}_{\text{fund}}, \mathrm{OF}_{\text{bank}}]^\top,$$

$$\mathbf{OF}_t^{DEX} = [\mathrm{OF}_{\text{LP}}, \mathrm{OF}_{\notin\{\text{Top10,PM,LP}\}}, \mathrm{OF}_{\text{Top10}\cap\text{LP}}, \mathrm{OF}_{\text{PM}}, \mathrm{OF}_{\text{Top10}}, \mathrm{OF}_{\text{Top10}\cap\text{PM}}]^\top.$$

The recursive identification assumes that OTC order flow is exogenous within each period, that DEX order flow can respond contemporaneously to OTC flows, and that price changes respond contemporaneously to both OTC and DEX flows. Within each block, participants are ordered by sophistication, with LPs first and more informed traders later.

The following subsections describe how this structure is modified in each robustness test.

(i) Liquidity Controls. In the first robustness test, we extend the baseline SVAR to control for liquidity provision measured directly from Uniswap V3 mint and burn events. These variables capture changes in the supply of liquidity on either side of the pool and at different distances from the prevailing market price, as detailed in Appendix B.3.

Let $Liquidity_{t,h}^{net,b}$ and $Liquidity_{t,h}^{net,a}$ denote the hourly net liquidity flows within 100 basis points of the market price (*best*) and beyond that range (*away*), respectively:

$$Liquidity_{t,h}^{net,b} = \sum_{k\in h}\left(mint_{(k)}^{b} - burn_{(k)}^{b}\right), \qquad Liquidity_{t,h}^{net,a} = \sum_{k\in h}\left(mint_{(k)}^{a} - burn_{(k)}^{a}\right).$$

These variables measure the net addition or withdrawal of liquidity on the EURC (ask) side relative to the USDC (bid) side, aggregated across all providers. A positive value of $Liquidity_{t,h}^{net,b}$ indicates that LPs have, on balance, increased liquidity close to the market price, while a positive $Liquidity_{t,h}^{net,a}$ indicates greater liquidity provision further away from it.

To test whether changes in liquidity affect the estimated price impact of order flow, we augment the DEX block of the SVAR to include both liquidity variables ordered before all DEX order flows:

$$\mathbf{OF}_t^{DEX,liq} = \left[Liquidity_{t,h}^{net,b},\ Liquidity_{t,h}^{net,a},\ \mathbf{OF}_t^{DEX}\right].$$

This ordering assumes that liquidity supply decisions can contemporaneously affect order flow and returns but not vice versa within the same interval. The full VAR stack becomes

$$Y_t^{liq} = [\mathbf{OF}_t^{OTC},\ \mathbf{OF}_t^{DEX,liq},\ \Delta p_t]^{\top}.$$

Impulse responses reported in Figure A12 show that controlling for liquidity provision leaves the estimated price impact of sophisticated traders and primary dealers largely unchanged, suggesting that our main results are not driven by liquidity adjustments near or away from the best price.

(ii) High-Frequency Specification. To capture finer timing between blockchain and benchmark price adjustments, we estimate the SVAR at a 5-minute frequency. Since CLS benchmark order flow is available only at hourly frequency, the high-frequency specification includes blockchain variables and returns only,

$$Y_t^{HF} = [\mathbf{OF}_t^{DEX}, \Delta p_t^{DEX}, \Delta p_t^{CLS}]^{\top},$$

where CLS returns are interpolated to 5-minute intervals while preserving the information set at the hourly horizon. The recursive structure within the DEX block follows the same ordering as in the baseline, with less informed agents ordered before more sophisticated ones. The results in

Figure A13 confirm that high-frequency shocks from Top 10 wallets and primary dealers continue to have economically and statistically significant effects on both DEX and benchmark prices.

(iii) Trade Size Heterogeneity (Quintile Specification). We next examine whether the informational content of order flow varies systematically with trade size. Instead of grouping wallets by economic role, we partition all DEX traders into five mutually exclusive quintiles based on their average transaction volume and construct aggregate order flow for each group. The order flow vector for DEX becomes

$$\mathbf{OF}_t^{DEX,Q} = [\mathrm{OF}_{\text{Group 1}}, \mathrm{OF}_{\text{Group 2}}, \mathrm{OF}_{\text{Group 3}}, \mathrm{OF}_{\text{Group 4}}, \mathrm{OF}_{\text{Group 5}}]^\top,$$

where Group 1 contains the smallest traders and Group 5 contains the largest. The full VAR stack is

$$Y_t^Q = [\mathbf{OF}_t^{OTC}, \mathbf{OF}_t^{DEX,Q}, \Delta p_t]^\top,$$

and the recursive ordering preserves the same OTC block as in the baseline, followed by DEX quintiles ordered from smaller to larger traders, and finally returns. Figure A14 reports impulse responses to standardized shocks in each quintile and shows that larger traders have stronger and more persistent price impact on both DEX and CLS returns, which is consistent with a higher degree of informational trading.

(iv) Intra-day Patterns. We also analyze intra-day variation in price impact by estimating hour-of-day specific responses to a one million EURC shock in DEX order flow. For each trading hour, we compute the implied impact on DEX and benchmark returns while holding the SVAR structure fixed. Figure A15 shows that the price impact of sophisticated traders and primary dealers is concentrated between 13:00 and 15:00 UTC, overlapping with core European and U.S. trading hours, whereas LPs exhibit no systematic intra-day pattern in their price effects.

(v) Just-in-time Liquidity. Finally, we investigate whether LPs engage in just-in-time (JIT) liquidity strategies, in which LPs mint liquidity immediately before a trade and burn it in the same block in order to collect fees while minimizing exposure to adverse selection. Table G1 lists all detected JIT sequences in the EURC/USDC pool, defined as mint–swap–burn triplets sharing the same block number. These events are concentrated in a single wallet and remain rare over the sample period. Their limited frequency implies that JIT behavior does not materially affect the aggregate price dynamics documented in the SVAR analysis.

H.1 SVAR tests

H.1.1 Controlling for Liquidity Provision

Figure A12: Hourly price impact of blockchain order flow with liquidity controls

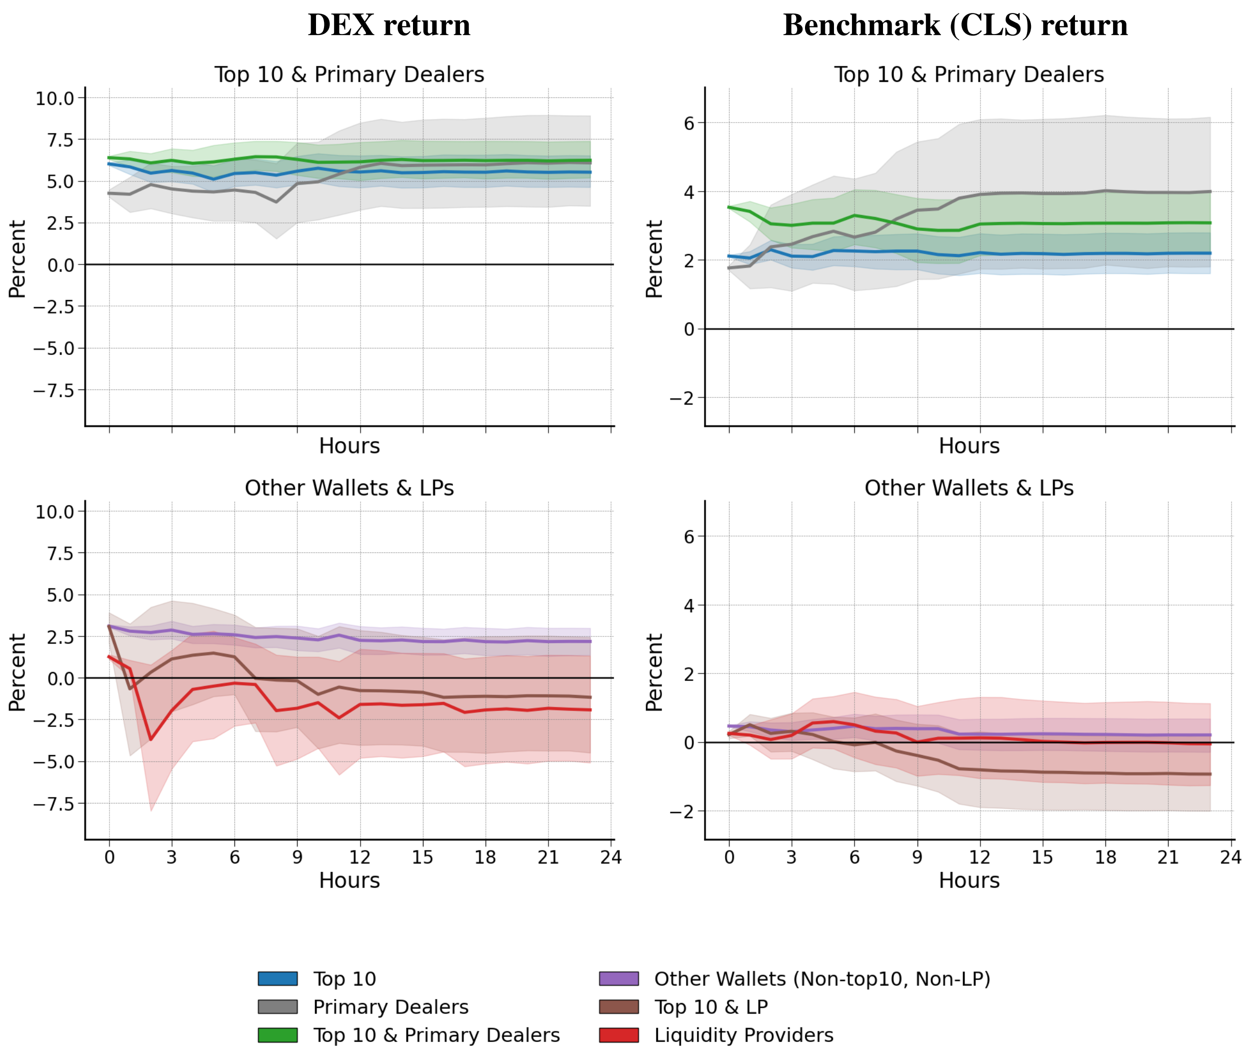

Note: This figure plots impulse responses of returns to shocks in blockchain order flow using hourly data, controlling for liquidity provision. The top row corresponds to transactions by Top 10 wallets and primary dealers, while the bottom row corresponds to other wallets and LPs. The left column shows EURC/USDC returns from Uniswap V3, and the right column shows EUR/USD returns from CLS. Liquidity provision is measured using net liquidity derived from mint and burn imbalances, where positive values indicate additional EURC liquidity in the pool. Responses are estimated using a structural VAR with 1,000 bootstrap replications. The sample covers 15 August 2022 to 30 April 2024.

H.1.2 High-Frequency

Figure A13: High-frequency price impact of blockchain order flow by trader groups

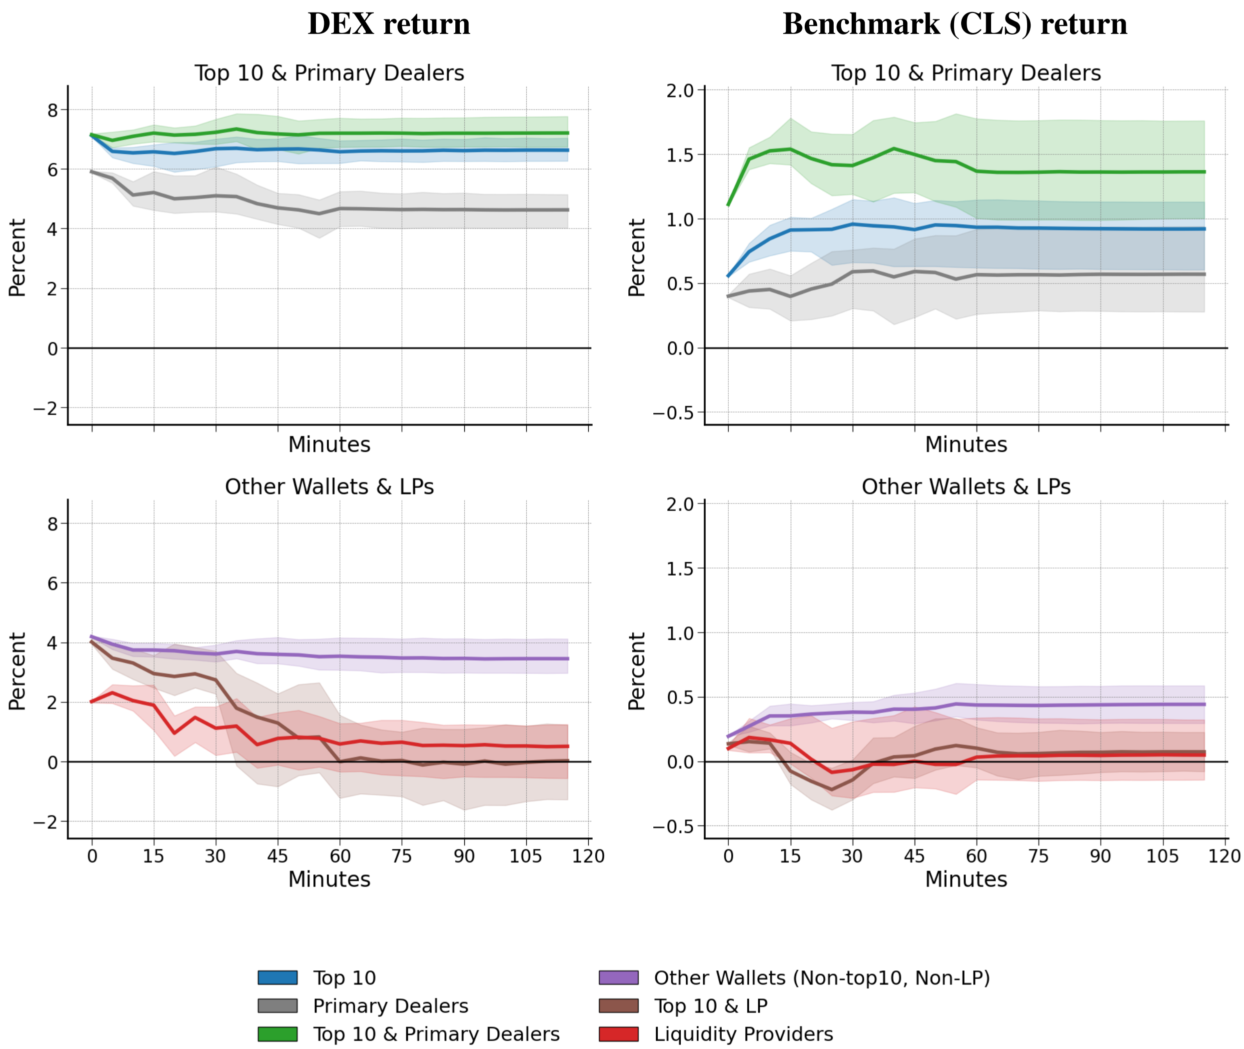

Note: This figure plots high-frequency (5-minute interval) impulse responses of returns to standardized shocks in blockchain order flow. The top row corresponds to transactions by Top 10 wallets and primary dealers, while the bottom row corresponds to other wallets and LPs. The left column shows EURC/USDC returns from Uniswap V3, and the right column shows EUR/USD returns from CLS. Responses are estimated using a structural VAR with 1,000 bootstrap replications. The sample covers 15 August 2022 to 30 April 2024.

H.1.3 Trading Volume

Figure A14: Price impact of (standardized) blockchain order flow by trading volume quintiles

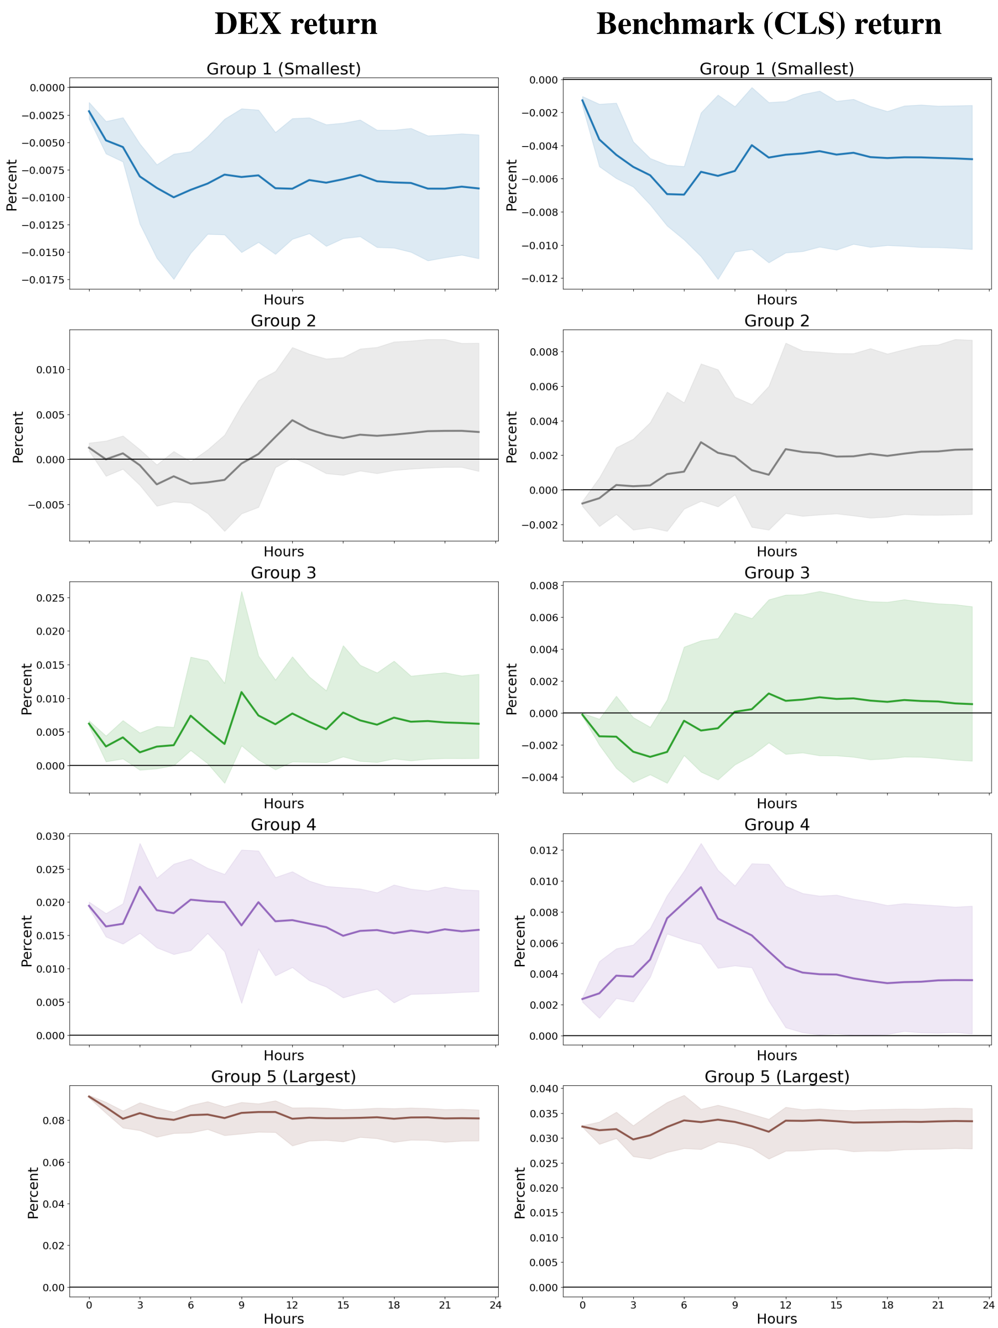

Note: This figure plots impulse responses of returns to a standardized shock in blockchain order flow, grouped by trader volume quintiles (Group 1 = smallest, Group 5 = largest). Shocks are normalized by the standard deviation of order flow within each quintile. The left column shows EURC/USDC returns from Uniswap V3, and the right column shows EUR/USD returns from CLS. Results are estimated using a structural VAR with 1,000 bootstrap replications. The sample covers 15 August 2022 to 30 April 2024.

H.2 Intra-day patterns

Figure A15: Price impact of blockchain order flow: intra-day patterns

Panel (a): EURC/USDC Return

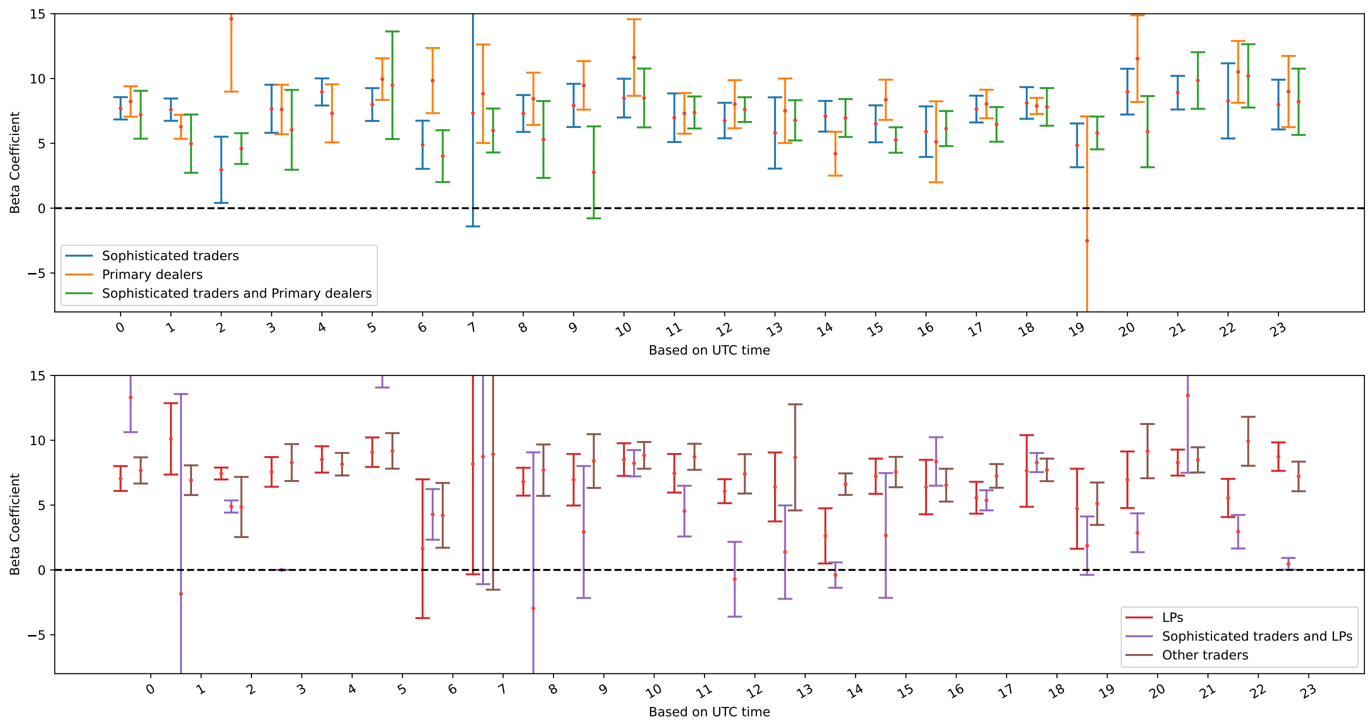

Panel (b): CLS Benchmark EUR/USD Return

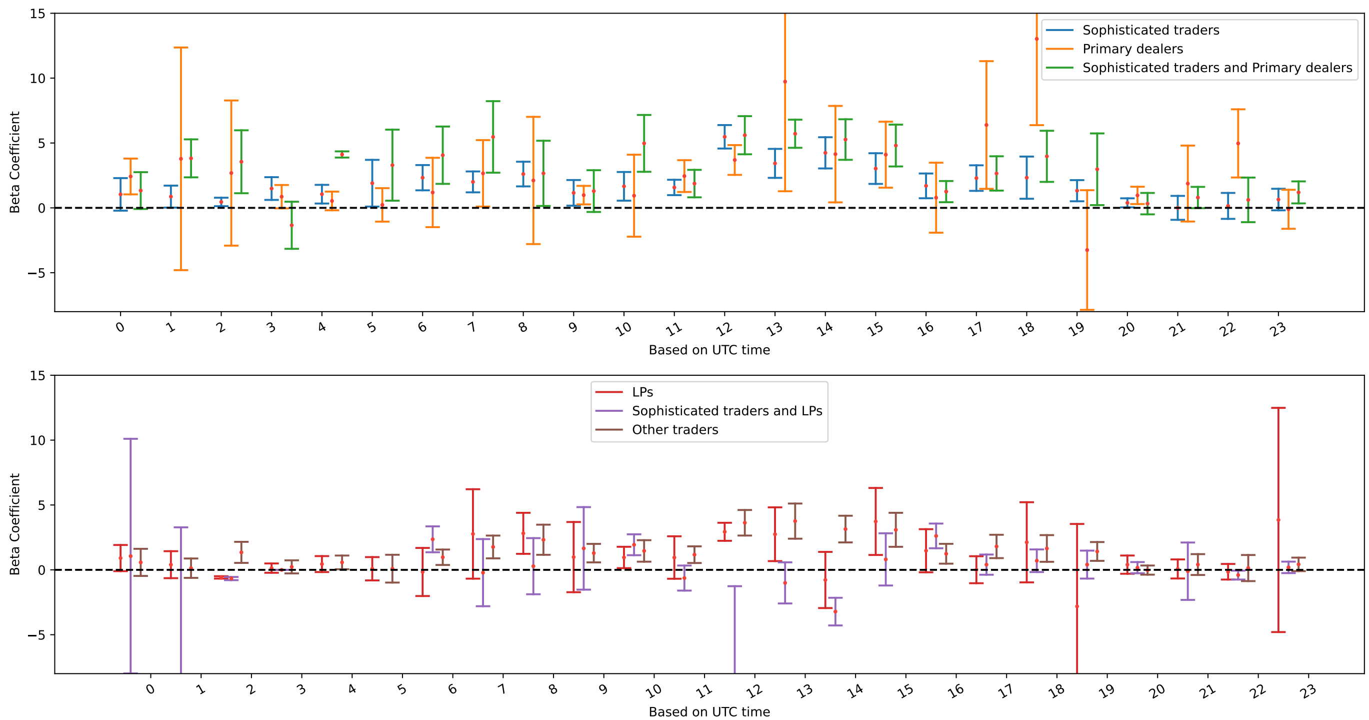

Note: This figure plots hourly price impact estimates in spot returns to a 1 Million EURC shock in blockchain order flow. Blockchain order flow measures net buyer transactions for purchasing EURC and is sourced from Uniswap V3 trade data. EURC/USDC returns are calculated using Uniswap V3 prices, and EUR/USD prices are sourced from CLS. Panel (a) shows the response of EURC/USDC returns, and Panel (b) shows the response of EUR/USD returns. The blockchain order flow is divided into six sub-categories: sophisticated traders (top 10 wallets), primary dealers, LPs, and intersections among these groups. The sample period is from 15 August 2022 to 30 April 2024.

H.3 Sophisticated Liquidity Providers (Just-in-time Liquidity)

Table G1: Transaction Details

Date (UTC)	Blk	Type	User	EURC	USDC	Lower P	Upper P	Price
2023-08-23 07:55	17976054	mint	ae13	50249.82	311076.93	1.09	1.09	
2023-08-23 07:55	17976054	swap	2cc4	-18956.61				1.09
2023-08-23 07:55	17976054	burn	ae13	-32048.08	-330930.63	1.09	1.09	
2023-08-30 09:07	18026424	mint	ae13	82322.59	7347.05	1.09	1.10	
2023-08-30 09:07	18026424	swap	6945	-56915.47				1.09
2023-08-30 09:07	18026424	burn	ae13	-28776.74	-65957.32	1.09	1.10	
2023-09-23 22:53	18201622	mint	ae13	64752.18	238260.06	1.07	1.07	
2023-09-23 22:53	18201622	swap	7cd3	-20246.88				1.07
2023-09-23 22:53	18201622	burn	ae13	-45252.44	-259148.72	1.07	1.07	
2023-10-05 18:33	18286118	mint	ae13	45404.15	7821.33	1.06	1.06	
2023-10-05 18:33	18286118	swap	3592	-9950.00				1.06
2023-10-05 18:33	18286118	burn	ae13	-36795.91	-16936.49	1.06	1.06	
2023-10-06 15:04	18292236	mint	ae13	45905.79	144510.45	1.06	1.06	
2023-10-06 15:04	18292236	swap	c128	-10162.90				1.06
2023-10-06 15:04	18292236	burn	ae13	-36178.39	-154826.68	1.06	1.06	
2023-10-08 00:20	18302152	mint	ae13	71135.53	303399.61	1.06	1.06	
2023-10-08 00:20	18302152	swap	10f2	-9865.26				1.06
2023-10-08 00:20	18302152	burn	ae13	-61490.61	-313649.24	1.06	1.06	
2023-10-11 10:23	18326578	mint	ae13	299169.38	12166.39	1.10	1.10	
2023-10-11 10:23	18326578	swap	aa20	-23186.98				1.10
2023-10-11 10:23	18326578	burn	ae13	-276435.77	-37067.49	1.10	1.10	
2023-10-14 08:03	18347311	mint	ae13	46293.22	12237.06	1.06	1.06	
2023-10-14 08:03	18347311	swap	f7d7	-9964.28				1.06
2023-10-14 08:03	18347311	burn	ae13	-37591.08	-21442.93	1.06	1.06	
2023-10-17 12:33	18370121	mint	ae13	49133.49	6172.49	1.06	1.06	
2023-10-17 12:33	18370121	swap	3592	-19338.40				1.06
2023-10-17 12:33	18370121	burn	ae13	-32279.83	-24072.91	1.06	1.06	
2023-11-03 13:02	18491700	mint	ae13	260626.98	51902.25	1.10	1.10	

Continued on next page

Table G1: Transaction Details (continued)

Date (UTC)	Blk	Type	User	EURC	USDC	Lower P	Upper P	Price
2023-11-03 13:02	18491700	swap	9593	-17213.04				1.10
2023-11-03 13:02	18491700	burn	ae13	-243688.26	-70495.13	1.10	1.10	
2023-11-03 13:13	18491757	mint	ae13	243720.52	69561.03	1.10	1.10	
2023-11-03 13:13	18491757	swap	9593	-20386.83				1.10
2023-11-03 13:13	18491757	burn	ae13	-223658.74	-91671.48	1.10	1.10	
2023-11-07 16:49	18521374	mint	ae13	59330.57	256372.22	1.07	1.07	
2023-11-07 16:49	18521374	swap	46f5	-18621.87				1.07
2023-11-07 16:49	18521374	burn	ae13	-41311.27	-275714.07	1.07	1.07	
2023-11-11 23:46	18552054	mint	ae13	147338.11	36400.63	1.08	1.09	
2023-11-11 23:46	18552054	swap	5319	-38379.19				1.08
2023-11-11 23:46	18552054	burn	ae13	-110425.65	-76438.77	1.08	1.09	
2023-11-30 00:25	18680832	mint	ae13	53340.95	424301.65	1.15	1.15	
2023-11-30 00:25	18680832	swap	b299	-3287.17				1.15
2023-11-30 00:25	18680832	burn	ae13	-50053.88	-428078.41	1.15	1.15	
2024-01-25 17:33	19085149	mint	ae13	208855.00	161529.61	1.12	1.13	
2024-01-25 17:33	19085149	swap	9593	-22789.61				1.12
2024-01-25 17:33	19085149	burn	ae13	-186138.52	-187076.72	1.12	1.13	
2024-01-25 19:18	19085666	mint	ae13	208451.09	66998.77	1.12	1.13	
2024-01-25 19:18	19085666	swap	9593	-18933.66				1.12
2024-01-25 19:18	19085666	burn	ae13	-189597.26	-88198.05	1.12	1.13	
2024-02-09 11:57	19190417	mint	ae13	323979.98	618278.92	1.15	1.15	
2024-02-09 11:57	19190417	swap	54a1	-23745.69				1.15
2024-02-09 11:57	19190417	burn	ae13	-300250.49	-645564.26	1.15	1.15	
2024-02-25 19:05	19306570	mint	ae13	61302.68	16123.49	1.09	1.09	
2024-02-25 19:05	19306570	swap	07d3	-27421.90				1.09
2024-02-25 19:05	19306570	burn	ae13	-36095.84	-43686.38	1.09	1.09	

Note: This table presents Just-in-time Liquidity (JIT) transactions in the EURC–USDC pool. The liquidity provider full address is `0xae2fc483527b8ef99eb5d9b44875f005ba1fae13`, abbreviated as 'ae13'. Each JIT event involves a mint, swap, and burn transaction occurring in the same block. Liquidity is posted at the specified price range (Lower and Upper Price). The sample period is from 15 August 2022 to 30 April 2024.