

Blockchain Currency Markets

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

This draft: March 2026

Journal of Financial and Quantitative Analysis, forthcoming

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 frictions such as gas fees and Ether volatility. Decentralized exchange rates closely track traditional currency markets through arbitrage and informed trading. Traders with substantial market share and access to primary markets exert 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 George G. Pennacchi (editor), an anonymous referee, 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, Tamaro 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, Tsinghua University, 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 Klops for assistance with the CLS Market Data.

1 Introduction

Decentralized finance (DeFi) enables financial services to operate on public blockchains through self-executing smart contracts,¹ allowing assets to trade without traditional intermediaries. A fast-growing segment of this activity involves fiat-referenced stablecoins such as USDC and EURC, which maintain one-to-one parity with the U.S. dollar and the euro.² These stablecoins trade on decentralized exchanges (DEXs), creating an on-chain analogue to traditional foreign exchange (FX) markets. Although still small relative to global FX turnover, such markets have attracted increasing attention as potential infrastructure for cross-border payments and settlement.³

This paper provides the first microstructure analysis of a *blockchain currency market* and assesses whether it is efficiently linked to the underlying currency market. We identify two channels supporting this linkage. The first is an *arbitrage* channel that keeps decentralized and traditional prices tightly aligned, although it is constrained by blockchain frictions such as gas fees and Ether price volatility. The second is an *information channel*, whereby order flow reflects private information about fundamentals. Participants who trade at large scale and have direct access to stablecoin issuance and redemption exhibit informational advantages consistent with the asymmetric-information paradigm in FX markets ([Rinaldo and Somogyi, 2021](#)).

A central goal of the paper is to examine whether—and how—blockchain and traditional FX markets interact. To do so, we construct a granular transaction-level dataset for the EURC/USDC pair on Uniswap V3, the dominant DEX for stablecoin trading, and combine it with price and order data representative of the global traditional currency market obtained from the Continuous Linked Settlement (CLS) system.⁴ A distinctive feature of blockchain markets is the transparency of wallet-level identities, which allows us to classify participants by economic role, distinguishing sophisticated traders, primary dealers with mint and redemption access, and liquidity providers (LPs). Such real-time identification is rarely available in traditional FX microstructure data, where dealer and client activity must typically be inferred indirectly ([Hortaçsu and Sareen, 2005](#); [Hagströmer and Menkveld, 2019](#)).

¹Smart contracts are self-executing programs stored on the blockchain that automatically enforce agreed terms without the need for a trusted intermediary. Their code and execution history are publicly verifiable by all network participants.

²Both tokens are fully backed by high-quality reserves and are redeemable 1:1 for the corresponding fiat currency, with backing verified through regular attestations.

³See BIS Innovation Hub (2023), “Project Mariana: Cross-border trading and settlement using wholesale CBDC and DeFi infrastructure.”

⁴CLS is the primary settlement infrastructure for wholesale FX markets and provides a representative benchmark for interbank EUR/USD pricing.

Using this dataset, we document several features of decentralized FX trading. EURC/USDC prices are closely linked to EUR/USD benchmarks, with average deviations of roughly 24 basis points that narrow during periods of active trading. These deviations covary with blockchain-specific frictions such as gas fees and slippage rather than with balance-sheet constraints that determine pricing in traditional FX markets.⁵ Transaction costs therefore vary substantially across users. Gas fees dominate for smaller traders, while slippage is the primary cost for larger traders. Although these costs exceed those faced by inter-dealer FX participants, they are comparable to those incurred by less-privileged over-the-counter (OTC) clients (Hau et al., 2021).

We then examine the arbitrage mechanism linking decentralized and traditional FX markets. Price differences between EURC/USDC and EUR/USD predict subsequent on-chain order flow, indicating that trading activity corrects cross-market price discrepancies. This mechanism is particularly visible during the USDC de-pegging episode on 11 March 2023, when concerns about reserves held at Silicon Valley Bank (SVB) generated sharp price dislocations across venues. In this episode, subsequent on-chain trading was dominated by sophisticated traders, who were the only group to systematically arbitrage across centralized exchanges and Uniswap pools. Similar dynamics are observed around Federal Reserve announcements. Following policy news, EURC/USDC prices adjust rapidly toward EUR/USD, accompanied by pronounced increases in on-chain trading volume that coincide with heightened activity in traditional FX markets.

Having established the arbitrage channel, we next investigate whether on-chain order flow also conveys private information about fundamentals. Using a structural vector autoregression (SVAR) framework following Hasbrouck (1991), we study the dynamic interaction between blockchain order flow and prices. Order flow from sophisticated traders and primary dealers generates significant and persistent price impacts, indicating that their trades contain information about fundamentals. In contrast, LPs exhibit little or weakly negative price impact, consistent with their role as liquidity suppliers rather than informed traders. These findings are robust to alternative trader classifications and controls for liquidity and just-in-time liquidity provision.

We next examine whether the permanent price effects reflect private information about fundamentals or arbitrage that aligns decentralized prices with traditional benchmarks. We use two complementary tests. We exploit a feature unique to blockchain markets by classifying transactions according to their routing, distinguishing public transactions from private transactions that bypass

⁵Gas 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.

the public mempool.⁶ Private transactions, which are concentrated among sophisticated traders and primary dealers, generate larger and more persistent price impacts, suggesting that these trades embed information relevant for future EUR/USD movements. We also decompose order flow into an arbitrage-driven component predicted by lagged DEX–CLS price differentials and a residual component that we interpret as informational order flow. Only the residual component produces significant price impact with respect to benchmark FX returns.

Taken together, the evidence shows that blockchain currency markets are closely integrated with traditional FX markets through both arbitrage and information channels, and that blockchain transparency offers a unique lens for studying price discovery at high frequency.

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

A first strand studies the design and functioning of DEXs. Recent work assesses market quality and efficiency in automated market makers (AMMs, decentralized trading protocols in which prices are determined algorithmically as a function of pool liquidity rather than through a limit order book) and examines liquidity provision and informed trading incentives in these environments (e.g., [Barbon and Rinaldo, 2024](#); [Capponi and Jia, 2025](#); [Lehar and Parlour, 2025](#)). Related contributions study AMMs in FX and equity settings ([Malinova and Park, 2024](#); [Foley et al., 2023](#)) and use transaction-level blockchain data to analyze market events and trading costs ([Liu et al., 2023](#); [Adams et al., 2023](#)). Relative to this literature, we use wallet-level data to identify informational advantages across trader types and to show how participants link decentralized and traditional FX markets through both arbitrage trading and the incorporation of fundamental information.

A second strand focuses on stablecoins and their price stability. Prior work shows that arbitrage stabilizes on-chain exchange rates across tokens and venues ([Lyons and Viswanath-Natraj, 2023](#); [Ma et al., 2025](#)), while other studies highlight fragility and run risk in collateralized designs ([Gorton et al., 2022](#); [Aldasoro et al., 2023](#); [Eichengreen et al., 2023](#)). We extend this literature by studying stablecoins within the informational efficiency of a blockchain currency market and by assessing the extent to which such markets can intermediate FX trading alongside traditional OTC infrastructure.

We also contribute to the FX microstructure literature on order flow and price formation ([Evans and Lyons, 2002](#); [Rinaldo and Somogyi, 2021](#); [Huang et al., 2025](#)). Traditional FX pricing models emphasize portfolio shifts and inventory management transmitted through inter-dealer trading

⁶The 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.

(Evans and Lyons, 2002), whereas pricing on Uniswap V3 is determined algorithmically through bonding curves that clear trades without intermediaries. Despite these structural differences, heterogeneous traders in blockchain currency markets play economically analogous roles. Analogous to the microstructure tests of Lyons (1995), which distinguish price adjustments driven by inventory control from those driven by private information, we test whether prices are driven primarily by arbitrage constraints and blockchain frictions or by the asymmetric information of sophisticated traders.

Our results indicate that sophisticated traders and primary dealers exhibit informational advantages, while LPs behave passively by reallocating liquidity in response to price movements rather than actively managing inventory risk. Consistent with the asymmetric-information paradigm in FX markets (Rinaldo and Somogyi, 2021), blockchain order flow predicts EUR/USD returns, while arbitrage activity links decentralized prices to traditional FX benchmarks.

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. Section 4 evaluates the informational content of blockchain order flow and its link to traditional FX markets. Section 5 concludes.

2 Data and Institutional Background

2.1 Market Structure of Blockchain and Traditional Currency Markets

Figure 1 summarizes the institutional structure of the traditional EUR/USD market and the blockchain EURC/USDC market studied in this paper. Traditional FX is organized around a two-tier dealer system in which banks intermediate customer trades and trade among themselves in an inter-dealer market that remains central to price discovery (King et al., 2012; Chaboud et al., 2023). Dealers supply liquidity, manage inventory risk, and process information embedded in customer and inter-dealer order flow, consistent with inventory and portfolio-shift models (Evans and Lyons, 2002; Bjønnes and Rime, 2005; Rinaldo and Somogyi, 2021; Huang et al., 2023). Electronic platforms such as Refinitiv and EBS support this structure by allowing dealers to post quotes on centralized limit order books, while non-bank financial institutions, corporates, and asset managers typically access liquidity through dealers.

Blockchain currency markets operate under a different structure. Stablecoins such as EURC

and USDC are issued by a centralized intermediary, Circle, but trade in a decentralized secondary market. In the primary market, Circle’s treasury mints and redeems tokens at par for participants who transact directly with the issuer, which we refer to as *primary dealers*. These dealers link primary and secondary markets by arbitraging deviations from par. When a stablecoin trades above par, dealers can mint tokens and sell at a premium; when it trades below par, they can purchase at a discount and redeem for fiat. These flows help stabilize the peg and connect issuance to secondary-market prices. Appendix A provides additional institutional detail.

In the secondary market, stablecoins are traded across a range of applications, including decentralized and centralized exchanges, lending and liquidity protocols, and cross-border payment systems (Adams et al., 2023). On decentralized exchanges, trades execute against automated market makers implemented as smart contracts, with liquidity supplied by LPs and demanded by traders. We describe the mechanics of AMMs and the roles of different participants further below.

Table 1 highlights several institutional differences between traditional FX and blockchain-based currency markets. For example, FX trades typically settle on a $T + 2$ basis through CLS or correspondent banking networks, whereas blockchain transactions settle atomically within each block.⁷ Trading frictions also differ, with FX execution shaped by credit limits and balance-sheet constraints, while on-chain execution depends on gas fees, priority rules, and network congestion.

[INSERT FIGURE 1 and TABLE 1 ABOUT HERE]

2.2 Automated Market Making in Uniswap

Uniswap is a decentralized exchange implemented as an AMM, in which trading occurs against on-chain liquidity pools managed by smart contracts. Version 1 launched in 2018, followed by V2 in 2020 and V3 in 2021.⁸ We use Uniswap V2 to introduce the basic pricing mechanism and Uniswap V3 to describe the concentrated-liquidity design used in the EURC/USDC market.

⁷Under the $T + 2$ convention, counterparties exchange currencies two business days after the trade date. Settlement risk is mitigated through the CLS system, which provides payment-versus-payment (PvP) settlement. Atomic settlement embeds PvP directly in the transaction layer, so both legs of a trade either execute together or fail together within a single block. On Ethereum, blocks are confirmed roughly every 12–15 seconds, implying near-instant finality relative to traditional settlement cycles.

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

Uniswap V2 and the bonding curve. In Uniswap V2, each pool holds reserves of two tokens, EURC and USDC, denoted by x and y , which satisfy the constant-product invariant

$$k = x \times y. \quad (1)$$

The marginal price of EURC in terms of USDC is $P_{x/y} = y/x$. Because reserves and transactions are public on-chain, price setting is deterministic and liquidity is continuously available as long as both tokens remain in the pool.

Swaps move reserves along the bonding curve defined by (1). If a trader purchases $\Delta x > 0$ EURC in exchange for USDC, post-trade reserves satisfy

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

Since $x' < x$, the marginal price rises mechanically, generating slippage that increases with trade size. Liquidity mints instead scale up k and shift the bonding curve outward. To keep $P_{x/y}$ unchanged, LPs add both tokens in proportion to the prevailing price so that y/x remains constant. Burns reverse this process.

Figure 2 illustrates the relationship between swaps, liquidity provision, and price impact. Panel (a) provides a conceptual illustration in which swaps move the pool along a given bonding curve from E_0 to E_1 , while liquidity mints shift the curve outward by increasing available reserves at the prevailing price, moving the pool from E_0 to E_2 . Panel (b) then proceeds with a numerical example showing how price impact varies with pool depth for a given trade size.⁹

Each pool in Panel (b) starts at $P_{x/y} = 1.10$ but differs in scale. For a low-liquidity pool with $(x, y) = (100, 110)$, a purchase of $\Delta x = 5$ raises the price to $P'_{x/y} \approx 1.219$, while the same trade yields $P'_{x/y} \approx 1.122$ and 1.111 in medium- and high-liquidity pools, respectively.

More generally, the percentage price impact of a trade of size Δx is

$$\text{Price Impact (\%)} = 100 \times \left[\frac{1}{\left(1 - \frac{\Delta x}{x}\right)^2} - 1 \right], \quad (3)$$

which depends on trade size relative to pool depth, with larger pools exhibiting flatter curves and

⁹We use *price impact* to denote the change in the pool price implied by the AMM after a hypothetical trade, also referred to as *pool slippage*. This differs from *transaction-level slippage*, which is measured using the execution price of an actual on-chain transaction relative to the prevailing pool price immediately before the trade.

smaller price responses.

Uniswap V3 and concentrated liquidity. Uniswap V3 generalizes the constant-product design by allowing LPs to concentrate liquidity within a user-defined price range $[p_a, p_b]$ rather than distributing it uniformly over $[0, \infty)$. This improves capital efficiency by ensuring that liquidity is active only when the market price lies within the specified range. V3 also introduces multiple fee tiers (0.01%, 0.05%, 0.3%, and 1%) that segment liquidity across pools and allow LPs to choose different risk–return profiles.¹⁰

Within any active range, trades are priced using *virtual reserves* that satisfy¹¹

$$\left(x + \frac{L}{\sqrt{p_b}}\right) (y + L\sqrt{p_a}) = L^2, \quad (4)$$

where L is the liquidity parameter governing effective depth within the active range. A higher L implies deeper liquidity and smaller price impact. When the range expands to $[0, \infty)$, (4) reduces to $x \times y = L^2$, so L coincides with the Uniswap V2 constant-product parameter.

[INSERT FIGURE 2 and FIGURE 3 ABOUT HERE]

Tick-based pricing and liquidity distribution. V3 discretizes prices into ticks indexed by $i \in \mathbb{Z}$, where $p_i = 1.0001^i$. Tick spacing determines which ticks can host liquidity. In the EURC/USDC 0.05% pool, tick spacing is 10, so liquidity is allocated in 10-tick intervals. Figure 3 plots tick-level liquidity for this pool at block 19,771,559 (April 30, 2024). Liquidity is concentrated near the market price and tapers toward the tails, resembling a tokenized limit order book. Further details on liquidity aggregation and price setting are provided in Appendix B.1 and Appendix B.2.

¹⁰Liquidity for a given trading pair may be distributed across multiple fee-tier pools. For example, ETH/USDC trades in both the 0.05% and 0.3% fee tiers, whereas the EURC/USDC pair trades exclusively in the 0.05% fee tier during our sample period. Lower-fee pools typically attract higher trading volume and tighter spreads, while higher-fee pools offer greater fee income per trade but require less frequent rebalancing by LPs.

¹¹In Uniswap V3, the pricing rule within an active range $[p_a, p_b]$ can be written in constant-product form using notional, or “virtual,” reserves. Specifically, the virtual balances are $x^{\text{virt}} = x + L/\sqrt{p_b}$ and $y^{\text{virt}} = y + L\sqrt{p_a}$, where x and y denote the actual token balances at the current price and L is the liquidity parameter. These virtual reserves determine marginal pricing and price impact but do not correspond to physical balances held by the pool.

2.3 Data

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

We source a benchmark EUR/USD rate from CLS, which provides a volume-weighted average price of interbank quotes at five-minute intervals. We aggregate these data to hourly and daily frequency. EURC/USDC prices are constructed as the last transaction price at each UTC hour (and day) from the Uniswap V3 EURC/USDC pool, obtained via the Subgraph API.¹²

Panel (a) of Figure 4 compares EURC/USDC and EUR/USD prices and plots their difference. Consistent with Adams et al. (2023), the DEX rate tracks the benchmark closely, with an average absolute deviation of 24 basis points. Deviations are larger early in the sample when liquidity is low, so our empirical analysis begins on August 15, 2022. A major episode is the March 2023 USDC de-pegging following concerns about reserve exposure to Silicon Valley Bank; during March 11–12, 2023, EURC/USDC traded at a relative premium versus EUR/USD.

[INSERT FIGURE 4 ABOUT HERE]

2.3.2 DEX Trading Volume and Liquidity Provision

Our Uniswap V3 dataset contains the complete history of swap transactions in the EURC/USDC market. Each swap is recorded at the wallet level, where a wallet corresponds to an Ethereum address controlling the associated tokens.¹³ We complement swap data with proprietary Kaiko data on liquidity provision, which record mint and burn transactions by LPs, including the quantities of EURC and USDC added or removed and the price ranges over which liquidity is allocated.¹⁴

Using wallet histories, we classify market participants into three base groups reflecting trading activity and access to issuance and liquidity provision. *Sophisticated traders* are defined as the top ten wallet addresses by trading volume in each month, consistent with the microstructure view that large trades are more likely to reflect informed or institutional participation (Easley and O’Hara, 1987; Barber et al., 2008). These addresses account for an average of 52 percent of total trading volume between 15 August 2022 and 30 April 2024. *Primary dealers* are wallets that transact directly with the EURC or USDC treasury, minting stablecoins by depositing fiat and

¹²API available at <https://thegraph.com/hosted-service/subgraph/uniswap/uniswap-v3>

¹³Technically, a wallet stores the private keys required to access and authorize transactions from a specific Ethereum address.

¹⁴For 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.

redeeming them at par.¹⁵ *Liquidity providers (LPs)* are wallets that supply or withdraw liquidity by minting and burning positions in the EURC/USDC pool. Each of the latter two groups accounts for approximately 7 percent of aggregate trading volume.

Table 2 reports summary statistics for seven trader groups. These consist of the three base groups, their pairwise intersections, and a residual category. The base groups contain 76 sophisticated traders, 68 primary dealers, and 90 LP addresses. Six addresses are classified as Top 10 primary dealers, seven as Top 10 LPs, and three as primary dealer LPs.¹⁶ The remaining 2,342 addresses form the residual category. Additional details on wallet characteristics and trading and liquidity provision behavior are provided in Appendix C.

[INSERT TABLE 2 ABOUT HERE]

Blockchain Order Flow. We construct signed order flow as net buyer-initiated EURC volume. Each swap reports amount0 (EURC) and amount1 (USDC) from the pool’s perspective. Since EURC is the base currency, amount0 < 0 indicates a buyer-initiated EURC trade, while amount0 > 0 indicates a sale of EURC to the pool. For interval t ,

$$OF_t = \sum_{k \in \mathcal{K}(t)} (\mathbb{1}[T_k = B] - \mathbb{1}[T_k = S]) \times V_k, \quad (5)$$

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

Panel (b) of Figure 4 plots cumulative order flow alongside prices, showing positive co-movement. Disaggregating by trader group reveals that Top 10 traders often take the opposite side of LPs and residual wallets. This pattern is consistent with an asymmetric-information interpretation in which sophisticated traders exploit short-lived mispricing, while LPs and residual wallets supply liquidity on the other side, a channel we test formally in Section 4.

Turning to LPs, their net EURC purchases during periods of EUR appreciation reflect inventory rebalancing rather than informed trading. As swap traders buy EURC from the pool, LP inventories become increasingly tilted toward USDC. LPs respond by withdrawing liquidity from stale price ranges, acquiring EURC through swaps, and re-minting liquidity at updated price bounds. This

¹⁵The USDC Treasury address used to retrieve transaction history is `0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48`, and the EURC Treasury address is `0x1abaea1f7c830bd89acc67ec4af516284b1bc33c`.

¹⁶This latter group, which recorded only six transactions, is excluded from our heterogeneous trading analysis.

behavior is documented at the transaction level in Appendix C.3.3 and is consistent with LPs passively absorbing order flow and rebalancing inventories, in a manner analogous to inventory management by traditional FX dealers (Lyons, 1995; Bjønnes and Rime, 2005).

Liquidity measurement. To characterize the supply-side response to trading activity, we measure net liquidity using Uniswap V3 mint and burn events following Klein et al. (2024). These events record additions and withdrawals of liquidity over price intervals $[P_a, P_b]$ at the block level. For each block k ,

$$\begin{aligned} \text{mint}_{(k)}^{\text{net}} &= \text{mint}_{(k)}^{\text{ask}} - \text{mint}_{(k)}^{\text{bid}}, \\ \text{burn}_{(k)}^{\text{net}} &= \text{burn}_{(k)}^{\text{ask}} - \text{burn}_{(k)}^{\text{bid}}, \\ \text{Liquidity}_{(k)}^{\text{net}} &= \text{mint}_{(k)}^{\text{net}} - \text{burn}_{(k)}^{\text{net}}. \end{aligned} \tag{6}$$

A positive $\text{Liquidity}_{(k)}^{\text{net}}$ indicates that more ask-side (EURC-side) liquidity is added than withdrawn in block k . To distinguish between active and passive liquidity provision, we classify positions within $\pm 1\%$ of the mid-price as *best* liquidity and positions outside this range as *away*. We aggregate block-level liquidity measures to the hourly frequency to align them with order flow and returns. Appendix B.3 provides full construction details.

2.3.3 Additional Data

CLS volume. To study traditional-market activity, we use the CLS FX Spot Flow dataset, which settles around 40% of global FX volume across spot, swaps, and forwards and covers 18 currencies.¹⁷ The data provide hourly buy and sell volumes between banks and three categories of price takers (funds, non-bank financials, corporates) and are widely used in FX microstructure research (Rinaldo and Somogyi, 2021; Kloks et al., 2023; Huang et al., 2023). We construct sector-level volumes for interbank, bank–funds, bank–non-bank financials, and bank–corporates, with interbank volume obtained by subtracting bilateral bank–client flows from the aggregate series.

Figure 5 compares hourly Uniswap EURC/USDC activity with EUR/USD CLS volumes by sector. Traditional trading is concentrated between 13:00 and 16:00 UTC and is driven mainly by interbank and fund–bank flows, coinciding with the overlap of London, Frankfurt, and New

¹⁷The currencies are AUD, CAD, DKK, EUR, HKD, HUF, ILS, JPY, MXN, NZN, NOK, SGD, ZAR, KRW, SEK, CHF, GBP, and USD, covering 33 currency pairs.

York trading hours and the 16:00 UTC WMR fix (Krohn et al., 2024). Blockchain trading occurs throughout the day, with peaks in the afternoon and around 09:00 UTC, consistent with 24/7 trading and retail participation. Average daily CLS volume is 28.42 billion EUR, compared to 0.423 million EURC on Uniswap, or about 0.0015% (0.15 bps) of total EUR/USD trading. Additional summary statistics and intraday liquidity patterns are reported in Appendix C.3.2 and Table 3.

Transaction costs and volatility. Gas fees compensate Ethereum validators for confirming transactions. We retrieve transaction-level gas fees for each swap from Etherscan (in ETH, converted to USDC) and use them to compute arbitrage bounds and trader-level transaction costs.¹⁸ To capture market-wide conditions, we use the EthVol index from T3 Index as a measure of expected 30-day implied volatility for Ether and a daily gas-fee index from CoinMetrics to study price efficiency at the daily frequency.¹⁹

Intermediary constraints. We use two measures of intermediary constraints. The first is the intermediary capital risk factor (ICRF) 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 captures 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.²⁰ The first principal component of standardized VLOOP series explains about 46% of total variation and captures global inter-dealer pricing distortions. Summary statistics are reported in Table 3.

[INSERT FIGURE 5 AND 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 begin by documenting how decentralized markets co-move with traditional FX benchmarks in both prices and quantities. Our baseline measure of price alignment is the absolute deviation

¹⁸Etherscan: <https://etherscan.io>.

¹⁹T3 Index (EthVol): <https://t3index.com>. CoinMetrics: <https://coinmetrics.io>.

²⁰Spot FX data retrieved from LSEG tick history.

between EURC/USDC and the CLS EUR/USD benchmark,

$$\Delta_t = |p_{\text{EURC/USDC},t} - p_{\text{EUR/USD},t}|. \quad (7)$$

The average deviation is 24 basis points, with a median of 16 basis points, and exceeds 200 basis points only during the March 2023 USDC de-pegging episode. These moments indicate that decentralized prices generally track the benchmark closely, although short-lived dislocations arise under stress.

Trading volumes on decentralized and traditional venues also display a strong connection. DEX activity follows the intraday pattern of CLS volumes, with a clear peak between 13:00 and 16:00 UTC when European and U.S. markets overlap. Using CLS volumes disaggregated into interbank, non-bank financial, and corporate sectors, and comparing them to DEX activity across sophisticated traders, primary dealers, and LPs, we find that decentralized trading mirrors the pattern of the interbank segment, which the literature identifies as the most informed part of the FX market (Rinaldo and Somogyi, 2021; Huang et al., 2023). 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 the core trading window when New York, London, and Frankfurt are open.

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

$$V_{N_{DEX},t} = \alpha + \sum_{i \in N_{CLS}} \beta_i V_{N_{CLS},t} + \varepsilon_t, \quad (8)$$

where $V_{N_{DEX},t}$ denotes DEX trading volume by participant type, measured in EURC, and $V_{N_{CLS},t}$ denotes CLS sector-level trading volume, measured in EUR millions. Table 4 shows that interbank trading is most strongly associated with DEX activity. For sophisticated traders, a €1 million increase in interbank CLS volume is associated with approximately 4.35 units of EURC trading on DEXs.

Figure 6 provides additional evidence. Panel (a) shows that DEX volumes for all trader categories peak between 13:00 and 16:00 UTC and decline outside these hours, with volumes falling by about 50 percent for sophisticated traders and 37 percent for primary dealers. For traders classified as both sophisticated and primary dealers, the drop reaches 74 percent. Panel (b) shows that weekend activity falls significantly across groups, with the steepest decline of 87 percent among sophisticated primary dealers. Together, these findings demonstrate that decentralized FX markets move closely

with traditional FX markets across 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 Frictions

In an efficient market, decentralized prices should track benchmark EUR/USD values tightly. Persistent deviations therefore reveal the frictions that limit arbitrage. To identify these frictions, we regress Δ_t on blockchain execution costs, stablecoin fundamentals, and traditional intermediary constraints at the daily frequency,

$$\Delta_t = \beta_0 + \beta_1 \text{gasfee}_t + \beta_2 \sigma_{ETH,t}^{IV} + \beta_3 R_{ETH,t} + \beta_4 |p_{\text{USDC/USD},t} - 1| + \beta_5 |p_{\text{EURC/EUR},t} - 1| + \beta_6 \text{VLOOP}_t + \beta_7 \text{ICRF}_t + \varepsilon_t, \quad (9)$$

where gasfee_t , $\sigma_{ETH,t}^{IV}$, and $R_{ETH,t}$ capture on-chain execution costs and crypto-market risk, $|p_{\text{USDC/USD},t} - 1|$ and $|p_{\text{EURC/EUR},t} - 1|$ measure stablecoin peg deviations, VLOOP is a standardized measure of FX triangular arbitrage violations, and ICRF is the intermediary capital risk factor of He et al. (2017), which proxies shocks to U.S. dealer balance-sheet constraints.

Table 5 shows that blockchain-specific frictions are the primary determinants of price deviations. Both gas fees and Ether volatility consistently explain variation in Δ_t . A \$1 increase in gas fees raises Δ_t by about 1.3 basis points, while a one basis point increase in implied Ether volatility increases deviations by 0.13 to 0.16 basis points. Ether volatility matters not because of dealer balance-sheet constraints, but because it raises uncertainty for traders holding cryptoasset portfolios, reducing their willingness to deploy capital for arbitrage.

Stablecoin fundamentals also play an important role. The significance of $|p_{\text{USDC/USD}} - 1|$ reflects spillovers from USDC market conditions into EURC/USDC pricing. Frictions in USDC/USD markets constrain arbitrage capital for participants benchmarking EURC/USDC to EUR/USD, and as the more liquid leg of the pair, USDC/USD conditions provide a more informative proxy for these constraints. This pattern is partly explained by the SVB de-pegging event of March 2023, during which Circle’s disclosure of \$3.3 billion in SVB deposits caused USDC to temporarily trade below par, generating sharp spikes in Δ_t .²¹

²¹For details on primary-market issuance, see Appendix A. Peak issuance in our sample is approximately 60 billion

Consistent with this mechanism, column (5) shows that a one basis point increase in the USDC/USD peg deviation is associated with a 0.75 basis point rise in Δ_t , whereas deviations in the EURC/EUR peg are positive but statistically insignificant. To assess whether this effect holds outside the de-pegging episode, column (6) re-estimates the full specification excluding the four days of the SVB episode (10–13 March 2023). The R^2 falls from 0.23 to 0.08 and the coefficient rises to 1.84 with a larger standard error, reflecting more limited variation in the USDC peg outside the event, but the coefficient remains statistically significant.

All other coefficients remain stable across columns (5) and (6), confirming that the core blockchain-friction results are not driven by the SVB episode. Traditional intermediary constraints such as VLOOP and ICRF are insignificant, suggesting that balance-sheet frictions binding intermediaries in OTC FX markets do not carry over to decentralized venues. Taken together, these results indicate that price inefficiencies in the EURC/USDC market arise primarily from blockchain-specific frictions (execution costs, on-chain congestion, and stablecoin fundamentals) rather than traditional intermediary constraints (Barbon and Ranaldo, 2024; Foley et al., 2023).

[INSERT TABLE 5 ABOUT HERE]

3.3 Trading Costs

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

We next examine transaction costs at the trade level. Total cost, expressed in basis points relative to trade size, is the sum of gas fees, LP fees, slippage, and private-routing fees. Gas fees are payments to validators for processing transactions on the public blockchain and are largely fixed per transaction, implying higher proportional costs for small trades. LP fees are constant at 5 basis points in the EURC/USDC pool. Slippage reflects the price impact of consuming liquidity along the bonding curve and is measured as the difference between the execution price and the pool price immediately before the trade. Private-routing fees arise only for privately routed transactions and represent payments to validators for authenticating these orders; they 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.

USDC, compared with about 80 million EURC.

Figure 7 shows substantial heterogeneity in trading costs across participants. Panel (a) indicates that larger traders face lower and more concentrated cost distributions, while smaller or residual wallets incur higher and more dispersed costs. Panel (b) decomposes median total costs by component. Gas fees are the dominant cost for most users, accounting for roughly 7 to 23 basis points across groups and representing a larger fraction of total costs for smaller traders. Slippage becomes more important for larger trades, reaching up to about 20 basis points, while LP fees are constant by design at 5 basis points. Overall, median total trading costs range from around 15–17 basis points for LP- and PM-related addresses to roughly 45 basis points for Top 10 LP wallets, with Top 10 traders facing median costs just above 30 basis points.

[INSERT FIGURE 7 ABOUT HERE]

These results have two main implications. First, trading costs define the effective bounds within which arbitrage can operate. Appendix D shows that once gas fees, slippage, LP fees, and private-routing fees are incorporated, the share of arbitrage-bound violations is approximately 16–19%, falling to 3–5% after accounting for centralized exchange fees and OTC bid–ask spreads. The EURC/USDC market is therefore best characterized as *constrained price efficient* rather than fully frictionless.

Second, decentralized trading costs compare favorably with those faced by many OTC clients. While on-chain costs exceed inter-dealer spreads (EUR/USD spreads averaged only 0.55 basis points in 2023 (Filippou et al., 2024)), they are of a similar order of magnitude to those paid by less-privileged OTC clients. Hau et al. (2021) show that FX derivatives clients at the 90th percentile face spreads of up to 50 basis points. Liquidity depth remains limited, particularly for large trades, but decentralized venues may nonetheless expand access to FX trading for smaller participants.

4 Informational Efficiency

4.1 Blockchain Order Flow and Arbitrage Trading

We examine how information is incorporated into prices in blockchain currency markets, beginning with the *arbitrage* mechanism. Arbitrage trades exploit price differences across venues and contribute to keeping the on-chain reference rate (EURC/USDC) aligned with the traditional benchmark (EUR/USD). Because transaction costs such as gas fees are largely fixed per transaction,

arbitrage incentives are strongest for traders who can execute at scale, making participation uneven across market participants.

To measure these responses, we test whether DEX order flow adjusts to lagged price differences between the DEX reference rate and the CLS benchmark rate. The analysis is conducted at the hourly frequency to capture short-horizon arbitrage dynamics. Specifically, we estimate Equation (10), regressing blockchain order flow on the lagged DEX–benchmark price differential, with controls that include the lagged EURC/USDC return,

$$OF_{i,t} = \alpha + \beta_1(p_{EURC/USDC,t-1} - p_{EUR/USD,t-1}) + \text{controls}_{i,t} + \epsilon_{i,t}. \quad (10)$$

[INSERT TABLE 6 ABOUT HERE]

Table 6 shows that benchmark deviations predict subsequent order flow primarily for 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.

4.1.1 Public Information

We next examine trading responses to public information. These episodes provide clean settings to study the arbitrage mechanism, as common shocks generate price dislocations across centralized and on-chain venues. We focus on two types of public information: the March 2023 USDC de-pegging episode and scheduled Federal Reserve announcements.

USDC De-Pegging Event. The USDC de-pegging on 11 March 2023 provides a natural setting to study arbitrage under stress. When it became public that \$3.3 billion of USDC reserves were held at Silicon Valley Bank, which had entered resolution, concerns about reserve accessibility caused USDC to trade at a discount and fall to 87 cents on centralized exchanges. Confidence recovered on 13 March after the Federal Deposit Insurance Corporation announced that all SVB deposits would be fully protected.²²

²²Further details on USDC reserves 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>.

[INSERT FIGURE 8 ABOUT HERE]

Figure 8 plots deviations of EURC/USDC from the EUR/USD benchmark together with disaggregated blockchain order flow. Sophisticated traders are net buyers of EURC throughout the episode, while primary dealers and smaller wallets exhibit net selling pressure. This behavior is consistent with arbitrage opportunities created by the USDC de-pegging. As USDC traded at a discount on centralized exchanges, sophisticated traders could purchase USDC on those venues and sell it into Uniswap V3 pools where its relative price was higher, thereby arbitraging the temporary dislocation. LPs behaved differently. Their swap and mint–burn activity did not respond to price discrepancies, and their inventories adjusted mechanically as traders executed swaps, consistent with passive liquidity provision. Appendix E.1 provides additional evidence of repeated cross-market arbitrage across multiple pools, shows that cost and routing advantages favored Top 10 wallets, and documents minimal liquidity re-positioning during the episode.

Monetary Announcements. Public information also arrives through scheduled macroeconomic releases. Appendix E.2 examines Federal Reserve FOMC announcements and shows that EURC/USDC deviations from CLS benchmarks remain small after releases, while trading volumes rise sharply in both markets. Prices in the two venues move together almost immediately, indicating active arbitrage as macroeconomic news is incorporated.

Taken together, the regression evidence and event studies show that arbitrage linking decentralized and traditional FX prices is concentrated among sophisticated traders and primary dealers, consistent with scale economies in blockchain execution costs, while other traders respond weakly to short-horizon price gaps and LPs remain largely passive.

4.2 Blockchain Order Flow and Fundamental Information

4.2.1 Permanent Price Impact

We study whether blockchain order flow contains information about exchange rate fundamentals beyond arbitrage by examining its *permanent* price impact. In asymmetric-information models, public news is incorporated rapidly, whereas private information is revealed gradually through order flow (Evans and Lyons, 2002). In traditional FX markets, this process operates primarily through inter-dealer trading. Blockchain markets lack such an inter-dealer tier, so informational trading is transmitted directly through trades rather than through dealer inventories. In this sense,

informed trading on chain more closely resembles the framework of Kyle (1985), in which traders with private signals reveal information through their trades.

The presence of informed trading is likely to differ across participants. Primary dealers, with direct access to EUR and USD funding and links to interbank markets, are plausibly more exposed to fundamental signals. Sophisticated traders, operating at scale and facing lower execution costs, can transmit such information across venues. By contrast, LPs primarily supply inventory and are exposed to adverse selection when prices move on fundamentals (Capponi and Jia, 2025). This heterogeneity motivates an empirical design that estimates permanent price impact separately by trader type.

We estimate a structural VAR 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 incorporated through institutional trading, we control for traditional FX order flow from CLS. This yields a block SVAR with traditional OTC order flow (\mathbf{OF}_t^{OTC}), blockchain order flow (\mathbf{OF}_t^{DEX}), and exchange rate changes (Δp_t):

$$\begin{bmatrix} \mathbf{OF}_t^{OTC} \\ \mathbf{OF}_t^{DEX} \\ \Delta p_t \end{bmatrix} = \alpha + \sum_{k=1}^L \mathbf{A}_k \begin{bmatrix} \mathbf{OF}_{t-k}^{OTC} \\ \mathbf{OF}_{t-k}^{DEX} \\ \Delta p_{t-k} \end{bmatrix} + \epsilon_t. \quad (11)$$

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

Identification uses a recursive Cholesky ordering $\mathbf{OF}_t^{OTC} \rightarrow \mathbf{OF}_t^{DEX} \rightarrow \Delta p_t$. This allows traditional OTC order flow to affect blockchain order flow and prices contemporaneously, while DEX flows do not contemporaneously affect OTC flows, and prices respond immediately to all order flows. The ordering reflects the hierarchical structure of FX trading venues and standard information transmission assumptions in the microstructure literature. Full details are provided in Appendix F.

Figure 9 reports impulse response functions over a 24-hour horizon. Panel (a) shows responses of EURC/USDC returns, while Panel (b) shows responses of benchmark EUR/USD returns from

CLS to shocks in DEX order flow by participant type. Order flow from sophisticated traders and primary dealers generates the largest and most persistent effects on benchmark returns, whereas flows from liquidity providers have negligible impact.

Quantitatively, a one million EURC shock to primary dealer order flow produces a cumulative 24-hour impact of approximately 3.95 percent on CLS benchmark EUR/USD returns, while an equivalent shock to sophisticated trader order flow produces an impact of about 2.2 percent. The combined group of Top 10 wallets that are also primary dealers generates a cumulative impact of 3.1 percent. In contrast, order flow from LPs and residual wallets has negligible permanent effects on benchmark returns, with estimated impacts of approximately -0.05 and 0.2 percent, respectively, consistent with passive liquidity provision and uninformed trading.

Scaling these responses to a one-standard-deviation shock in daily order flow implies permanent impacts of roughly 1.6 basis points for sophisticated traders and 2.9 basis points for primary dealers.²³ These effects do not rely on inter-dealer inventory adjustment mechanisms. Instead, they are consistent with persistent price adjustment driven by heterogeneity in trader information and execution, in line with informed trading models such as [Kyle \(1985\)](#).

[INSERT FIGURE 9 ABOUT HERE]

Robustness. Appendix [G](#) reports a full set of robustness checks. We re-estimate the baseline SVAR under alternative specifications, including controls for net liquidity provision (within and beyond $\pm 1\%$ of the best price), a higher-frequency (5-minute) SVAR, and trader-volume quintiles. The main results are stable. Price impacts remain concentrated among top wallets and primary dealers, while LPs and smaller wallets exhibit weaker and less persistent effects. We also examine intraday variation (Appendix [G.2](#)) and find that price impacts are concentrated between 13:00 and 15:00 UTC, when major global financial markets are simultaneously open. Finally, we examine just-in-time liquidity provision ([Capponi et al., 2024a](#)). Appendix [G.3](#) identifies one wallet, ending in ‘ae13’, that consistently displays this behavior in the EURC/USDC pool.²⁴ This behavior is rare

²³Permanent price impact estimates are expressed in percent returns per one million EURC order flow shock. A one-standard-deviation shock in daily order flow is approximately 7,230 EURC. When scaled to comparable units, these impacts are smaller than those typically reported for traditional FX markets, which is economically intuitive given the much smaller scale of decentralized trading. For example, [Evans and Lyons \(2002\)](#) report price impacts of roughly 50 basis points per USD 1 billion of order flow in inter-dealer FX markets. The sample average daily trading volume in the EURC/USDC market is EURC 0.423 million (standard deviation EURC 0.674 million).

²⁴For example, on 23 August 2023, wallet ‘ae13’ deposited 50,249 EURC and 311,077 USDC within a narrow price range and withdrew the position shortly after a large trade by another user, a pattern consistent with strategic fee capture with minimal exposure.

in our sample and does not materially affect our estimates of permanent price impact.

4.3 Private Information and Permanent Price Impact

The permanent price effects documented above raise an identification issue as to whether these effects reflect private information about exchange rate fundamentals or arbitrage activity that mechanically aligns decentralized prices with traditional benchmarks. We address this distinction using two complementary tests. First, we exploit variation in transaction routing by comparing the price impact of privately routed and publicly broadcast trades. Second, we decompose blockchain order flow into an arbitrage-driven component predicted by cross-venue price differentials and a residual component orthogonal to these differentials, which we interpret as informational order flow. All figures and tables referenced in this subsection are reported in Appendix H.

Private Transactions. First, we exploit variation in transaction routing by separating *private* transactions, which bypass the public mempool, from *public* transactions and comparing their price effects within a joint SVAR. Private routing can reflect either strategic execution to avoid frontrunning and MEV extraction when conducting large trades, or an attempt to conceal fundamental information during execution. These mechanisms have different implications: MEV-related routing should co-move with contemporaneous cross-venue price gaps, whereas information-concealing routing should be associated with more persistent effects on benchmark prices.

We identify private transactions using Blocknative mempool archives (Appendix C.2.4). Such transactions are concentrated among Top 10 and Top 10 \cap LP addresses, while primary dealers and liquidity providers transact almost exclusively through the public mempool. We re-estimate the SVAR by splitting blockchain order flow into public and private components for Top 10 traders, primary dealers, and LPs, ordering each private component immediately after its corresponding public component.²⁵

Figure H.1 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 \cap LP addresses. For benchmark CLS EUR/USD returns, the estimated permanent price impact at a

²⁵Formally, if the baseline SVAR ordering of DEX order flow groups is $[OF_1, OF_2, OF_3, \dots]$, the revised ordering is $[OF_{1,\text{public}}, OF_{1,\text{private}}, OF_{2,\text{public}}, OF_{2,\text{private}}, OF_{3,\text{public}}, OF_{3,\text{private}}, \dots]$. This ordering reflects the information structure of blockchain execution. Public transactions are broadcast through the mempool and are therefore observable to private agents, who can condition on these signals before execution. The remainder of the SVAR structure follows Equation (11), with traditional OTC order flow ordered first, DEX order flow ordered second, and exchange rate changes last.

24-hour horizon of a one million EURC shock from Top 10 wallets is 2.75 percent for private transactions, compared with 1.15 percent for public transactions. For Top 10∩LP flows, the corresponding estimates are 1.25 percent and −0.60 percent, respectively.

Tables H.1 and H.2 examine whether public and private order flows respond to DEX–CLS price differentials, which proxy for arbitrage incentives. Public transactions display stronger and more systematic sensitivity to these price gaps than private transactions, consistent with an arbitrage channel that contributes to price alignment across venues. By contrast, private transactions, particularly those initiated by Top 10 traders, show little or no response to DEX–CLS price differentials, consistent with information-motivated trading or execution considerations. The only exception is the Top 10∩LP group, whose private transactions exhibit a weak response to price differentials, consistent with LPs using private routing to mitigate frontrunning and MEV extraction during large repositioning trades.

We interpret private transactions that are unresponsive to DEX–CLS price differentials as reflecting the trading of fundamental information rather than arbitrage. While private routing may also facilitate MEV-related strategies or high-frequency arbitrage not fully captured at the hourly frequency, we find supporting evidence that arbitrage flows occur through the public mempool: during the March 2023 USDC de-pegging episode, all observed EURC/USDC transactions by the sophisticated trader in Table E.1 were publicly broadcast,²⁶ indicating that large traders exploiting cross-pool dislocations did not rely on private routing in this episode. At higher frequencies, such as within-block MEV strategies,²⁷ arbitrage may still occur through private channels.

Arbitrage versus Informational Order Flow. Second, we separate informational effects from arbitrage by decomposing blockchain order flow into a component predicted by lagged DEX–CLS price differentials and a residual component orthogonal to these differentials. We estimate this decomposition using Equation (10), regressing net EURC order flow on lagged cross-venue price gaps. The fitted values capture arbitrage-driven order flow, while the residual proxies for information-driven trading.

Figure H.2 reports impulse responses of CLS benchmark EUR/USD returns to these components.

²⁶Appendix Table E.1 shows that the large EURC/USDC transactions of the sophisticated trader during the USDC de-pegging episode were all publicly broadcast. This provides transaction-level evidence that public routing was used during an episode with strong cross-pool arbitrage incentives.

²⁷MEV (Maximal Extractable Value) refers to the profit extractable by reordering, inserting, or censoring transactions within a block. Common strategies include sandwich attacks and atomic arbitrage across liquidity pools executed within a single block. See Capponi et al. (2025) for analysis of MEV and frontrunning in DEX settings.

The residual component generates persistent and statistically significant price effects, particularly for large traders and the intersection of Top 10 traders and primary dealers, whereas the arbitrage component has no significant impact on benchmark returns. Together with the routing evidence, these results indicate that the permanent price effects documented above are driven primarily by the incorporation of fundamental information rather than by cross-venue arbitrage.

5 Conclusion

Decentralized finance introduces a trading architecture based on smart contracts and automated liquidity provision, offering a market structure that differs fundamentally from traditional dealer-intermediated FX trading. Using granular blockchain and CLS data, we study the informational efficiency of the EURC/USDC market on Uniswap V3 and its connection to the traditional EUR/USD market.

We document two main findings. First, decentralized prices and trading activity track traditional FX benchmarks closely. Deviations are generally small and are determined primarily by blockchain-specific frictions such as gas fees, and while on-chain transaction costs exceed inter-dealer spreads, they are comparable to the costs faced by many OTC clients outside preferred pricing tiers.

Second, information is incorporated into decentralized prices through both arbitrage and asymmetric-information channels. Arbitrageurs quickly eliminate cross-market price differences and anchor DEX prices to traditional benchmarks, including during stress episodes and around macroeconomic announcements. In contrast, persistent price effects arise from trades executed by sophisticated investors and primary dealers, whose order flow contains private information about exchange-rate fundamentals. Liquidity providers play a largely passive role, supplying liquidity but contributing little to price discovery.

These results highlight that, despite their modest size, blockchain-based currency markets already function as informationally linked extensions of traditional FX markets. As blockchain infrastructure matures and frictions decline, decentralized venues may increasingly serve as complementary platforms for trading and settlement, broadening access to FX markets and potentially reshaping the microstructure of cross-border finance.

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.
- Barber, Brad M, Terrance Odean, and Ning Zhu**, “Do retail trades move markets?,” *The Review of Financial Studies*, 2008, 22 (1), 151–186.
- Barbon, Andrea and Angelo Ranaldo**, “On The Quality Of Cryptocurrency Markets: Centralized Versus Decentralized Exchanges,” *Available at SSRN 3984897*, 2024.
- 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**, “Liquidity provision on blockchain-based decentralized exchanges,” *The Review of Financial Studies*, 2025, 38 (10), 3040–3085.
- , —, 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 **Kanye Ye Wang**, “Maximal extractable value and allocative inefficiencies in public blockchains,” *Journal of Financial Economics*, 2025, 172, 104132.
- , —, 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.
- Easley, David and Maureen O’Hara**, “Price, trade size, and information in securities markets,” *Journal of Financial economics*, 1987, 19 (1), 69–90.

- 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.

Figures

Figure 1: Structure of Traditional and Blockchain Market

Note: This figure presents a schematic comparison of traditional and blockchain-based currency markets. In traditional FX markets, liquidity is provided by dealer banks operating in an interdealer market and a dealer–customer market, where dealers trade with corporates, funds, and non-bank financial institutions. In the blockchain market, issuance occurs in a primary market in which Circle, through its treasury operations, mints EURC and USDC tokens and distributes them to primary dealers. These tokens are subsequently traded in secondary markets, including centralized exchanges with limit order books and decentralized exchanges such as Uniswap, where EURC and USDC trade against each other. Secondary-market participants include liquidity providers and sophisticated traders.

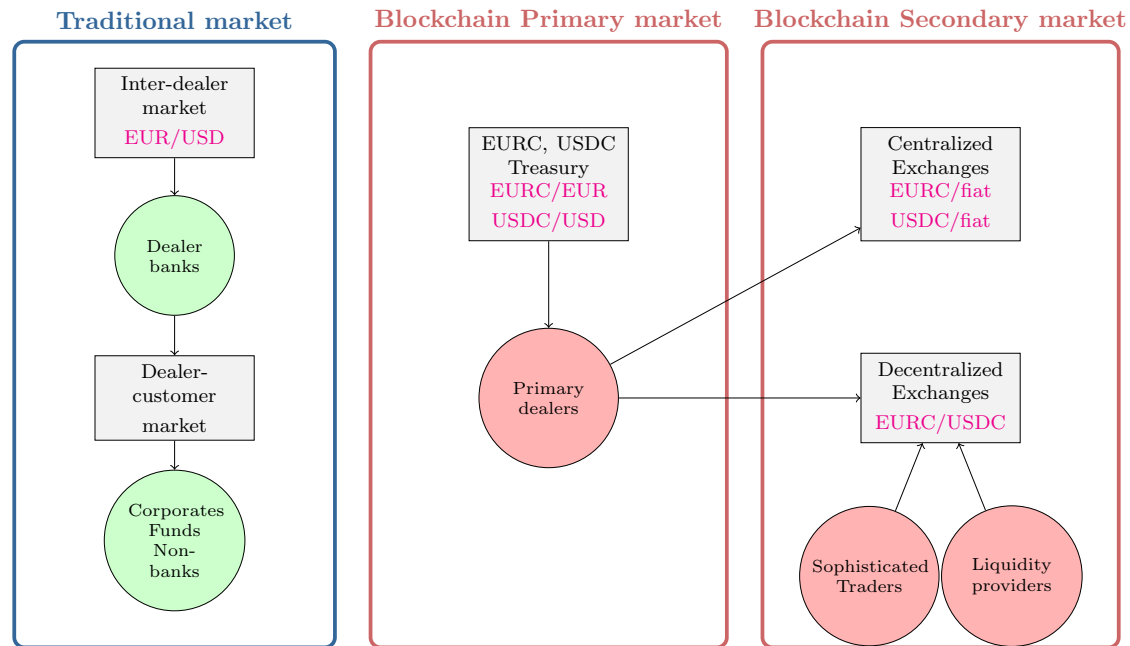



Figure 2: EURC/USDC Bonding Curves and Slippage

Note: Panel (a) illustrates bonding curves and liquidity provision in Uniswap. Starting from an initial equilibrium E_0 , a swap trade that purchases EURC moves the pool along the bonding curve to E_1 , while an increase in liquidity at the prevailing price shifts the equilibrium from E_0 to E_2 . Panel (b) illustrates the relationship between pool size and price impact in a constant-product automated market maker for the EURC/USDC pair. The left subpanel shows bonding curves for different liquidity levels, while the right subpanel shows percentage price impact for trades of varying size (in EURC). Price impact is lower in more liquid pools, reflecting reduced slippage.

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



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

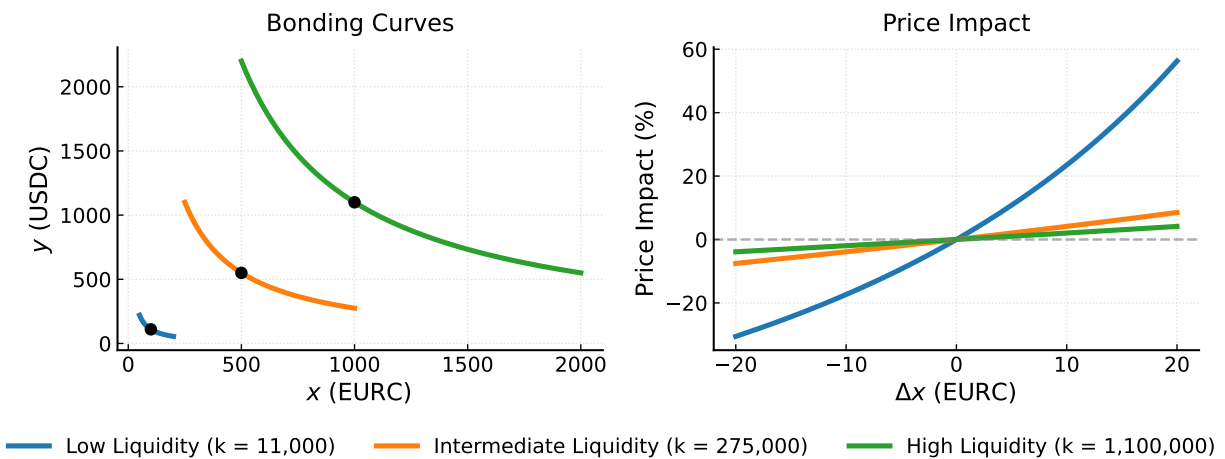


Figure 3: **Snapshot of EURC/USDC Liquidity.**

Note: This figure shows the tick-level distribution of liquidity around the prevailing pool price for the EURC/USDC 0.05% Uniswap V3 pool. The horizontal axis reports tick distance from the current market price (tick 0), where each tick corresponds to a discrete price interval in log base $\sqrt{1.0001}$. The pool has a fixed tick spacing of 10, implying price intervals of approximately 10 basis points. Ticks to the left of zero represent liquidity below the current price and correspond to buy limit orders for EURC, while ticks to the right represent liquidity above the current price and correspond to sell limit orders for EURC. Liquidity at each tick is expressed in thousands of EURC.


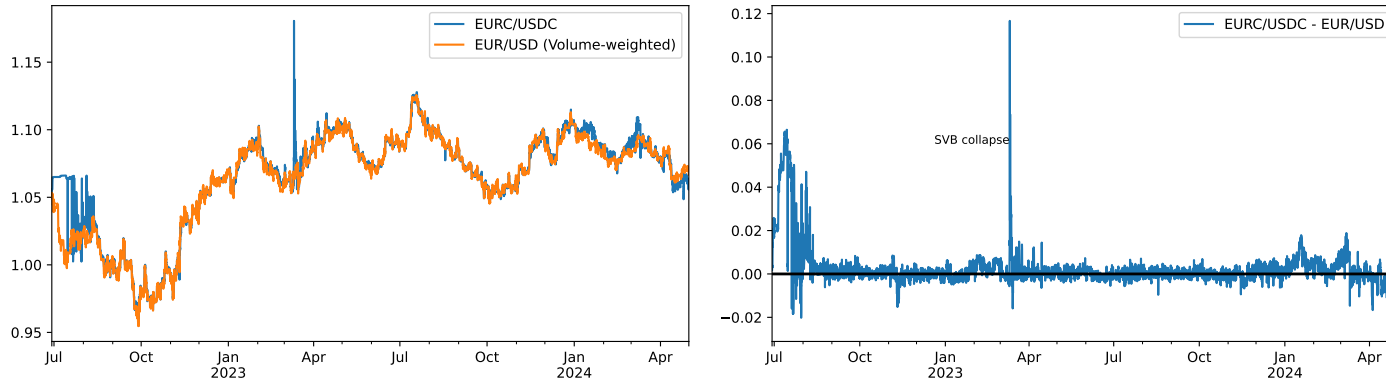


Figure 4: EURC/USDC Prices

Note: This figure plots EURC/USDC and EUR/USD prices at the hourly frequency. EURC/USDC prices are sourced from Uniswap V3, and EUR/USD prices are sourced from CLS. Panel (a) shows the EURC/USDC price, the benchmark EUR/USD price, and their cross-market price difference. Panel (b) reports cumulative blockchain order flow and prices in the EURC/USDC market. The left subpanel shows aggregate order flow, prices, and liquidity provision in EURC and USDC, while the right subpanel disaggregates cumulative order flow by trader group, including sophisticated traders (Top 10 wallets), primary dealers (PM), liquidity providers (LPs), their intersections, and all remaining traders. The sample period for Panel (a) is 28 June 2022 to 30 April 2024, and for Panel (b) is 15 August 2022 to 30 April 2024.

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



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

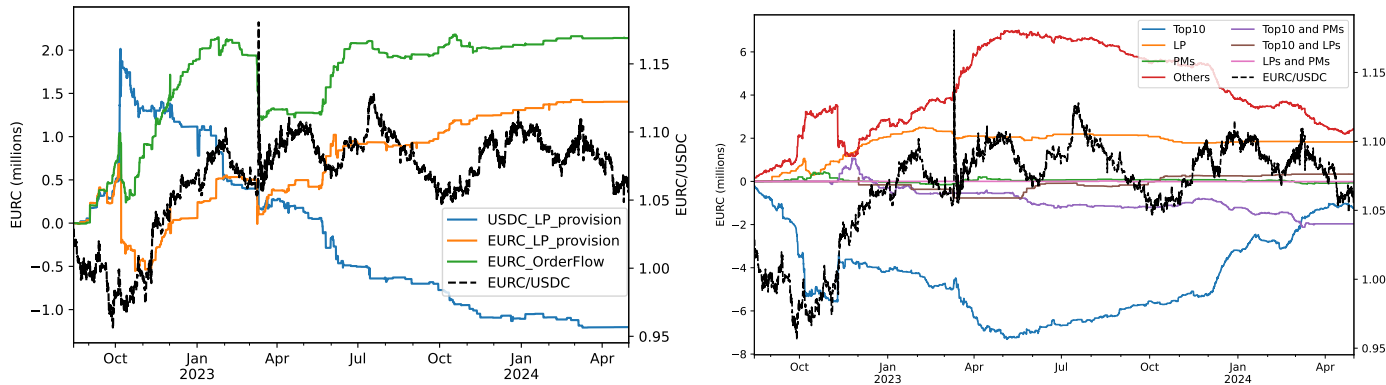



Figure 5: Hourly FX Trading Volume

Note: This figure plots trading volume at the hourly frequency. Panel (a) reports trading volume on Uniswap V3 for the EURC/USDC market, expressed in thousands of EURC. Panel (b) reports trading volume on CLS for the EUR/USD market, disaggregated by counterparty sector (bank–bank, bank–fund, bank–corporate, and non-bank financial–bank) and expressed in EUR millions. The sample period is 15 August 2022 to 30 April 2024.

Panel (a): DEX Trading Volume



Panel (b): CLS Trading Volume

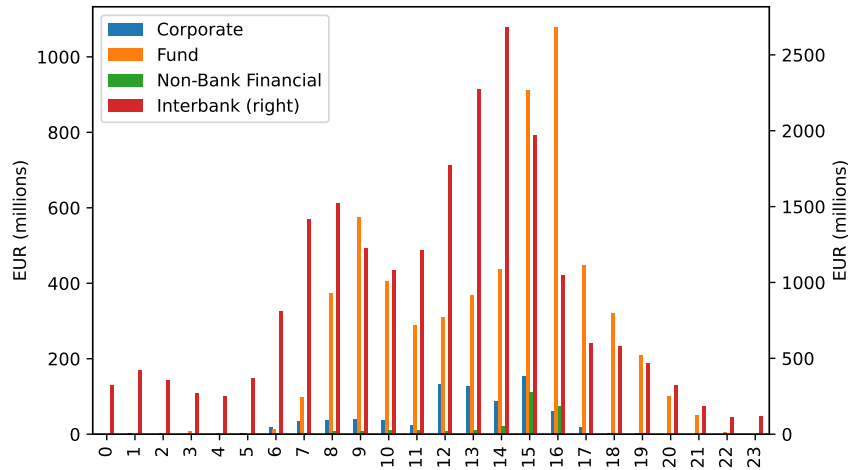
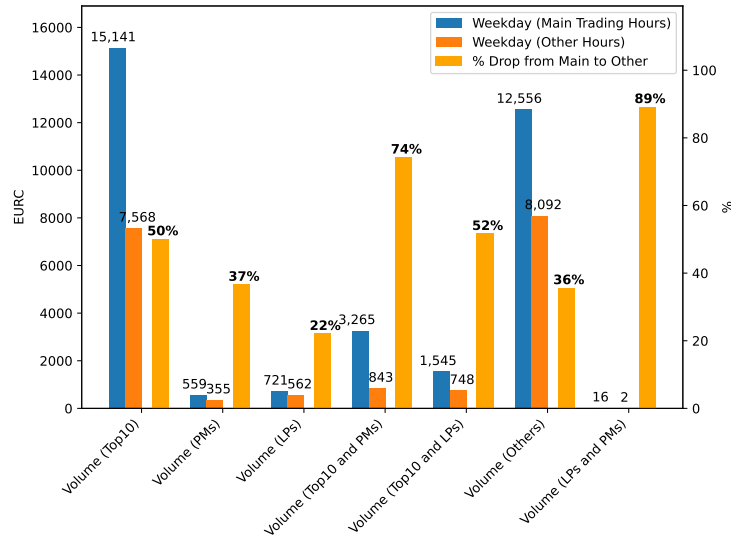


Figure 6: Weekend and Weekday Volume by Trader type

Note: This figure plots average trading volume, distinguishing between weekday and weekend trading by trader group. Panel (a) compares weekday trading volume during traditional primary-market opening hours (13:00–16:00 UTC) with other weekday hours. Panel (b) compares average trading volume across weekdays and weekends. Volumes are expressed in EURC. Blockchain trading volume is disaggregated by trader group following Table 2. The sample period is 15 August 2022 to 30 April 2024.

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



Panel (b): Weekday vs Weekend trading





Figure 7: Trading Costs

Note: This figure presents trading cost measures for the EURC/USDC market at the hourly frequency. Panel (a) reports the interquartile range of total transaction costs across trader groups. Total costs combine gas fees (Ethereum transaction fees converted to USD), private fees (payments to validators for privately routed transactions), liquidity provider fees (5 basis points for the EURC/USDC pool), and price slippage, all expressed in basis points. Panel (b) decomposes median transaction costs into these components by trader group. Trader classifications follow Table 2. The sample period is 15 August 2022 to 30 April 2024.

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



Panel (b): Trading Costs: Median Decomposition




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

Note: This figure plots cumulative blockchain order flow around the USDC de-pegging event at the hourly frequency. Price deviations are measured as the difference between EURC/USDC prices from Uniswap V3 and EUR/USD prices from CLS. Order flow (*OF*) measures net buyer transactions of EURC and is sourced from Uniswap V3 trade data. Cumulative order flow is disaggregated by trader group following Table 2. The sample period is 10 March 2023 to 12 March 2023.


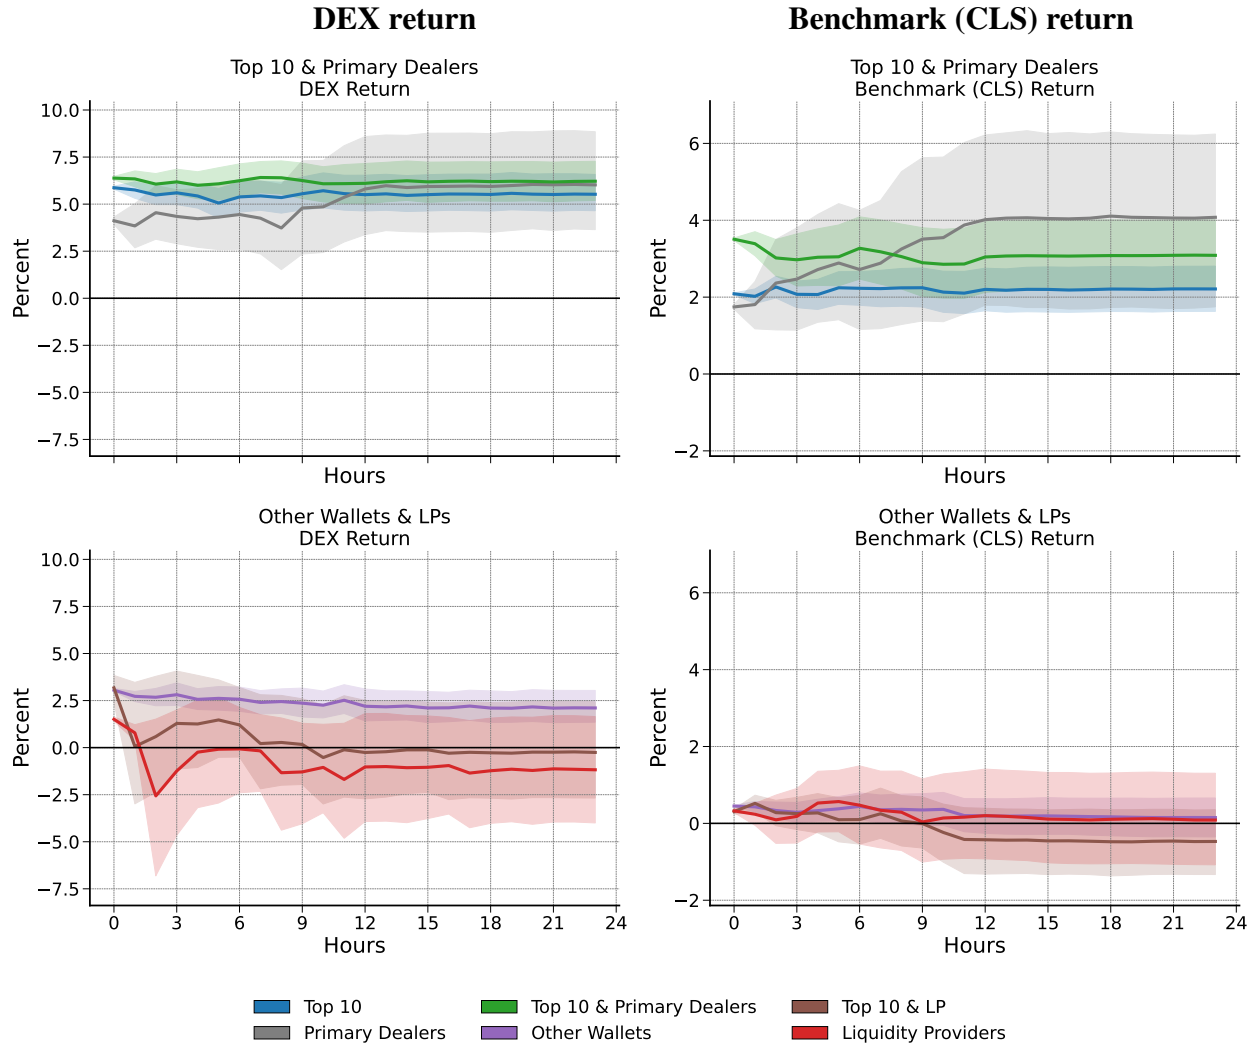


Figure 9: Price impact of blockchain order flow by trader groups

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



Tables

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

Note: AMM denotes automated market maker and CLS denotes Continuous Linked Settlement.

Dimension	Traditional FX Market	Blockchain-Based Market (Uniswap)
Market structure	OTC dealer market with inter-dealer and dealer–client tiers.	Decentralized exchange protocol using algorithmic smart contracts and liquidity pools.
Price formation	Dealer quotes reflect order flow, inventories, and competition across venues.	Prices determined mechanically by the AMM invariant (e.g., $xy = k$) or concentrated-liquidity curves in V3.
Market makers / liquidity providers	Dealers supply liquidity, manage inventories, face adverse selection, and charge bid–ask spreads; balance-sheet regulation shapes capacity.	LPs supply token pairs over price ranges, earn fees, and bear inventory and adverse-selection risk (impermanent loss); constraints stem from volatility, range choice, and gas costs.
Trading costs and frictions	Execution costs reflect spreads, credit limits, capital charges, and latency.	Execution costs driven by gas fees, slippage, priority rules, and network congestion.
Issuance	Fiat currency supply is determined by sovereign monetary policy. FX trading occurs entirely in secondary OTC markets.	Stablecoins are minted and redeemed at par, creating a direct arbitrage link between issuance and secondary-market prices.
Settlement	Typically T+2 via CLS or correspondent banking networks.	Atomic settlement at the blockchain’s block frequency.
Transparency and data	Data are proprietary and aggregated; dealer identities and trade direction are often unobserved.	All transactions are public at the wallet level, with clear trade direction and real-time data availability.

Table 2: Trader classification

Note: Panel (a) reports summary statistics for the number of transactions (Tx) by trader group and transactions per unique address ($Tx/N_{addresses}$). Panel (b) reports summary statistics for EURC volume per transaction. Wallets are classified as sophisticated traders (Top10), primary dealers (PM), and liquidity providers (LP), with intersections reported where applicable; the residual group comprises wallets not in these categories. The sample period is 15 August 2022 to 30 April 2024.

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
$Top10 \cap PM$	✓	✓	×	6	534	89.00
$Top10 \cap LP$	✓	×	✓	7	249	35.57
$PM \cap LP$	×	✓	✓	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
$Top10 \cap PM$	26,373	10,664	100	20,000	25,000	30,000	95,990
$Top10 \cap LP$	44,665	62,339	100	4,290	31,212	50,000	343,333
$PM \cap LP$	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

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

Note: Panel (a) reports EUR/USD trading volume from CLS (EUR billions), shown in aggregate and by counterparty category. Panel (b) reports EURC/USDC trading volume from Uniswap, disaggregated by trader group as defined in Table 2. Panel (c) reports prices and additional blockchain and FX variables, including ETH returns and volatility and an index of gas fees. All variables are measured at the daily frequency. The sample period is 15 August 2022 to 30 April 2024.

	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								
PEURC/USDC	625	1.067	0.035	0.962	1.058	1.078	1.091	1.128
PEUR/USD	625	1.066	0.035	0.960	1.058	1.077	1.090	1.124
$ \text{PEUR/USD} - \text{PEURC/USDC} $ (bps)	625	23.991	25.633	0.267	8.100	16.119	29.710	275.410
σ_{ETH} (%)	625	0.656	0.197	0.337	0.494	0.645	0.770	1.255
GasFee	625	6.055	4.533	0.972	3.082	4.786	7.427	29.459
R_{ETH} (%)	624	0.075	3.158	-19.415	-1.256	0.035	1.525	16.571

Table 4: DEX and CLS Volume correlations

Note: This table reports regressions of CLS volume on DEX trading volume. DEX volume aggregates EURC buy and sell transactions from Uniswap V3, is measured in EURC, and is disaggregated by trader group following Table 2. CLS volume is measured in EUR millions and reported in aggregate and by counterparty category. Regressions are estimated at the hourly frequency. Newey–West (HAC) standard errors are reported in parentheses. ***, ** and * denote significance at the 1, 5, and 10 percent levels, respectively. The sample period is 15 August 2022 to 30 April 2024.

	V_{top10}	V_{PM}	V_{LP}	$V_{top10 \cap PM}$	$V_{top10 \cap LP}$	$V_{LP \cap PM}$	$V_{\#top10, PM, LP}$
	1	2	3	4	5	6	7
Interbank	4.3478*** (0.7699)	0.1984*** (0.0413)	0.3286** (0.1309)	0.8337*** (0.1027)	0.4106* (0.2480)	-0.0001 (0.0006)	3.2545*** (0.7452)
Corporate-Bank	1.5545 (1.5778)	-0.0026 (0.1899)	0.3532 (0.3098)	0.5860* (0.3324)	-0.4185** (0.1945)	-0.0018 (0.0013)	2.2923 (1.9787)
Fund-Bank	1.1120*** (0.3931)	0.0353 (0.0308)	0.0166 (0.0419)	0.2303*** (0.0651)	0.0369 (0.0844)	0.0017 (0.0017)	0.9016*** (0.3087)
Non-Bank Financial-Bank	2.3239 (3.7332)	0.3554 (0.3012)	-0.0312 (0.1768)	0.7064 (0.6986)	0.0518 (0.1026)	-0.0002 (0.0002)	6.8670 (7.7577)
constant	3261.9288*** (529.2313)	113.7215*** (35.7043)	190.3928** (92.8885)	111.9940 (68.7749)	379.6192*** (136.5073)	2.7742 (2.3496)	4390.3679*** (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

Table 5: Determinants of EURC-USDC Peg Deviations

Note: This table reports OLS regressions of absolute peg deviations $|p_{EURC/USDC} - p_{EUR/USD}|$. Gas fees measure average transaction costs on the Ethereum network (USD). σ_{ETH}^{IV} denotes 30-day implied volatility for Ether, and R_{ETH} is the return on ETH. VLOOP is the standardized first principal component of G10 triangular arbitrage violations, and ICRF is the intermediary capital risk factor of He et al. (2017). Stablecoin peg deviations are measured relative to fiat reference values. Variables are measured at the daily frequency, and peg deviations, returns, and volatility measures are expressed in basis points. Column (6) re-estimates the full specification in column (5) excluding the four days surrounding the USDC de-pegging event of 10–13 March 2023 (Silicon Valley Bank failure). Newey–West (HAC) standard errors are reported in parentheses. ***, ** and * denote significance at the 1, 5, and 10 percent levels, respectively. The sample period is 15 August 2022 to 30 April 2024.

EURC/USDC – EUR/USD Peg Deviations						
	1	2	3	4	5	6
gasfee	1.3283*** (0.5036)				1.1787** (0.4718)	1.1731** (0.4717)
σ_{ETH}^{IV}		0.1349* (0.0803)			0.1582** (0.0723)	0.1620** (0.0714)
R_{ETH}		0.0050 (0.0042)			0.0011 (0.0031)	0.0011 (0.0031)
$ p_{USDC/USD} - 1 $			0.6438*** (0.0931)		0.7532*** (0.0272)	1.8449** (0.8712)
$ p_{EURC/EUR} - 1 $			0.1797 (0.1137)			
VLOOP				-1.0602 (0.8011)	-0.3698 (0.7163)	-0.4074 (0.7152)
ICRF				32.9633 (111.2642)	49.2961 (98.4555)	49.6054 (100.3747)
constant	15.9478*** (2.9195)	15.0894*** (5.3608)	20.6290*** (2.1497)	23.9748*** (1.8659)	5.1853 (5.9025)	3.9709 (5.6597)
R-squared	0.0552	0.0139	0.2805	0.0019	0.2294	0.0789
No. observations	625	624	429	625	624	620

Table 6: Determinants of EURC/USDC Order Flow

Note: This table reports regressions of EURC order flow on lagged price differences between DEX and CLS exchange rates. Order flow (OF) measures net buyer transactions of EURC, is measured in EURC millions, and is sourced from Uniswap V3. Order flow is disaggregated by trader group following Table 2. $P_{DEX} - P_{CLS}$ denotes the DEX–CLS price difference. Regressions are estimated at the hourly frequency. Newey–West (HAC) standard errors are reported in parentheses. ***, ** and * denote significance at the 1, 5, and 10 percent levels, respectively. The sample period is 15 August 2022 to 30 April 2024.

	$OF_{top10,t}$	$OF_{PM,t}$	$OF_{LP,t}$	$OF_{top10 \cap PM,t}$	$OF_{top10 \cap LP,t}$	$OF_{LP \cap PM,t}$	$OF_{\#top10,PM,LP,t}$
	1	2	3	4	5	6	7
$P_{DEX,t-1} - P_{CLS,t-1}$	-0.1454** (0.0609)	-0.0097 (0.0072)	-0.0207* (0.0121)	-0.1374*** (0.0297)	-0.0032 (0.0085)	-0.0003 (0.0002)	-0.2247*** (0.0354)
DEXReturn $_{t-1}$	-0.0077** (0.0035)	-0.0002 (0.0002)	0.0003 (0.0006)	-0.0012 (0.0010)	0.0002 (0.0002)	-0.0000 (0.0000)	-0.0008 (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)	
$OF_{\#top10,PM,LP,t-1}$							0.1332 (0.0863)
constant	0.0001 (0.0002)	0.0000 (0.0000)	0.0001** (0.0001)	-0.0000 (0.0001)	0.0000 (0.0000)	-0.0000 (0.0000)	0.0003* (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

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, focusing on Treasury transactions and their implications for stablecoin supply.
2. Appendix **B** outlines Uniswap V3 liquidity mechanics, including tick-level aggregation, virtual reserves, and the construction of net liquidity measures from mint and burn flows.
3. Appendix **C** reports detailed statistics on trader activity and liquidity provision, including volumes, trading frequency, participant types, tick-level liquidity distribution, intraday adjustments, and JIT behavior.
4. Appendix **D** provides supplementary evidence on market efficiency and the construction of arbitrage bounds between EURC/USDC and EUR/USD markets.
5. Appendix **E** presents transaction-level evidence from the March 10–12, 2023 USDC de-pegging episode and event-study responses to Federal Reserve announcements.
6. Appendix **F** details the SVAR identification strategy, including recursive Cholesky ordering and matrix construction for OTC and blockchain order flows.
7. Appendix **G** reports robustness tests, including intraday price impact, liquidity controls, and evidence on JIT liquidity provision.
8. Appendix **H** analyzes the relative roles of private information and arbitrage in explaining permanent price impact.

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.²⁸ The supply of USDC and EURC is shown in Figure A.1. 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.²⁹

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)

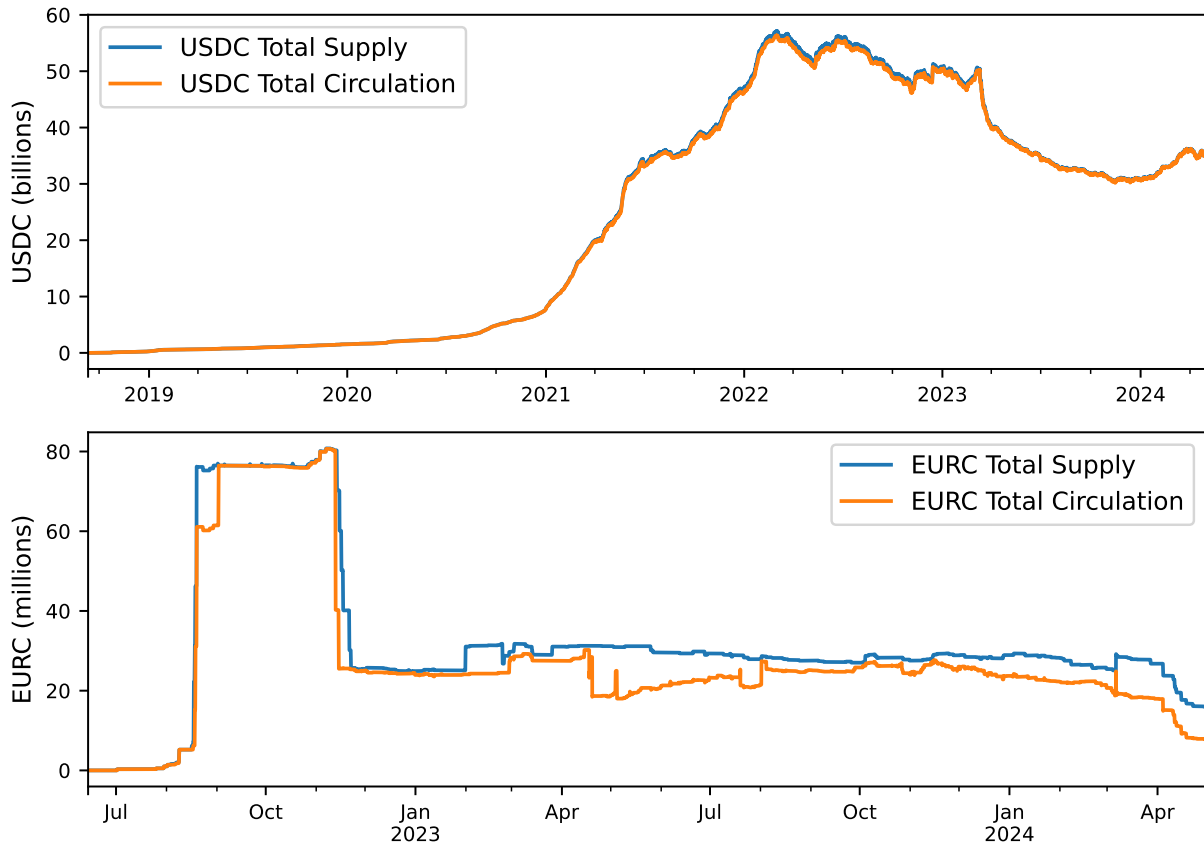
²⁸The USDC Treasury address we use to retrieve the transaction history is "0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48". The EURC Treasury address is "0x1abaea1f7c830bd89acc67ec4af516284b1bc33c"

²⁹One 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.

from the Treasury to the secondary market.

Figure A.1: Primary Market Issuance

Note: This figure plots the total supply of USDC and EURC, together with the amounts in circulation (net of Treasury holdings). The top panel reports the total supply of USDC and the bottom panel reports the total supply of EURC. For the main sample shown in the top and bottom panels, the sample period is 28 June 2022 to 30 April 2024. For the longer-run series, we extend the sample back to the initial issuance dates: 10 September 2018 for USDC and 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 $[i, i + l]$,

$$L_{EURC,i} = \frac{L_i}{\sqrt{z_i}} - \frac{L_i}{\sqrt{p_{i+l}}}, \quad (12)$$

$$L_{USDC,i} = L_i(\sqrt{z_i} - \sqrt{p_i}), \quad (13)$$

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 governing trades within the active tick. Prices are computed using *virtual reserves* rather than actual token balances. Let x denote the quantity of EURC and y the quantity of USDC. The price of EURC in USDC is

$$P_{x/y} \equiv P = \frac{L_{USDC,v}}{L_{EURC,v}}, \quad (14)$$

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

$$L_{EURC,v} = \frac{L}{\sqrt{P}}, \quad L_{USDC,v} = L\sqrt{P}, \quad L_{EURC,v}L_{USDC,v} = L^2. \quad (15)$$

For a liquidity provider supplying liquidity over a price interval $[P_a, P_b]$, the liquidity

parameter L relates to the *actual* token balances (x, y) held in the pool at price P according to

$$L = \frac{x + y}{\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right) + \left(\sqrt{P} - \sqrt{P_a}\right)}. \quad (16)$$

A trade that buys $\Delta x > 0$ units of EURC using USDC—reducing the pool’s EURC balance and increasing its USDC balance—satisfies

$$(L_{EURC,v} - \Delta x)(L_{USDC,v} + \Delta y) = L^2, \quad (17)$$

where $\Delta y > 0$ denotes the corresponding increase in USDC. The post-trade price is therefore

$$P' = \frac{L_{USDC,v} + \Delta y}{L_{EURC,v} - \Delta x} = \frac{L^2}{(L_{EURC,v} - \Delta x)^2}. \quad (18)$$

For small trades, slippage is approximately proportional to $\Delta x/L_{EURC,v}$, implying that a larger L reduces price impact locally.

Virtual and Actual Reserves. At price P , the actual token balances associated with a liquidity position spanning $[P_a, P_b]$ are

$$x = L\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right), \quad y = L\left(\sqrt{P} - \sqrt{P_a}\right). \quad (19)$$

By contrast, the virtual reserves $(L_{EURC,v}, L_{USDC,v})$ depend only on (L, P) and not on the liquidity bounds (P_a, P_b) . They represent notional reserves that replicate the constant-product pricing curve locally at price 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, virtual reserves exceed actual balances, providing greater effective depth at the prevailing price.

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

$$L_{\text{narrow}} = \frac{210}{\left(\frac{1}{\sqrt{1.10}} - \frac{1}{\sqrt{1.2}}\right) + \left(\sqrt{1.10} - \sqrt{1.0}\right)} = \frac{210}{(0.953 - 0.913) + (1.049 - 1.000)} = \frac{210}{0.089} \approx 2,349. \quad (20)$$

The virtual reserves at the initial price are

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

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

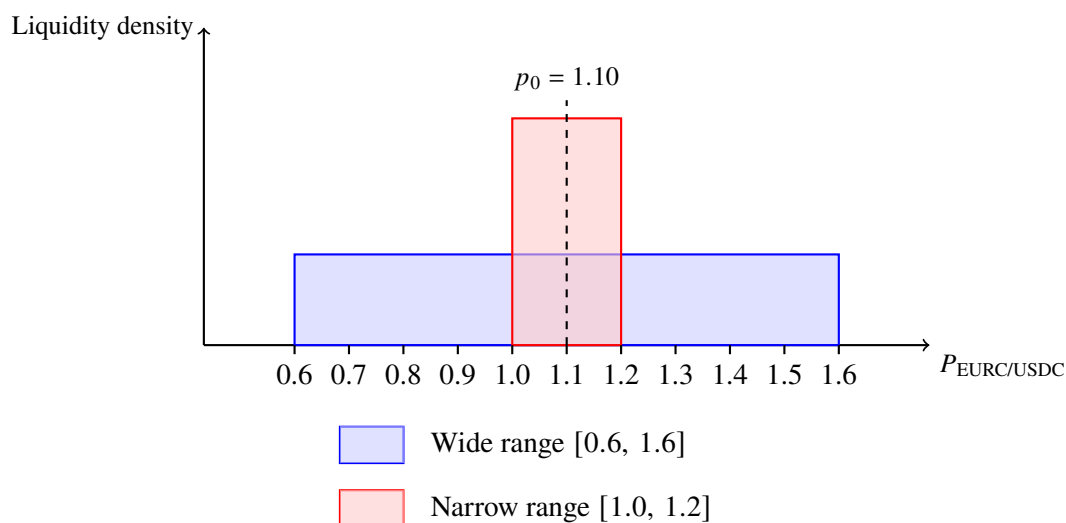
$$P'_{\text{wide}} = \frac{480.4^2}{(458.1 - 5)^2} \approx 1.124, \quad \text{slippage} \approx 2.22\%.$$

For comparison, in the V2 pool with $(x, y) = (100, 110)$, virtual and actual reserves coincide. A trade of 5 EURC gives $P'_{V2} \approx 1.219$, 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 B.1: Liquidity distribution along the price axis

Note: 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 measure liquidity changes separately on the ask (EURC) and bid (USDC) sides of the pool. Let Δx denote the amount of EURC added to (or removed from) the pool and Δy the corresponding amount of USDC.

Net liquidity added in block k is defined as the difference between ask- and bid-side

contributions,

$$mint_{(k)} = mint_{(k)}^{ask} - mint_{(k)}^{bid}, \quad (21)$$

where

$$mint_{(k)}^{ask} = \Delta x \cdot P_{x/y}, \quad (22)$$

$$mint_{(k)}^{bid} = \Delta y. \quad (23)$$

Here, $P_{x/y}$ denotes the contemporaneous price of EURC in USDC, so that ask-side liquidity is expressed in USDC units. When multiple mint events occur within the same block, contributions are summed prior to computing $mint_{(k)}$.

Net liquidity withdrawn is defined analogously,

$$burn_{(k)} = burn_{(k)}^{ask} - burn_{(k)}^{bid}, \quad (24)$$

with

$$burn_{(k)}^{ask} = \Delta x \cdot P_{x/y}, \quad (25)$$

$$burn_{(k)}^{bid} = \Delta y, \quad (26)$$

where Δx and Δy now denote token amounts removed from the pool. A positive value of $mint_{(k)}$ or $burn_{(k)}$ therefore reflects a relative shift toward ask-side (EURC) liquidity.

Distance from the Current Price. Liquidity provision in Uniswap V3 is range-specific, so we classify each mint and burn by its proximity to the prevailing market price. Let i_M denote the active tick at block k . For a liquidity position spanning ticks $[i_a, i_b]$, we compute the distance of both endpoints from i_M .

Positions within $\pm 1.0\%$ of the current price—corresponding to approximately ten tick intervals around i_M in the EURC/USDC pool—are classified as *best*, while positions outside this

band are classified as *away*. We further distinguish whether liquidity lies above the current price (ask side) or below it (bid side).

Some positions span both regions. For each mint event by LP i over a price interval $[P_a, P_b]$, we first compute the total liquidity parameter implied by the provider's token balances (x_i, y_i) ,

$$L_i = \frac{x_i + y_i}{\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}}\right) + \left(\sqrt{P} - \sqrt{P_a}\right)}. \quad (27)$$

This L_i measures the effective liquidity supplied by LP i across the full range.

When $[P_a, P_b]$ overlaps both the best and away regions, we partition L_i across regions using the Uniswap V3 price–tick mapping, allocating liquidity in proportion to the token quantities implied by each sub-interval.³⁰

After this decomposition, we aggregate across all LPs active in block k to obtain total liquidity additions and withdrawals within each region. For example, ask-side liquidity added within the best and away regions is

$$\text{mint}_{(k)}^{b,ask} = \sum_i \Delta x_i^{(b)} P_{x/y}, \quad \text{mint}_{(k)}^{a,ask} = \sum_i \Delta x_i^{(a)} P_{x/y}, \quad (28)$$

with analogous expressions for the bid side and for burn events.

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

$$\text{Liquidity}_{(k)}^{net} = \text{mint}_{(k)} - \text{burn}_{(k)}. \quad (29)$$

³⁰For example, if the current price is $P_M = 1.10$ and the best region is $[1.089, 1.111]$, an LP providing liquidity over $[1.08, 1.12]$ contributes to both regions. The corresponding EURC quantities satisfy

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

with analogous expressions for $\Delta y_i^{(b)}$ and $\Delta y_i^{(a)}$. By construction, $\Delta x_i^{(b)} + \Delta x_i^{(a)} = \Delta x_i$ and $\Delta y_i^{(b)} + \Delta y_i^{(a)} = \Delta y_i$.

A positive value indicates that, on net, more EURC-side liquidity has been added than withdrawn in block k .

We decompose net liquidity flows by distance from the market price. For each block k , we compute

$$Liquidity_{(k)}^{net,b} = mint_{(k)}^b - burn_{(k)}^b, \quad Liquidity_{(k)}^{net,a} = mint_{(k)}^a - burn_{(k)}^a.$$

Block-level flows are then aggregated to the hourly frequency,

$$Liquidity_{t,h}^{net,b} = \sum_{k \in h} Liquidity_{(k)}^{net,b}, \quad Liquidity_{t,h}^{net,a} = \sum_{k \in h} Liquidity_{(k)}^{net,a}. \quad (30)$$

These series measure net additions or withdrawals of EURC-side liquidity relative to the USDC side, separately for liquidity supplied near and away from the prevailing market price, and constitute the liquidity controls used 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 supplementary appendices that provide variable definitions, construction details, and descriptive statistics for trading and liquidity provision on Uniswap V3.

Glossary and Example Transactions (Appendix C.1). This appendix provides a glossary of all on-chain variables used in the analysis, illustrated through example swap and liquidity (mint) transactions from the EURC/USDC pool. It explains how execution prices, pool 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 prices and transaction-level slippage.

Trading and Liquidity Provision Statistics (Appendix C.2). This appendix reports descriptive statistics on trading activity, wallet characteristics, and liquidity provision. It documents the distribution of trading volume across wallets, the concentration of liquidity provision among LPs, wallet classification by activity patterns, and the prevalence of private transactions that bypass the public mempool.

Liquidity Provision and Inventory Management (Appendix C.3). The final appendix focuses on liquidity provision and LP behavior. It reports statistics on active LPs, aggregate liquidity, and its concentration across providers, and examines mint and burn activity. It also documents how large LPs adjust liquidity positions following order-flow imbalances by withdrawing and re-minting liquidity across price ranges, consistent with passive inventory management.

C.1 Glossary: Example Swap and Liquidity Transaction

C.1.1 Swap Transaction

Table C.1: Example Swap Transaction and Variable Glossary: EURC/USDC on Uniswap V3

Note: amount0 and amount1 are reported from the pool’s perspective. A negative EURC value indicates that the trader buys EURC and removes it from the pool, increasing the USDC balance; a positive value indicates that the trader sells EURC to the pool. Prices are derived from sqrtPriceX96 as $P = \left(\frac{\text{sqrtPriceX96}}{2^{96}}\right)^2$ (for equal token decimals), which is equivalent to the tick-based representation $P = 1.0001^{\text{tick}}$.

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.

C.1.2 Liquidity Provision Transaction

Table C.2: Example Mint Transaction and Variable Glossary: EURC/USDC on Uniswap V3

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


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.

C.2 Trading Statistics

C.2.1 Trading Volume

Figure C.1: Summary statistics of trading volume

Note: This figure reports monthly summary statistics of trading activity, including the number of active addresses, total trading volume, and the share of trading volume attributable to sophisticated traders (top 10 wallets). The sample period is 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, focusing on wallet age, the number of tokens transferred, and average transaction frequency. We compare these characteristics across seven mutually exclusive trader groups, including sophisticated traders, primary dealers (PMs), LPs, and combinations thereof. Table C.3 reports the corresponding summary statistics.

Sophisticated traders exhibit relatively old wallets (median age 742 days) and the highest activity levels, with a median of 0.68 transactions per day and 54 tokens transferred. Primary dealers display slightly younger wallets (median age 613 days) and lower token transfer activity (median of 20), but remain relatively active, with a median transaction frequency of 0.58 per day. LPs have the oldest wallets in the sample (median age 813 days) but transact less frequently (0.32 transactions per day), consistent with a more passive or liquidity-provision-oriented role, despite moderate token transfers (median of 30).

Wallets belonging to multiple trader categories tend to be substantially more active than single-category wallets. In particular, sophisticated traders who also act as LPs show very high token transfer counts (median of 88) and transaction frequency (1.13 per day), indicating intensive on-chain engagement. By contrast, wallets not classified into any of the main trader types are the least active, with low transaction frequency (0.24 per day) and few tokens transferred (median of 14).

Overall, these statistics highlight clear heterogeneity in on-chain behavior across trader types. Sophisticated traders and PMs are characterized by higher activity and relatively mature wallets, while LPs operate with lower transaction frequency, consistent with a more passive trading role.

C.2.3 Persistence of Wallet Classifications

Table C.4 reports the persistence of wallet classifications over the sample, measured by the share of months in which an address is active within a given group. Liquidity providers (LPs) are

highly stable: 86 of 90 addresses (96%) are present in every month of their trading history, and all LPs appear in at least half of the months. Primary dealers (PMs) display similar persistence, with 62 of 68 addresses (91%) trading throughout the sample.

In contrast, Top10 wallets exhibit substantially higher entry and exit. Only 17 of 76 addresses (22%) are present in all months, while 24 (32%) appear in at most 30% of months. Intersection groups are generally stable: all $PM \cap LP$ wallets are observed continuously, and 3 of 7 $Top10 \cap LP$ wallets trade in every month. Addresses outside the main categories are also highly persistent, with 2283 of 2342 wallets (98%) present throughout.

Weighting by group size, approximately 95% of addresses remain active over the full sample, indicating that the classification scheme is broadly stable, with turnover concentrated primarily among Top10 wallets (see Figure C.2).

Table C.3: Blockchain characteristics by address type

Note: This table reports summary statistics of blockchain characteristics, including wallet age (days since first observed trading), the number of tokens transferred, and trading frequency (transactions per day). Statistics are reported for seven trader groups. The sample period is 15 August 2022 to 30 April 2024.

Panel (a): Sophisticated traders								
	count	mean	std	min	0.25	0.50	0.75	max
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								
	count	mean	std	min	0.25	0.50	0.75	max
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								
	count	mean	std	min	0.25	0.50	0.75	max
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								
	count	mean	std	min	0.25	0.50	0.75	max
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								
	count	mean	std	min	0.25	0.50	0.75	max
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								
	count	mean	std	min	0.25	0.50	0.75	max
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)	2317	706.86	492.31	1.00	406.00	585.00	944.00	2834.00
Number of Tokens Transferred	2317	64.62	251.78	1.00	4.00	14.00	46.00	7631.00
Frequency (transactions per day)	2317	2.28	16.88	0.00	0.07	0.24	0.92	558.01


Table C.4: Persistence of wallet classifications

Note: Persistence is measured as the share of months in which an address is active within a given category over the sample period. Columns report the number of wallets present in all months (= 100%), more than 70%, 50%, or 30% of months, and at most 30% of months. Categories are mutually exclusive.

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 \cap PM	6	0	0	3	6	0
Top10 \cap LP	7	3	3	3	5	2
PM \cap LP	3	3	3	3	3	0
$\notin \{Top10, PM, LP\}$	2342	2283	2307	2320	2339	3

Figure C.2: Persistence of Wallet Classifications

Note: This figure reports the distribution of wallet persistence across disjoint bins defined by the fraction of months in which an address is active. Each bar sums to 100% within a wallet category, allowing 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 on-chain inclusion. We classify transactions using Blocknative mempool archives,³¹ based on two criteria. First, the transaction must be confirmed on-chain. Second, the `timePending` value must equal zero, indicating that the transaction did not wait in the public mempool.³²

Private activity is concentrated among Top10 and Top10∩LP addresses, while PMs and LPs transact almost exclusively through the public mempool.

Table C.5 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 small subset of wallets. Among Top10 addresses, 3,270 of 4,439 transactions (74%) are routed privately. Reliance on private submission is even higher for Top10∩LP wallets, 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 about 3% of their trades routed privately. Addresses outside the Top10, PM, or LP categories split their activity more evenly, with roughly 40% of trades submitted privately.

Panel (b) combines statistics on block position and transaction size. Private trades enter blocks substantially earlier, with a median position of 7 compared to 82 for public transactions, consistent with direct relay to validators. Private transactions are also larger: the median private trade size is approximately 11,100 EURC, compared to 6,400 EURC for public trades, and the upper tail is thicker, with maximum sizes exceeding 1 million EURC. Overall, these patterns indicate that private submission is primarily used by large Top10 and Top10∩LP wallets to obtain execution priority and accommodate block-sized trades, while smaller traders and PMs rely on the public mempool.

³¹See Blocknative documentation at <https://docs.blocknative.com/data-archive/mempool-archive>.

³²Blocknative matches the vast majority of confirmed transactions to their mempool records. We identify 17 confirmed transactions for which Blocknative provides no mempool entry. Since these transactions were included on-chain without any evidence of public broadcast, we classify them as private. Such cases arise when transactions are submitted through private routing channels that bypass the public mempool.

Table C.5: Summary statistics of private and public transactions

Note: This table reports descriptive statistics for private and public transactions. Panel (a) reports the number of transactions by account type and the share of private flows. Panel (b) reports statistics on block positions (ordering within blocks) and transaction volumes (in EURC equivalent) per trade. For each metric, separate columns are reported for private and public transactions. The sample period is 15 August 2022 to 30 April 2024.

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%
$\notin \{\text{Top10, PM, LP}\}$	3,642	5,474	9,116	40%


Panel (b): Block positions and volumes per transaction (EURC equivalent)				
	BlockPos _{Private}	BlockPos _{Public}	Vol/TX _{Private}	Vol/TX _{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

C.3 Liquidity Provision Statistics

C.3.1 Summary Statistics

Figure C.3: Summary statistics of liquidity provision

Note: This figure reports monthly summary statistics on liquidity provision. It shows the number of addresses, aggregate liquidity provision, and the share of liquidity supplied by the top five LPs. The sample period is 1 July 2022 to 30 April 2024.

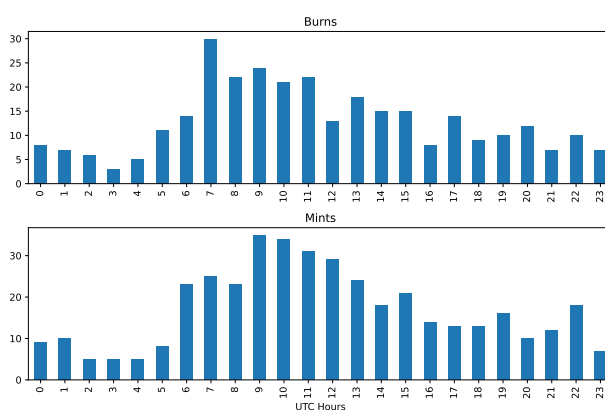


C.3.2 Intra-day patterns


Figure C.4: Intra-day LP Mints and Burns

Note: This figure reports hourly liquidity provision, classified into mints (liquidity additions) and burns (liquidity withdrawals). Panel (a) reports LP transaction counts for mints and burns. Panel (b) reports mint and burn volumes, disaggregated into ‘best’ and ‘away’ regions. Liquidity is classified as ‘best’ when provided within $\pm 1\%$ of the prevailing market price, corresponding to active liquidity at the best bid and ask. Liquidity outside this range is classified as ‘away’ and becomes active only following sufficiently large price movements. The sample period is 15 August 2022 to 30 April 2024.

Panel (a): Number of transactions



Panel (b): Volume



C.3.3 Liquidity Provision as Inventory Management

LPs actively manage inventory in response to order flow that depletes their holdings of a given token. When swap traders purchase EURC from the pool, LP inventories become tilted

toward USDC. To restore target exposure, LPs may withdraw liquidity, purchase EURC through swaps, and re-deposit liquidity at updated price ranges. Table C.6 illustrates these dynamics using selected transactions.

We document this behavior among large LPs during Q4 2022, when the EUR/USD exchange rate appreciated from 0.98 to 1.07. The three most active LPs—wallets *75e6*, *2488*, and *2200*, defined by swap sizes exceeding 10,000 EURC—exhibit a consistent sequence of actions. As order flow reduces their EURC exposure, these LPs withdraw USDC from stale price ranges, swap into EURC, and re-mint liquidity around the prevailing market price, restoring inventory balance and reallocating liquidity symmetrically.

Illustrative example: wallet *75e6* on 22 Nov 2022 Negative values indicate a reduction of the corresponding token in the pool, so negative EURC swap amounts reflect purchases of EURC. On 22 Nov 2022, wallet *75e6* executes a swap of –19,638 EURC at price 1.028, followed by a mint of 19,638 EURC and 11,708 USDC within bounds 1.018–1.044, while simultaneously burning –40,701 USDC from a stale range of 0.976–1.021. The wallet then executes a second swap of –25,765 EURC at the same price and re-mints 25,749 EURC and 17,870 USDC at the updated bounds. This sequence shows how *75e6* withdraws surplus USDC, acquires EURC through swaps, and re-mints liquidity to rebalance its position. Comparable patterns for wallets *2488* and *2200*, reported in Table C.6, follow the same logic with smaller trade sizes.

Table C.6: Transaction Details for Selected Wallets (75e6, 2488, 2200)

Note: This table reports transactions for wallets ending 75e6 (0x33e0...75e6), 2488 (0xe75b...2488), and 2200 (0x3231...2200). For swap rows, a negative amount indicates a purchase of EURC; a positive amount indicates a sale. Mint and burn rows report net liquidity changes and price bounds. The sample period is October 2022 to May 2023.

Date (UTC)	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	19493.647	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	14630.775	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
2022-11-11 19:16	mint	2488	14424.185	11632.420	1.029	1.031	
2022-11-11 19:18	mint	2488	4181.239	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	

Continued on next page

Table C.6: Transaction Details for Selected Wallets (cont.)

Date (UTC)	Type	User	EURC	USDC	Lower Price	Upper Price	Price
2022-11-05 11:54	swap	2200	-6290.000				0.996
2022-11-05 11:56	mint	2200	6295.919	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	4570.372	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	1345.637	12083.623	1.000	1.097	
2023-05-02 09:23	burn	2200	0.000	-13553.126	1.000	1.097	

Appendix D: Market Efficiency

26

D.1 Volume Correlations

Table D.1: Correlation Matrix of Trading Volume: Liquid Trading Hours vs Off Hours

Note: This table reports pairwise Pearson correlation coefficients of trading volume across decentralized and traditional FX market segments. Panel (a) reports correlations during active trading hours (Europe and New York, 07:00–22:00 UTC), while Panel (b) reports correlations during off hours. On-chain volumes correspond to EURC/USDC trading on Uniswap V3 and are disaggregated by participant group: top 10 wallets ($V(\text{top10})$), primary dealers ($V(\text{PM})$), liquidity providers ($V(\text{LP})$), 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. The sample period is 15 August 2022 to 30 April 2024.

Panel (a): Trading Hours (Europe + New York, 07:00–22:00 UTC)												
	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	
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)	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	
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	

D.2 Arbitrage Bounds

This section describes how we construct arbitrage bounds based on observed transaction costs. The bounds compare price differences between decentralized and traditional markets to the costs required to execute an arbitrage trade, including liquidity fees, slippage, gas costs, and validator payments. When observed deviations exceed these costs, we classify them as violations of the arbitrage bound, yielding cost-based measures of the frictions faced by traders in practice.

Finite liquidity and transaction fees limit arbitrage and allow deviations from efficient prices to persist (Barbon and Rinaldo, 2024). A standard way to quantify such 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 $Z \leftrightarrow X$:

$$\Delta = |1 - P_{XY}P_{YZ}P_{ZX}|, \quad (31)$$

where P_{AB} denotes the price of currency A in units of currency B . A triangular trade is profitable only if Δ exceeds transaction costs, including liquidity fees, slippage, gas costs, and validator payments on private transactions.

Using 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 (32).³³ The sample for these measures begins on 1 March 2023, reflecting the availability of centralized exchange price data. We consider three trading paths: Δ_1 converts EURC to USDC on DEX, to USD on Kraken, and back to EURC via EURC/USD on Coinbase; Δ_2 converts EURC to USDC on DEX, to EUR on Coinbase via USDC/EUR, and back to EURC

³³Centralized 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 USDC/USD market.

via EURC/EUR; and Δ_3 traces a four-currency path EURC \rightarrow USDC \rightarrow USD \rightarrow EUR \rightarrow EURC.

$$\begin{aligned}\Delta_1 &= \left| 1 - \frac{P_{\text{EURC/USDC}} \cdot P_{\text{USDC/USD}}}{P_{\text{EURC/USD}}} \right| \\ \Delta_2 &= \left| 1 - \frac{P_{\text{EURC/USDC}} \cdot P_{\text{USDC/EUR}}}{P_{\text{EURC/EUR}}} \right| \\ \Delta_3 &= \left| 1 - \frac{P_{\text{EURC/USDC}} \cdot P_{\text{USDC/USD}}}{P_{\text{EUR/USD}} \cdot P_{\text{EURC/EUR}}} \right|\end{aligned}\tag{32}$$

These measures capture inefficiencies between decentralized and centralized markets arising from liquidity, gas fees, validator payments, and execution costs. While conceptually distinct, they are closely related to the baseline deviation measure Δ used in the main body of the paper, defined as the absolute price difference between EUR/USD and EURC/USDC.

Panel (a) of Table D.2 summarizes the distribution of triangular arbitrage deviations. Median values are around 0.2 percent, with mean deviations between 0.3 and 0.4 percent—slightly larger than the baseline deviation in the main text, which averages 24 basis points. Large deviations are concentrated in periods of stress, most notably the March 2023 USDC de-pegging episode, when deviations reach 7–8 percent.

Panels (b) and (c) report the incidence of arbitrage-bound violations under alternative cost assumptions. Panel (b) accounts for on-chain frictions, including gas fees (converted to USD), a fixed 5 basis-point liquidity-provider fee, slippage, and private validator fees.³⁴ Under these assumptions, violations occur in 16.1 percent, 17.6 percent, and 18.6 percent of observations for Δ_1 , Δ_2 , and Δ_3 , respectively.

Panel (c) further incorporates off-chain intermediation costs, such as CEX taker fees and OTC bid–ask spreads.³⁵ Including these costs reduces the share of violations to 3.4 percent for Δ_1 , 3.1 percent for Δ_2 , and 4.6 percent for Δ_3 .

³⁴Execution prices are computed from realized token exchanges (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.

³⁵Coinbase charges 60 bps for trades under \$10,000, declining 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. \(2024\)](#).

Overall, 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 interpret these findings as evidence of *constrained efficiency*, reflecting the costs and frictions of decentralized trading and settlement, particularly during stablecoin stress episodes (Gromb and Vayanos, 2002).

Table D.2: Triangular arbitrage conditions and transaction costs: violations of the upper bound

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) reports absolute percentage deviations and the estimated upper bound (Δ_{UB}). Panel (b) reports the share of observations for which the arbitrage metric exceeds estimated on-chain transaction costs, including gas fees (ETH converted to USD), a 0.05% liquidity-provider fee, slippage, and private validator fees. Panel (c) further includes centralized-exchange (CEX) taker fees and bid–ask spreads for traditional FX trades. Binary indicators equal 1 when deviations exceed the corresponding bounds. Gas fees and transaction costs are winsorized at the top 1%. The sample period is 1 March 2023 to 30 April 2024.

	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

Appendix E: Public Information

E.1 USDC De-Pegging Event

This section provides supplementary evidence on trading behavior, liquidity provision, and execution costs during the March 2023 USDC de-pegging episode. Table E.1 reports transaction-level activity of the most active sophisticated investors during the event. The table documents repeated purchases of USDC from centralized exchanges, followed by rapid sales of USDC to decentralized exchange pools, including EURC/USDC, illustrating how a small number of large wallets transmitted USDC selling pressure into on-chain markets.

Table E.2 complements this evidence by reporting liquidity provision and swap activity by LPs over the same period. The transactions show infrequent withdrawals from stale price ranges and re-minting at updated bounds as prices moved sharply, consistent with largely passive LP behavior during the de-pegging episode. Figure E.1 further visualizes these dynamics by presenting tick-level liquidity snapshots for the EURC/USDC pool before, during, and after the de-pegging. As swap traders absorbed EURC reserves, liquidity became highly asymmetric at the peak of the event, with provision concentrated on the USDC side. Following stabilization of the USDC peg, liquidity gradually rebalanced across both sides of the pool.

Finally, Table E.3 reports gas fees across trader types in the full sample and during the de-pegging window. This table provides direct evidence on blockchain execution costs faced by different market participants and shows how gas fees varied across trader types and intensified during the crisis period.

E.1.1 USDC De-Pegging Event: Sophisticated Investor Trades

Table E.1: Transactions of Sophisticated Investor during USDC De-Pegging Event (2023-03-10 to 2023-03-12)

Note: This table reports swap transactions by a sophisticated investor with wallet ID `0xd64137f743432392538a8f84e8e571fa09f21c37`, abbreviated as wallet 1c37, during the USDC de-pegging event from 10–12 March 2023. This wallet was the largest single source of USDC selling pressure. The “From” and “To” columns refer to transfers of USDC. The sample period is 10–12 March 2023.

Date (UTC)	Hash	From	To	USDC	From Name	To Name	Public/Private
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	Public
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	Public
2023/03/10 13:34:59	3601	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 13:43:35	5de7	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 14:11:11	ae67	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 14:24:47	6aa6	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
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	Public
2023/03/10 14:36:59	ebaf	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 14:43:35	021d	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 15:03:47	3c82	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
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	Public
2023/03/10 15:55:11	e414	1c37	02ce	5000	trader	Uniswap V3: USDC-PRIME 2	

Continued...

Date (UTC)	Hash	From	To	USDC	From Name	To Name	Public/Private
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	Public
2023/03/10 19:55:47	c9e0	1c37	73d6	7500	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 21:30:35	5f29	1c37	73d6	10000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 21:34:47	419e	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
2023/03/10 21:40:47	8acd	1c37	73d6	20000	trader	Uniswap V3: EURC-USDC	Public
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	
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	Public
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	Public
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	

Continued...

Date (UTC)	Hash	From	To	USDC	From Name	To Name	Public/Private
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	
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	

E.1.2 USDC De-Pegging Event: Liquidity Provision


Table E.2: Liquidity provision during the USDC de-pegging event

Note: This table reports all transactions by LPs during the USDC de-pegging event from 10–12 March 2023. The “Type” column identifies swaps, mints, and burns. For swaps, a negative EURC value indicates a purchase; for mints/burns, negative values indicate withdrawals. Wallets: 4f35 = 0x767f...4f35; 4f8f = 0xf550...4f8f; ebb3 = 0x2516...ebb3.


Date (UTC)	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

Figure E.1: Snapshots of EURC/USDC Liquidity Around the USDC De-Pegging Event.


Note: 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.




(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

E.1.3 Gas Fees across Trader Types: Full Sample and USDC De-Pegging Event

Table E.3: Gas Fees in USDC per 10,000 EURC Transacted

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 USDC de-pegging period from 10–12 March 2023. Trader groups are defined as in Table 2, including intersection groups such as $\text{Top10} \cap \text{PM}$ and a residual category of traders not belonging to any primary group. Median gas fees are substantially lower for Top10 wallets relative to other groups, particularly during the de-pegging episode.

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	25,916.65
$\text{Top10} \cap \text{PM}$	534	5.10	12.42	0.388	1.594	2.981	5.317	251.50
$\text{Top10} \cap \text{LP}$	249	168.98	643.03	0.203	3.362	15.831	83.927	8,430.81
$\text{PM} \cap \text{LP}$	6	42.83	66.52	4.321	6.983	11.090	40.961	173.97
$\notin \{\text{Top10}, \text{PM}, \text{LP}\}$	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
$\text{Top10} \cap \text{LP}$	1	15.97	–	15.97	15.97	15.97	15.97	15.97
$\notin \{\text{Top10}, \text{PM}, \text{LP}\}$	240	191.04	835.42	0.611	12.547	32.820	102.360	8,665.39

E.2 Monetary Announcements

An efficient market should incorporate public news into prices rapidly. We test this hypothesis by examining how blockchain prices (EURC/USDC) respond to scheduled Federal Reserve announcements, using high-frequency timestamps of FOMC releases at 2pm Eastern Time. We study the intra-day response of the on-chain EURC/USDC price on Uniswap V3 and the off-chain EUR/USD price on CLS.

For each announcement, we center the data on the release time and compute average

responses over a ± 6 hour window. To benchmark normal intra-day variation, we construct a placebo group by aligning the same clock time (2pm ET) on non-announcement weekdays. Figure E.2 reports the aggregate event-study response, while Figure E.3 plots the reaction for each of the 13 announcements between August 2022 and March 2024.

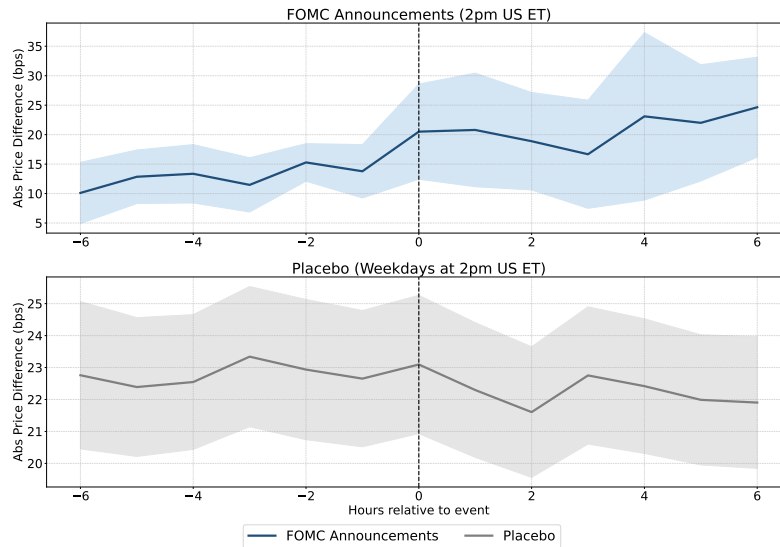
Panel (a) of Figure E.2 shows the absolute price difference between EURC/USDC and EUR/USD. While placebo days provide a baseline for typical intra-day dispersion, announcement days capture deviations around public news arrivals. Panel (b) plots trading volume in both markets and shows a clear spike at the announcement time, consistent with intensified price discovery. Importantly, the EURC/USDC price closely tracks the off-chain EUR/USD response, indicating rapid incorporation of public information on-chain.

To formally assess whether announcement-day price differences differ from normal conditions, we perform Welch t-tests comparing the absolute price difference between announcement and placebo days at each hourly interval. Table E.4 reports the results. Price differences are significantly lower in the pre-announcement window, consistent with anticipatory trading or liquidity provision, but are statistically indistinguishable from placebo values in the post-announcement period. The minimal post-announcement deviations coincide with elevated trading volume at the announcement time, suggesting that increased on-chain activity helps maintain tight price alignment between decentralized and traditional FX markets when public news arrives.

Figure E.2: Event Study Around Federal Reserve Monetary Announcements

Note: This figure presents event studies around Federal Reserve monetary announcements at 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 volume in both markets: DEX trading volume is reported in thousands of EURC, and CLS volume is reported in EUR millions. The placebo group corresponds to weekdays at 2pm ET, excluding announcement days. The sample period is 15 August 2022 to 30 April 2024.

Panel (a): Absolute Price Difference (bps)



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

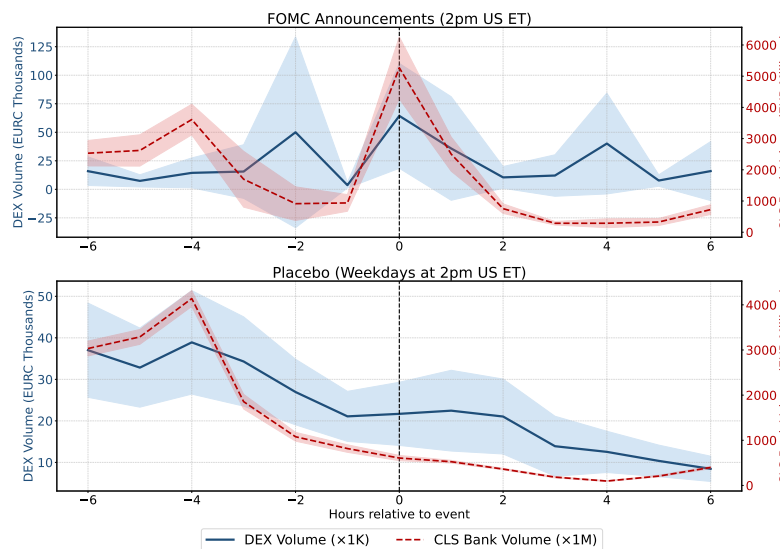


Figure E.3: Federal Reserve Monetary Announcements

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

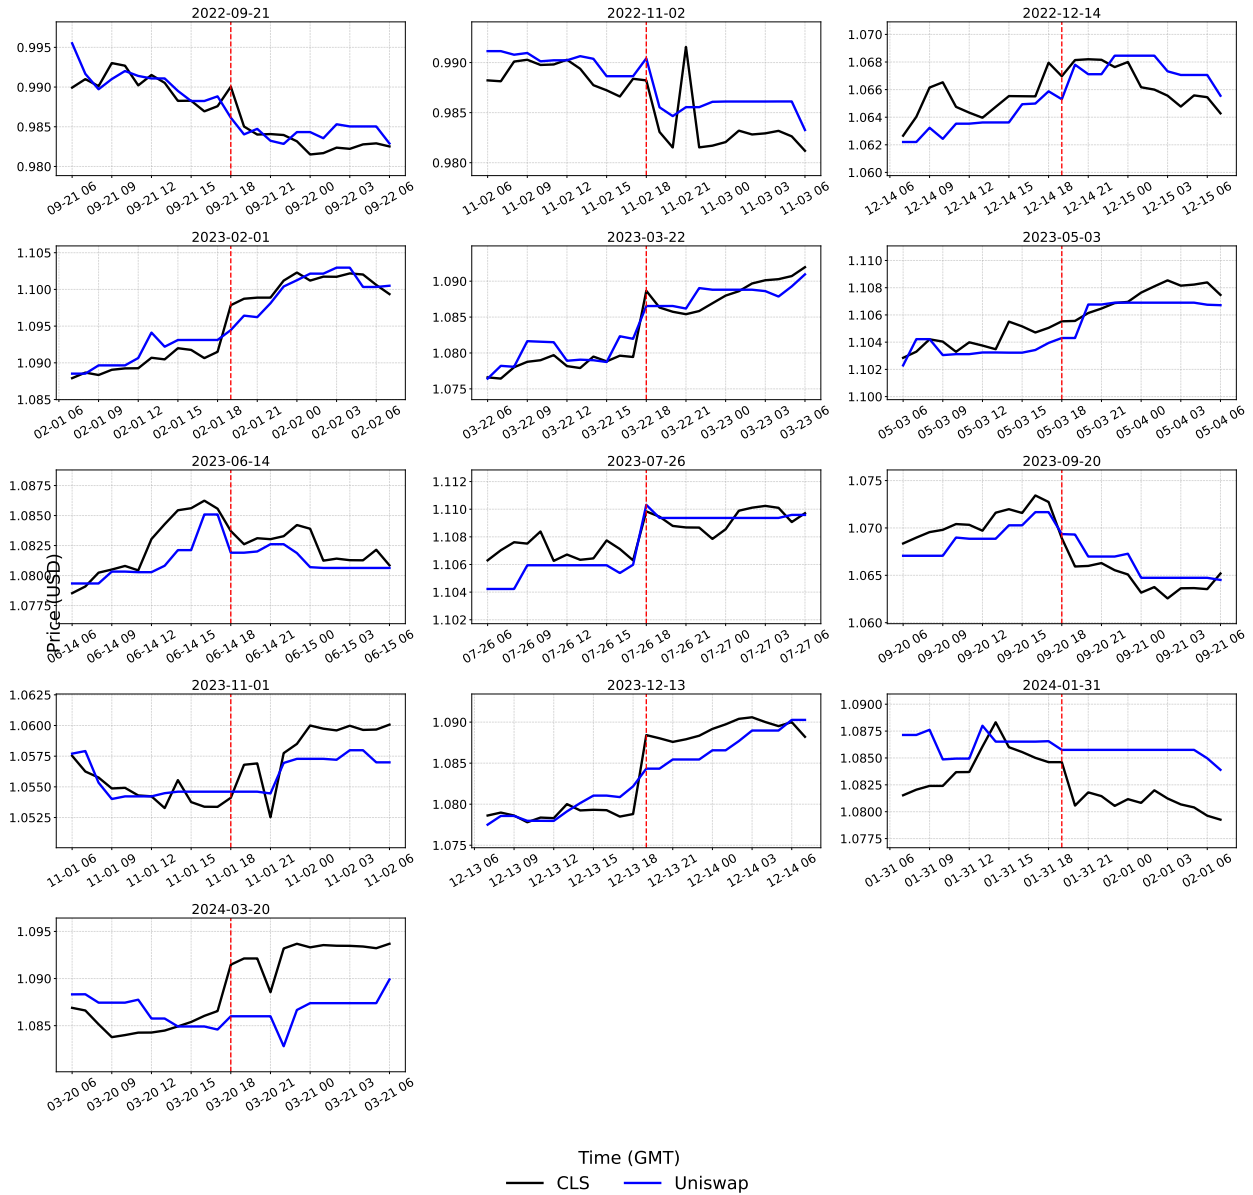


Table E.4: Welch T-test: Absolute Price Difference – FOMC vs Placebo

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.

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

Appendix F: SVAR Identification Assumptions

This appendix provides further detail on the identification strategy for the 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 SVAR to examine the contemporaneous and dynamic relationship between sector-level order flow and exchange rate changes. The structural form is given by

$$AY_t = A_0 + \sum_{j=1}^L A_j Y_{t-j} + \varepsilon_t, \quad (33)$$

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 applying A^{-1} to both sides:

$$Y_t = C_0 + CY_{t-1} + B\varepsilon_t, \quad (34)$$

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

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

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

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

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

1. OTC Block A^{OTC} .

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

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: Robustness Tests

This appendix reports robustness checks for the SVAR analysis described in Section IV.B.1 and Appendix F. The first three exercises build on the same baseline specification and identification scheme and modify the model along a single dimension at a time. We assess robustness to liquidity provision, higher-frequency dynamics, trader heterogeneity, intra-day variation in price impact, and just-in-time (JIT) liquidity behavior by LPs.

Variable definitions, recursive ordering, and Cholesky identification follow Appendix F. Each robustness exercise alters either the information set, sampling frequency, or composition of the DEX block, while preserving the hierarchical structure between OTC order flow, DEX order flow, and prices.

(i) Liquidity Controls. In the first robustness test, we extend the baseline SVAR to control for liquidity provision measured from Uniswap V3 mint and burn events, as defined in Appendix B.3. These variables capture changes in liquidity supply near and away from the prevailing market price.

To assess whether liquidity adjustments affect the estimated price impact of order flow, we augment the DEX block to include both liquidity measures ordered before all DEX order flows,

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

This ordering allows liquidity supply to affect order flow and returns contemporaneously, but not vice versa within the same interval. The resulting VAR stack is

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

Impulse responses in Figure G.1 show that controlling for liquidity provision leaves the estimated price impact of sophisticated traders and primary dealers largely unchanged.

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

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

with CLS returns interpolated to 5-minute intervals. The recursive ordering within the DEX block follows the baseline specification. Figure G.2 shows that shocks from Top 10 wallets and primary dealers remain economically and statistically significant.

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

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

with Group 1 containing the smallest traders and Group 5 the largest. The VAR stack becomes

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

where DEX quintiles are ordered from smaller to larger traders. Figure G.3 shows that larger traders exert stronger and more persistent price impact on both DEX and CLS returns.

(iv) Intra-day Patterns. We also examine intra-day variation in price impact by estimating hour-of-day-specific responses to a one million EURC shock in DEX order flow. Figure G.4 shows that price impact from sophisticated traders and primary dealers is concentrated between 13:00 and 15:00 UTC, overlapping with core European and U.S. trading hours, while LPs exhibit no systematic intra-day pattern.


(v) Just-in-time Liquidity. Finally, we investigate just-in-time (JIT) liquidity strategies, in which LPs mint liquidity immediately before a trade and burn it within the same block. Table G.1 lists all detected JIT sequences in the EURC/USDC pool, defined as mint–swap–burn triplets within a single block. These events are rare and concentrated in a single wallet, suggesting that JIT behavior does not materially affect the aggregate price dynamics documented in the SVAR analysis.

G.1 SVAR tests

G.1.1 Controlling for Liquidity Provision

Figure G.1: Hourly price impact of blockchain order flow with liquidity controls

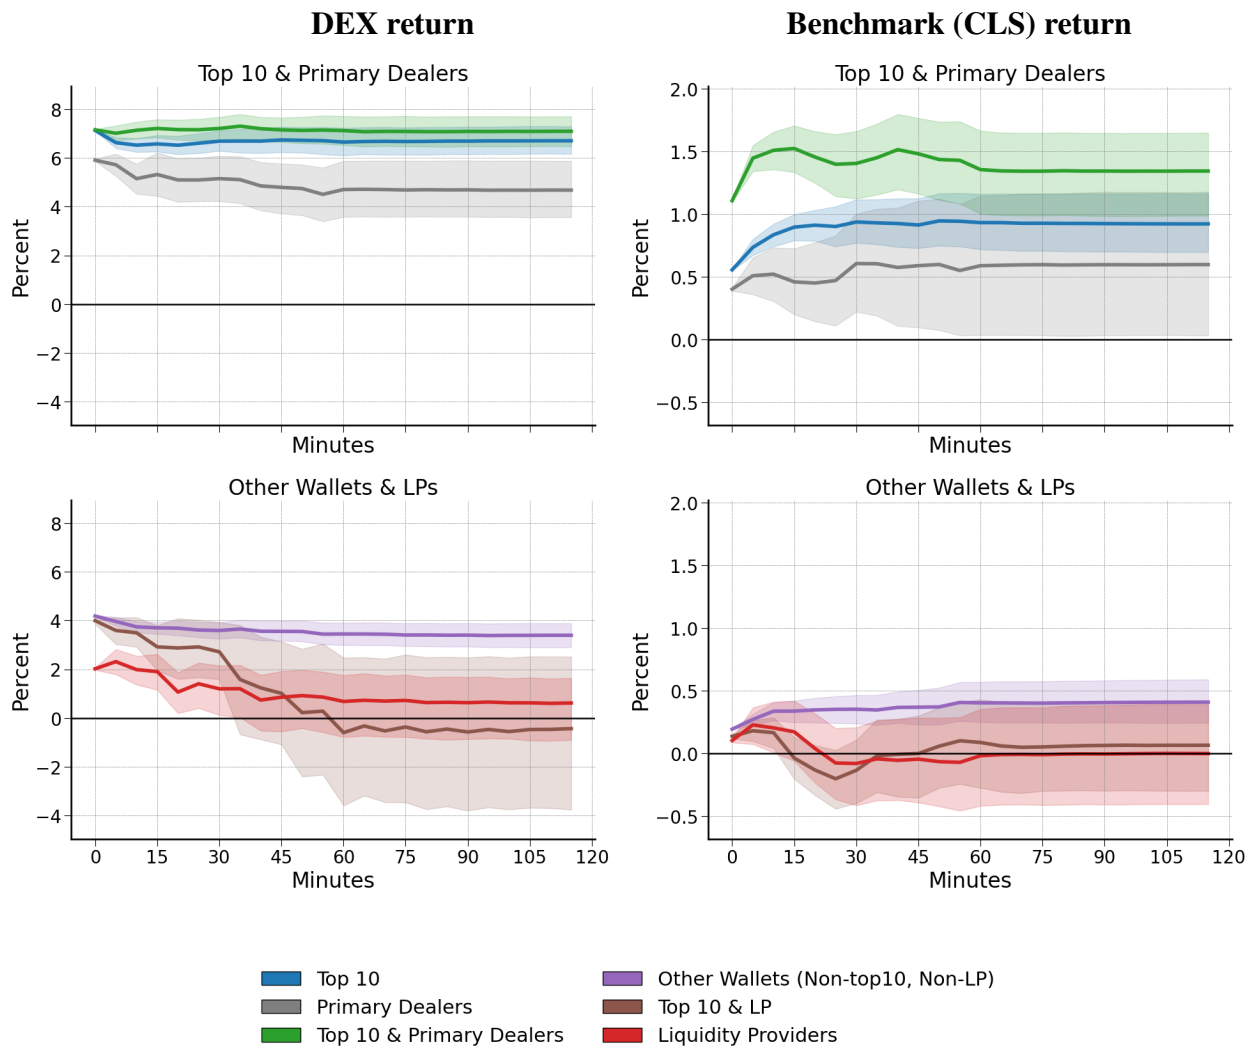
Note: This figure plots impulse responses of returns to a one million EURC shock in blockchain order flow at the hourly frequency, controlling for liquidity provision. The top row corresponds to transactions by Top 10 wallets and primary dealers, while the bottom row corresponds to other wallets and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



G.1.2 High-Frequency

Figure G.2: High-frequency price impact of blockchain order flow by trader groups

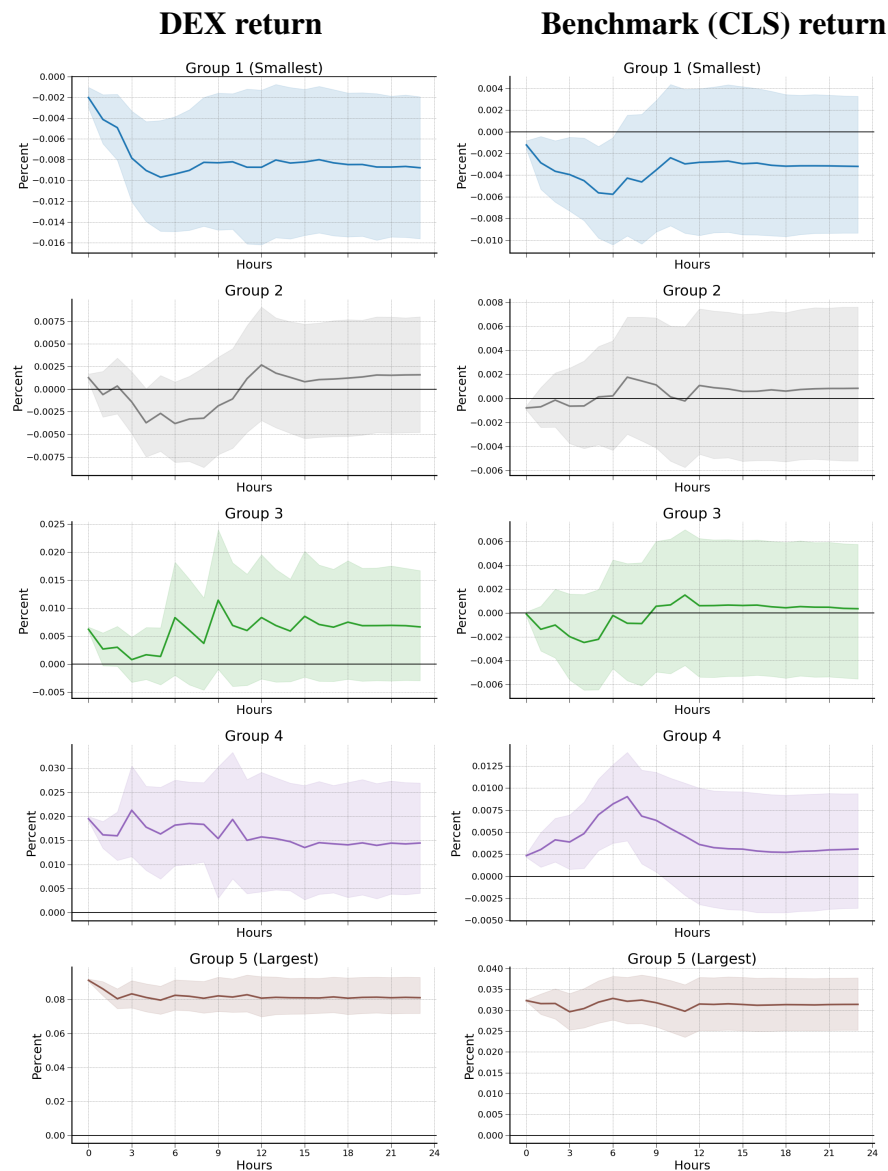
Note: This figure plots high-frequency (5-minute) impulse responses of returns to a one million EURC shock in blockchain order flow. The top row corresponds to transactions by Top 10 wallets and primary dealers, while the bottom row corresponds to other wallets and LPs. The left column reports EURC/USDC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



G.1.3 Trading Volume

Figure G.3: 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 traders; Group 5 = largest traders). Shocks are normalized by the standard deviation of order flow within each quintile. The left column reports EURC/USDC returns from Uniswap V3, and the right column reports EUR/USD returns from CLS. Responses are estimated using a structural VAR with 1,000 bootstrap replications. The sample period is 15 August 2022 to 30 April 2024.



G.2 Intra-day patterns

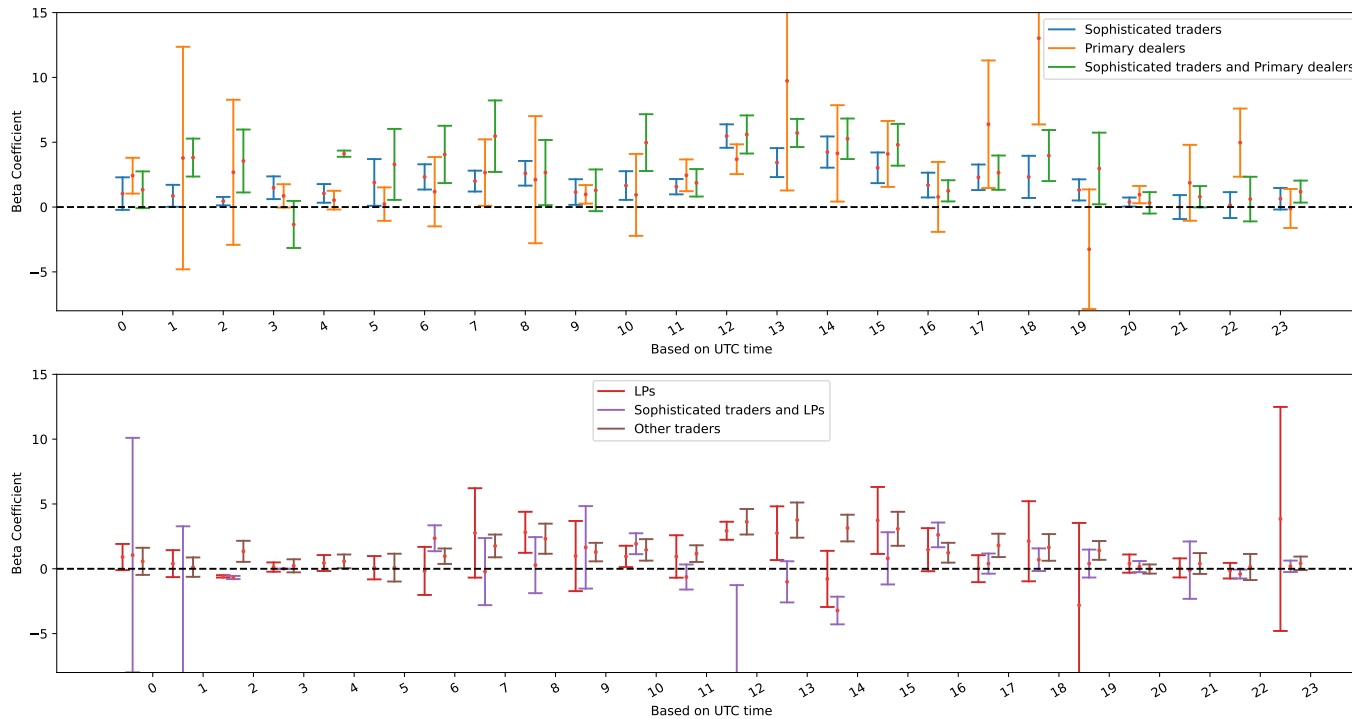
Figure G.4: Price impact of blockchain order flow: intra-day patterns

Panel (a): EURC/USDC Return



Panel (b): CLS Benchmark EUR/USD Return

Note: This figure plots hourly price-impact estimates of spot returns in response to a one million EURC shock to blockchain order flow. Order flow measures net EURC buyer-initiated transactions sourced from Uniswap V3 trade data. EURC/USDC returns are computed from Uniswap V3 prices, while EUR/USD returns are sourced from CLS. Panel (a) reports the response of EURC/USDC returns, and Panel (b) reports the response of EUR/USD returns. Blockchain order flow is disaggregated into six trader categories: sophisticated traders (Top 10 wallets), primary dealers, LPs, and their intersections. The sample period is 15 August 2022 to 30 April 2024.



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

Table G.1: Transaction Details

Note: This table reports just-in-time (JIT) liquidity transactions in the EURC–USDC pool. The LP wallet is 0xae2fc483527b8ef99eb5d9b44875f005ba1fae13 (abbreviated ae13). Each JIT event is a mint–swap–burn triplet in a single block. The sample period is 15 August 2022 to 30 April 2024.

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	
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	

Continued on next page

Table G.1: Transaction Details (continued)

Date (UTC)	Blk	Type	User	EURC	USDC	Lower P	Upper P	Price
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	

Appendix H: Private Information

H.1 Public versus Private Transactions

Figure H.1: 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 reports responses of EURC/USDC returns from Uniswap V3, and the right column reports responses of EUR/USD returns from CLS. Results are estimated using a SVAR with 1,000 bootstrap replications. Private transactions correspond to trades routed off-chain prior to execution, while public transactions occur directly on-chain. Trading groups include the top 10 wallets, the intersection of the top 10 wallets and LPs, and all remaining wallets. The sample period is 15 August 2022 to 30 April 2024.

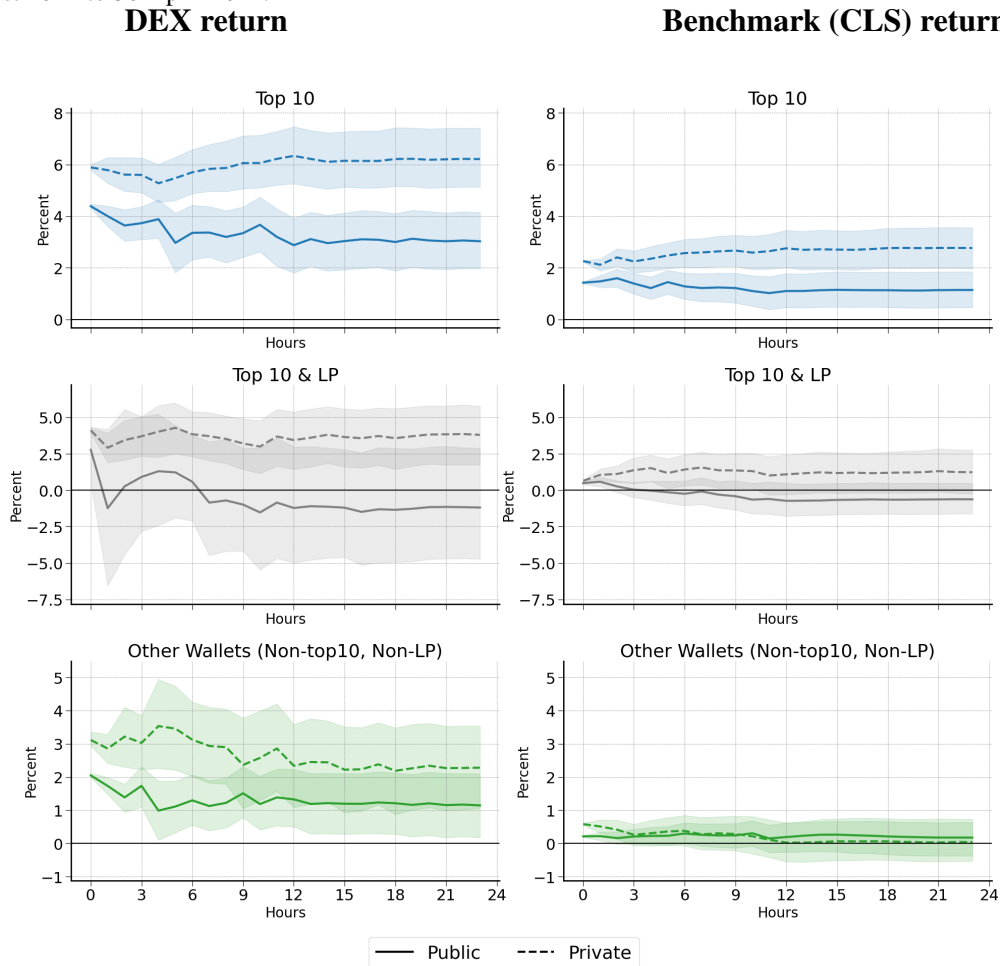


Table H.1: Determinants of EURC/USDC Order Flow: Public Transactions

Note: This table reports regressions of *public* order flow on the price difference between DEX and CLS exchange rates. *OF* measures net EURC buyer-initiated transactions using Uniswap V3 data. $P_{DEX} - P_{CLS}$ denotes the price difference between DEX and CLS exchange rates. Order flow is decomposed into sub-categories, including top-10 wallets, traders with access to primary markets, and liquidity providers (LPs). Standard errors are Newey–West (HAC) and reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. The sample period is 15 August 2022 to 30 April 2024.

	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*** (0.0334)	-0.0145** (0.0060)	-0.0227* (0.0120)	-0.1371*** (0.0297)	0.0034 (0.0079)	-0.0003 (0.0002)	-0.1936** (0.0979)
DEXReturn _{t-1}	-0.0009 (0.0009)	-0.0002 (0.0002)	0.0003 (0.0006)	-0.0012 (0.0010)	0.0000 (0.0003)	-0.0000 (0.0000)	-0.0009 (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*** (0.0001)	0.0000 (0.0000)	0.0002*** (0.0001)	-0.0000 (0.0001)	-0.0001** (0.0000)	-0.0000 (0.0000)	0.0003 (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

Table H.2: Determinants of EURC/USDC Order Flow: Private Transactions

Note: This table reports regressions of *private* order flow on the price difference between DEX and CLS exchange rates. *OF* measures net EURC buyer-initiated transactions using Uniswap V3 data. $P_{DEX} - P_{CLS}$ denotes the price difference between DEX and CLS exchange rates. Order flow is decomposed into sub-categories, including top-10 wallets, traders with access to primary markets, and liquidity providers (LPs). Standard errors are Newey–West (HAC) and reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. The sample period is 15 August 2022 to 30 April 2024.

	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 (0.0432)	0.0049 (0.0044)	0.0020* (0.0012)	-0.0004 (0.0004)	-0.0063* (0.0038)		-0.0342 (0.1159)
$DEXReturn_{t-1}$	-0.0055** (0.0023)	-0.0000 (0.0000)	-0.0000 (0.0000)	0.0000 (0.0000)	0.0002* (0.0001)		-0.0022 (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*** (0.0001)	-0.0000 (0.0000)	-0.0000 (0.0000)	0.0000 (0.0000)	0.0001*** (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

H.2 Decomposing Order Flow into a Arbitrage vs Residual Component

Figure H.2: Price impact of blockchain order flow: information versus arbitrage trading (EUR/USD CLS Return)

Note: This figure plots impulse responses of spot EUR/USD returns to a one million EURC shock in blockchain order flow, estimated using an SVAR with 1,000 bootstrap replications. Blockchain order flow measures net EURC buyer-initiated transactions from Uniswap V3 trade data, while EUR/USD prices are sourced from CLS. To isolate informational content from arbitrage activity between decentralized and traditional markets, order flow is decomposed by regressing it on the lagged DEX–CLS price difference, yielding an arbitrage-driven component and a residual component. The top panel reports responses to the residual (information) component, and the bottom panel reports responses to the predicted (arbitrage) component. Results are shown for blockchain order flow sub-categories: sophisticated traders (Top 10 wallets), primary dealers, and their intersection. The sample period is 15 August 2022 to 30 April 2024.

