

Does Market Efficiency Impact Capital Allocation Efficiency? The Case of Decentralized Exchanges

Evgeny Lyandres and Alexander Zaidelson*

March 15, 2024

Abstract

We examine the effect of market efficiency on the efficiency of capital allocation in the setting of decentralized exchanges of crypto assets. Utilizing data on nearly 100 million trades in concentrated liquidity pools on two leading blockchains, we construct a highly granular, capital-market-based measure of capital allocation efficiency. We also design and implement a method of identifying market-efficiency-restoring arbitrage transactions among all blockchain transactions and construct arbitrage-based granular measures of market efficiency. We find that market efficiency has positive, economically and statistically significant, and causal impact on capital allocation efficiency.

Keywords: Market efficiency; capital allocation; decentralized exchanges; concentrated liquidity

*Lyandres: Coller School of Management, Tel Aviv University, Israel. Email: lyandres@tauex.tau.ac.il. Zaidelson: SCRT Labs, Israel Email: zaidelson@gmail.com. We thank seminar participants at Aarhus University, Bar Ilan University, Montpellier Business School, Monash University, National University of Singapore, Shanghai University of Finance and Economics, Shenzhen University, Sun Yat-sen University, University of California at Santa Barbara, Wayne State University, 2023 Hong Kong University Finance Forum, 2023 Quarterly Journal of Finance Forum, 2024 Crypto Asset Lab Conference, 2024 International Conference on Computational Finance, 2024 Future of Financial Information Conference, 2024 Fintech Chair at Paris Dauphine University Conference, and 2024 FMA Asia/Pacific Conference for valuable comments and suggestions. Evgeny Lyandres is grateful to the Henry Crown Foundation of Business Research in Israel, the Vicky and Joseph Safra Research Institute for Banking and Financial Intermediation, and Chaire Fintech at University Paris Dauphine - PSL for providing financial support for this project.

“[England’s financial] organization is so useful because it is so easily adjusted. Political economists say that capital sets towards the most profitable trades, and that it rapidly leaves the less profitable non-paying trades. But in ordinary countries this is a slow process, ... In England, however, ... capital runs as surely and instantly where it is most wanted, and where there is most to be made of it, as water runs to find its level.”

Walter Bagehot, 1873

1. Introduction

There is vast evidence that developed financial markets facilitate growth (see Levine (1998), Levine et al. (2000), Durnev et al. (2004), and Acemoglu et al. (2005) for country-level evidence; Jayaratne and Strahan (1996) for the U.S.-state-level evidence; Rajan and Zingales (1998) and Beck and Levine (2002) for industry-level evidence; and Demircuc-Kunt and Maksimovic (1998) for firm-level evidence.)

A plausible channel for this relation is the ability of developed financial markets to facilitate efficient allocation of capital to its most profitable uses. Efficient capital allocation can be due several reasons: efficiency of prices (e.g., Tobin (1969)), screening efficiency (e.g., Diamond (1984), Boyd and Prescott (1986), and King and Levine (1993)), and reduced agency costs (e.g., Jensen (1986)).

Empirically, there does not seem to be a consensus regarding the empirical relation between financial market development and capital allocation efficiency. On the one hand, Wurgler (2000) finds that elasticity of industry investment to value added is higher in countries with developed capital markets and Rajan and Zingales (1998) report that external-finance-dependent industries grow faster in financially developed markets. On the other hand, bank-centered financial systems may be as good at efficiently allocating capital as developed public financial markets. For example, Beck and Levine (2002) show that external-finance-dependent industries do not grow faster in market-based systems than in bank-based financial systems.

In this paper we examine the impact of one important component of financial market development — market efficiency — on the efficiency of allocation of capital in this market. Our main result is that there seems to be a positive, statistically and economically significant, and likely causal link between market efficiency and capital allocation efficiency.

Our setting is decentralized exchanges (DEXes) of crypto assets. The market for crypto currencies is still in its infancy: The overall crypto market capitalization is around 1.4 trillion dollars as of November 2023, compared with roughly 100 trillion dollars capitalization of global equity markets. Yet, this market and in particular its decentralized part, which is characterized by a unique trading technology in which assets are exchanged without reliance on financial intermediaries, present a unique laboratory for investigating our research question. The reason is that in this market it is possible to measure both market efficiency and capital allocation efficiency at the level of granularity unattainable in traditional financial markets.

Utilizing data on close to 100 million trades in crypto assets facilitated by so-called ”concentrated liquidity pools” over a 2.5-year period on two blockchains — Ethereum and Polygon — we estimate a measure of capital allocation efficiency for every liquidity pool every week. There are nearly 50,000 pool-week observations on Ethereum blockchain and nearly 25,000 on Polygon blockchain. The main reason for examining two blockchains is that one of them (Polygon) serves as a control sample in an identification test of the causal relation between market efficiency and capital allocation efficiency on the other (Ethereum).

Our capital allocation efficiency measure is based on examining how useful concentrated liquidity is in reducing the price impact of trades. For every pool-week, we estimate the liquidity of a hypothetical non-concentrated-liquidity pool that would lead to the lowest average normalized difference between the price impact that would have occurred in this hypothetical pool and the realized price impact in the real, concentrated-liquidity pool. Our estimation procedure is able to match trade outcomes quite precisely. We define capital efficiency as the ratio of liquidity in the hypothetical pool with non-concentrated liquidity to the liquidity of the concentrated-liquidity pool. This pool-week measure exhibits large cross-sectional (across pools), time-series, and cross-chain

variation.

Building on Fama (1970), who argues that arbitrage activity is the mechanism that is at the core of market efficiency and that leads to adherence to the law of one price, our market efficiency measure is based on the proportion of efficiency-restoring arbitrage transactions in a pool during a week. The higher the fraction of arbitrage transactions the shorter the expected time to closure of gaps in prices across markets and to restoration of market efficiency. We develop a procedure for identifying efficiency-restoring arbitrage transactions with various degrees of complexity, among all transactions recorded in the two blockchains.

Having constructed granular measures of both capital allocation efficiency and market efficiency, we proceed to examining the relation between the two. Our baseline result is that there is a strong and robust positive association between market efficiency and capital allocation efficiency, after controlling for numerous other factors that may impact the desirability of liquidity concentration: non-market-efficiency restoring arbitrage transactions, overall bot activity in the pool, utilization of the pool's liquidity, relative average transaction size, volatility of exchange rate of the pool's assets, the overall state of the market, and the costs of performing transactions on the blockchain ("gas costs").

To examine whether the relation between market efficiency and capital allocation efficiency is causal, we estimate the response of capital allocation efficiency on Ethereum to an event that had a significant (albeit likely temporary) effect on the arbitrage activity and, by implication, on market efficiency, in liquidity pools on Ethereum. The event is "Shapella" upgrade of Ethereum blockchain that disrupted economic relations between "block builders", who are crucial for performing effective arbitrage on the blockchain by optimally placing proposed arbitrage transactions within blocks, and "validators", who receive proposed blocks of transactions from the builders and add them to the blockchain. We find that the effect of a temporary reduction in market efficiency on capital allocation efficiency in Ethereum pools relative to matched pools on Polygon (that did not experience the shock) is profound: Capital efficiency is 0.3-0.4 standard deviations lower in the weeks following the shock.

Our paper’s contribution is twofold. First, we contribute to the economic literature examining the effects of financial development. Our findings indicate that market efficiency — a crucial feature of developed financial markets — leads to more efficient allocation of capital. To our knowledge, this is the first evidence, using a (trade)-outcome-based granular measure of efficiency of capital allocation, of the impact of market efficiency on capital allocation efficiency. The observed positive effect is consistent with theoretical arguments in Tobin (1969), Diamond (1984), Boyd and Prescott (1986), King and Levine (1993), and Jensen (1986). While the setting for our analysis is a relatively niche and rapidly developing market — on-chain trading in crypto assets, there is nothing in the characteristics of this market that would suggest that our findings would lack external validity.

Our second contribution is to the emerging literature studying decentralized exchanges. Lehar and Parlour (2023) and Foley et al. (2023) perform comprehensive empirical analyses of trading and liquidity provision on DEXes. Capponi and Jia (2023) investigate theoretically and empirically market microstructure of DEXes. Aoyagi (2020), Milionis et al. (2022), Heimbach et al. (2022), and Lehar et al. (2023) perform theoretical and empirical analyses of liquidity provision in non-concentrated and concentrated liquidity pools. and Caparros et al. (2023) studies empirically liquidity repositioning strategies of liquidity providers in the presence off fixed transaction (gas) costs, which is also a focus of investigation of Lehar et al. (2023). Barbon and Rinaldo (2023) examine empirically transaction costs and price efficiency on DEXes. Park (2023) analyzes several tradeoffs involved in a DEX pool design. Hasbrouck et al. (2022) and Frtisch et al. (2023) model optimal choice of pool fees in a DEX. Angeris et al. (2023) perform a theoretical analysis of price stability on DEXes. Malamud and Rostek (2017) compare welfare on a decentralized exchange to that on its centralized counterpart.

We contribute to this literature by providing the first analysis of capital allocation efficiency in concentrated liquidity provision, which has become the dominant way of facilitating trades in crypto assets on the blockchain. We develop a new, capital-market-based measure of capital allocation efficiency, and document that while concentrated liquidity provision is useful in facilitating efficient trading, concentrating liquidity does not always leads to efficient outcomes. We also examine

various determinants of capital allocation efficiency in concentrated liquidity pools.

The remainder of the paper proceeds as follows. In the next section, we discuss the basics of decentralized exchanges and concentrated liquidity provision. Section 3 describes the data and summary statistics. Section 4 develops and analyzes our measure of capital allocation efficiency. Arbitrage-based measures of market efficiency are developed in Section 5. The empirical analysis of the relation between market efficiency and capital allocation efficiency is presented in Section 6. Section 7 concludes.

2. Setting

2.1. Decentralized exchanges: Overview

The setting of our analysis is decentralized exchanges of crypto assets (DEXes). Traditional (centralized) crypto exchanges (CEXes), such as Binance, Coinbase, and Kraken among over 200 spot crypto exchanges as of November 2023¹ rely on order book technology, akin to that employed by stock exchanges. DEXes, facilitating trades in crypto assets on public blockchains, largely rely on an entirely different type of trading technology, referred to as "Automated Market Making (AMM)". This alternative trading technology has become possible thanks to two characteristics of public blockchains. The first complete transparency of all transactions on a blockchain. The second is the ability of smart contracts to facilitate trades that are recorded directly on the blockchain ("on-chain") without requiring custody of users' funds at any point by intermediaries, such as CEXes.²

AMM-based trading, first introduced to the crypto trading market by Bancor and Uniswap protocols in 2017-2018,³ gained popularity during the summer of 2020, often referred to as the "Decentralized Finance (DeFi) summer". The popularity of AMM-based trading may be attributed to three main elements. The first is the problematic nature of CEXes that involves centralized

¹See leading aggregators, such as Coinmarketcap (www.coinmarketcap.com) and Coingecko (www.coingecko.com).

²See Cong and He (2019) for a general discussion and analysis of smart contracts and John et al. (2022) for an analysis of smart contracts in decentralized finance.

³See Hertzog et al. (2017) and Adams et al. (2020) for Bancor and Uniswap white papers.

custody of users' funds in a largely unregulated environment,⁴ as was clearly illustrated by the monstrous collapse in November 2022 of FTX – the third largest crypto exchange at the time. Second, unregulated crypto exchanges are plagued by fake trading, hurting individual traders.⁵ Third, as we discuss below, AMM-based trading reduces the extent of adverse selection that noise traders are exposed to, as a counterparty to any AMM-based trade is a blind pool of liquidity, as opposed to a (potentially better informed) trader.

Figure 1 presents the evolution of overall trading volume on DEXes.

[FIGURE 1 HERE]

The peak of trading on DEXes occurred at the turn of 2021 and 2022, in which monthly trading volume routinely exceeded \$US 200 billion. Trading volume on DEXes has decreased drastically since. This reduction largely coincided with the "crypto winter", which began with a series of large-scale scandals, such as the collapse of a popular blockchain Terra along with its fully algorithmic stable coin UST in May 2022 and defaults of large crypto investment funds and lending platforms (e.g., Three Arrows Capital and Celcius) soon thereafter, and was further deepened by the aforementioned collapse of FTX.

In relative terms, DEXes have been gaining importance even throughout the crypto winter, with the fraction of trading volume on DEXes relative to that on CEXes reaching 0.5 in November 2023,⁶ up from 0.05-0.10 throughout 2021-2022. Roughly half of all DEX trading volume routinely occurs on Ethereum, the oldest and most popular blockchain enabling smart contracts and, as a result, on-chain trading — one of the two blockchains we analyze in this study. Much of the other half of on-chain trading has been happening on Ethereum Virtual Machine (EVM)-compatible "layer-2" blockchains, which tend to be less secure but allow for a much larger throughput of transactions and are much cheaper as a result.⁷ Our analysis extends to one of the largest layer-2 chains — Polygon — responsible for roughly 5% of DEX trading volume.

⁴See Yadav (2022) for a discussion of regulation of crypto exchanges.

⁵See Cong et al. (2023) and Amiram et al. (2023) for analyses of fake (wash) trading on CEXes.

⁶See DeFillama (www.defillama.com), the leading aggregator of on-chain (DeFi) data

⁷See Schaffner (2021) for a discussion of various features of layer-2 blockchains.

2.2. *Constant product automated market making*

Automated market making allows trading in two assets using a "liquidity pool" containing these assets.⁸ Once liquidity has been deposited into a pool, a trader can withdraw any amount of one of the pool's assets, "asset out" (up to the total amount of that asset in the pool) by depositing a deterministic amount of the other asset into the pool. This deterministic "amount in" is a function of just two arguments — the amounts of the two assets in the pool at the time of the trade and the amount of "asset out" the trader wishes to withdraw from the pool. Similar to traditional, order-book-based trading, AMM-based trading features two main types of traders — noise traders and arbitrageurs, typically referred to as "Maximum Extractable Value (MEV)" bots; this classification of traders will play a crucial role in our analysis. Importantly, the outcome of the trade is only a function of trade size and the state of liquidity in the pool at the time of the trade, and is never a function of a trader's identity or type.

The economics of liquidity provision to AMM-based DEXes consists of the following elements. Liquidity providers to a pool receive fractional trading fees for every trade involving the assets in the pool in either direction.⁹ Swap fees remain in the pool, allowing the pool to grow over time. Any time a liquidity provider deposits assets into a pool ("establishes a new liquidity position"), she receives newly minted "liquidity pool (LP) tokens", in proportion to the newly established position relative to overall pool size. LP tokens constitute an on-chain, verifiable representation of every liquidity position. A liquidity provider into a pool is able to withdraw her liquidity from the pool partially or fully at any time by sending her LP tokens to the smart contract governing the pool and "burning" them.

Liquidity provision is not a risk-free proposition. In addition to the usual risk of holding the two assets deposited into the pool, liquidity providers (LPs) are subject to the "impermanent loss" risk. Impermanent loss occurs whenever an arbitrageur makes a profitable trade in the pool in response

⁸Typically, liquidity pools consist of two assets. However, protocols such as Balancer and Curve have developed AMM-based trading technologies in liquidity pools consisting of multiple assets. See Martinelli and Mushegian (2019) and Egorov (2019) for Balancer and Curve white papers, respectively.

⁹See Hasbrouck et al. (2022) for a model of optimal fees in AMM.

to exogenous changes in the two assets' relative values. In other words, impermanent loss is a loss compared to a situation in which there were no trades in the pool following exogenous changes in prices of the pool's assets. Impermanent loss may be reversed if/when the exchange rate of the two assets returns to the one at which the liquidity was provided. However, impermanent loss becomes permanent if the exchange rate at the time the liquidity is withdrawn by an LP from the pool differs from the exchange rate at the time of establishment of her liquidity position. Impermanent loss is a one-sided risk — i.e. expected impermanent loss is always positive and is increasing in the volatility of the exchange rate between the pool's assets.¹⁰

Trading curves in AMM pools take various configurations. The earliest and still the most common type of AMM is "constant product market making", in which trading in assets X and Y is defined by a constant: $X \times Y = K$. Consider an example. Assume that there are two liquidity providers who deposit liquidity into a pool of assets X and Y at 1:2 exchange rate.¹¹ Assume that the first (second) LP deposits 50 (100) and 100 (200) units of assets X and Y into the pool, respectively, resulting in the following pool composition: $X = 150$ and $Y = 300$. The two LPs' positions and the resulting trading curve in the pool are depicted in Figure 2a.

[FIGURE 2 HERE]

Under constant product trading curve, the amount of asset Y , ΔY , that a trader receives from the pool in return to sending (selling) ΔX units of asset X to the pool equals

$$\Delta Y = Y - \frac{X \times Y}{X + \Delta X(1 - f)} = \frac{\Delta X(1 - f) \times Y'}{X + \Delta X(1 - f)} \quad (1)$$

where f is the pool's fractional trading fee. Assume that in our example, a trader wishes to sell 15 units of asset X , i.e. $\Delta X = 15$. Assume also that the trading fee in the pool is 0.003, i.e. 0.3% of an asset sent to the pool stays in the pool, whereas the rest is exchanged for the other asset (asset out). The resulting amount of asset out is: $\Delta Y = \frac{15 \times 0.997 \times 300}{150 + 15 \times 0.997} = 27.198$. This trade is illustrated in

¹⁰See Milionis et al. (2022) and Deng et al. (2023) for theoretical analyses of impermanent loss in liquidity provision.

¹¹Liquidity is typically provided at the "correct" exchange rate corresponding to either an exogenously given exchange rate on another DEX or CEX or the exchange rate based on the market's assessment of the two assets' valuations. Providing liquidity at any other rate would lead to an arbitrage opportunity and result in impermanent loss to liquidity providers.

Figure 2b, where pool composition moves along the trading curve from the black dot to the green one as a result of the trade.

Constant product trading curve satisfies the following intuitive property: The price impact of a trade (i.e. the difference between the marginal exchange rate and the average exchange rate in a trade) is an increasing function of the ratio of the amount of asset sold to the pool (ΔX) and pool size ($X \times Y = K$). Price impact can be computed as:

$$1 - \frac{\Delta Y}{\frac{Y\Delta X}{X}} - f = \frac{\Delta X(1-f)^2}{X + \Delta X(1-f)}, \quad (2)$$

and in our example equals 9.04%, in addition to 0.3% pool fee.

The equilibrium pool size is a result of the following tradeoff. On the one hand, a larger pool leads to a lower price impact, increasing the pool’s attractiveness to traders. On the other hand, for a given trading volume, the return to liquidity provision is decreasing in pool size.¹² The curvature of the trading function involves the following tradeoff. On the one hand, the flatter the trading curve the lower the trading costs due to lower price impact. On the other hand, the flatter the trading curve the larger the impermanent loss to liquidity providers, as a flatter curve leads to larger arbitrage trades and higher arbitrage profits in response to exogenous changes in prices of pool’s assets, requiring higher pool fees to compensate LPs for the higher impermanent loss risk. Optimal trading curve balances these two considerations. While constant product curve dominates trading on DEXes, other curves, such as the combination of constant product and constant sum, introduced by Curve protocol, $X \times Y + \xi \times (X + Y) = K$, which is flatter around the spot exchange rate, have become popular.¹³

2.3. Concentrated liquidity and capital allocation efficiency

A problem with the constant product curve, sometimes referred to as "Uniswap V2" or simply "V2", as most constant-product pools are governed by either the second version of open-source

¹²See Aoyagi (2020) and Lehar and Parlour (2023) for theoretical analyses of the equilibrium size of AMM-based liquidity pools.

¹³See Capponi and Jia (2023) and Angeris et al. (2023) for models of optimal design of trading curves on an AMM.

smart contracts deployed by the largest and most successful DEX protocol — Uniswap — or their clones (“forks”), is that most of the liquidity provided to pools is idle, resulting in relatively low capital efficiency. Consider the example illustrated in Figure 2. As a rule of thumb, the price impact of a trade roughly equals the ratio of trade size and pool’s liquidity. Thus, trades making use of a large fraction of the pool’s liquidity rarely occur — generally, and in our sample that we describe below. In May 2021, Uniswap came out with revolutionary new smart contracts allowing for concentrated liquidity provision, usually referred to as “Uniswap V3” or “V3”.¹⁴ The idea is that optimal liquidity should not be distributed evenly along the full trading curve, as in Figure 2. Instead, optimal liquidity allocation would be characterized by higher liquidity deployed to the portion of the trading curve in which the majority of the trades occur, i.e. around the spot exchange rate.

More formally, when providing concentrating liquidity, an LP sets the lower and upper bounds of the exchange rate between which her liquidity would be concentrated. As long as the exchange rate in the pool remains between the two bounds, the LP’s effective (“virtual”) liquidity is larger than her actual liquidity deployed to the pool. This larger de-facto liquidity results in larger share of the pool and higher fraction of the pool’s trading fees. By the time the exchange rates leaves these predetermined bounds, the entire LP’s liquidity has been swapped to the cheaper of the two assets and becomes idle until the exchange rate returns to the bounds or until the liquidity position is rebalanced.

Consider the following modification of the example above. Assume that the second LP continues to provide standard, non-concentrated liquidity of $X = 100$ and $Y = 200$, illustrated by the blue dot in Figure 3a. Non-concentrated liquidity provision into a concentrated liquidity pool is equivalent to setting the lower and upper exchange rate bounds at 0 and ∞ , respectively.

[FIGURE 3 HERE]

The first LP decides to concentrate her liquidity in a way that would result in five-fold leveraged liquidity within a certain range and idle liquidity outside that range. The first LP’s real liquidity

¹⁴See Adams et al. (2021) for detailed explanation of Uniswap V3 AMM as well as Barbon and Rinaldo (2023) for a discussion of differences between Uniswap V2 and V3 AMM.

is still $X = 50$, $Y = 100$, resulting in $K = 50 \times 100 = 5,000$. To achieve the desired level of five-fold liquidity concentration, define virtual liquidity, $K^* = K \times 5^2 = 125,000$. Denote by $p_b = \left(\frac{Y}{X}\right)_b > p_a = \left(\frac{Y}{X}\right)_a$ the upper and lower bounds (exchange rates) of virtual liquidity provision. Virtual concentrated liquidity is given by

$$\left(X + \left(\frac{K^*}{p_b}\right)^{0.5}\right) \times \left(Y + (K^* \times p_a)^{0.5}\right) = K^*. \quad (3)$$

Assume that the chosen upper bound of the exchange rate is $p_b = 2.5$. Plugging $X = 50$, $Y = 100$, and $K^* = 125,000$ into (??) results in lower bound, p_a , of 1.228, i.e. five-fold leveraged liquidity is possible in the range of exchange rates (1.228, 2.5). The first LP's real liquidity is denoted by the red dot in Figure 3a and the resulting virtual trading curve is depicted by the upper red curve, which is a linear transposition of the real (lower red) trading curve that fits the chosen exchange rate bounds. The resulting total trading curve in the pool is depicted in green. Note that the pool's trading curve is discontinuous, due to the fact that the first LP's liquidity is active within the exchange rate bounds and is idle outside of the bounds.

Assume now that a trader swaps 15 units of asset X for Y in the concentrated liquidity pool. Assume temporarily that the pool's post-trade exchange rate would still be inside the first LP's concentrated liquidity bounds — a conjecture we will verify. This trade, illustrated by the movement of the pool's liquidity from the green dot to the black one along the green trading curve in Figure 3b, results in the following amount of asset Y received by the trader: $\Delta Y = \frac{15 \times 0.997 \times 700}{350 + 15 \times 0.997} = 28.684$. Note that trading occurs over the virtual trading curve: $(100 + 50 \times 5) \times (200 + 100 \times 5) = 245,000$ instead of the real trading curve in the non-concentrated liquidity example: $(100 + 50) \times (200 + 100) = 45,000$. The resulting price impact of the trade is 4.09%, compared with the price impact of 9.04% in the non-concentrated liquidity pool. The post-trade exchange rate in the pool, $\frac{300 - 28.684}{150 + 15} = 1.644$ is within the concentrated liquidity bounds, consistent with the conjecture that the trade is possible along the virtual trading curve.¹⁵ This example is helpful in defining capital allocation efficiency

¹⁵In reality, the exchange rate plane is divided into "ticks" using a fine grid, and virtual trading curve within each tick is a combination of virtual liquidity of all non-idle liquidity positions in that tick. A trade can occur within a tick or across multiple ticks.

in our setting:

Def.: *Capital allocation efficiency is the ratio of a) the value of liquidity that needs to be deposited into a non-concentrated liquidity pool to achieve the same price impact of a trade as the realized price impact in the concentrated liquidity pool, and b) the value of liquidity of that concentrated liquidity pool.*

We will use this definition of capital allocation efficiency in our empirical analysis. A non-concentrated liquidity pool has capital allocation efficiency of one. Tighter exchange rate bounds result in higher capital allocation efficiency if the spot exchange rate is within the bounds. In our example, the capital allocation efficiency of the concentrated liquidity pool equals $\frac{350 \times 2 + 700}{150 \times 2 + 300} = 2.33$.

Concentrating liquidity does not always lead to capital efficiency improvement. Consider an example illustrated in Figure 3c, in which the current spot rate is outside the concentrated liquidity bounds: $p_b = 1.25$ and $p_a = 0.466$. In this example, the first LP's liquidity is idle, and the total pool's trading curve consists of just the second LP's liquidity. The resulting amount out in this trade is $\Delta Y = \frac{15 \times 0.997 \times 200}{100 + 15 \times 0.997} = 26.019$ and the price impact is 13.27%, compared to price impact of 9.04% in the case of non-concentrated liquidity. Inefficient concentration of liquidity results in capital efficiency that is lower than one, $\frac{100 \times 2 + 200}{150 \times 2 + 300} = 0.67$ in this example.

3. Data and summary statistics

Our primary data consist of all trades in all concentrated liquidity ("V3") pools on two blockchains. The first is Ethereum, the largest and most important blockchain supporting smart contracts and, by implication, decentralized finance. The second is Polygon — one of the leading "layer 2" scaling solutions for Ethereum, developed with the goal of reducing costs of transacting on the blockchain (at a cost of reduced security). The sample period is between May 6, 2021 — the date of establishment of the first concentrated liquidity pool on Uniswap — and September 17, 2023, for a total of 125 weeks of trading data. Concentrated liquidity pools have captured a large fraction of overall trading volume on DEXes over the last two years, as is evident in Figure 4.

The reason for performing the analysis on two separate blockchains is twofold. First, due to orders-of-magnitude differences in costs of transacting on Ethereum (where transactions tend to be quite expensive) and Polygon (where transactions are cheap), the distributions of trades on the two blockchains are dissimilar. There are lower economies of scale on Polygon, leading to significantly smaller transactions on average. The vastly different distributions of transactions have important implications for estimation of capital allocation efficiency, measurement of market efficiency, and the relation between the two. Thus, examining transactions on Polygon in addition to Ethereum enhances external validity of our analysis. Second, (a subset of) liquidity pools on Polygon serves as a control sample in an identification test that we perform that is based on a significant event affecting market efficiency that occurred on Ethereum blockchain but not on Polygon.

These trading data are vast, encompassing over 30 million transactions in concentrated liquidity pools on Ethereum and over 45 million transactions on Polygon. The trading data was generated from the raw blockchain data using a series of data transformations. The raw blockchain data for Ethereum, Polygon and several other chains is made available to the public in the form of SQL tables available through publicly available datasets hosted on Google BigQuery analytical platform. The data extraction and transformation processes required to create and update the tables are developed and maintained by Blockchain ETL project (<https://github.com/blockchain-etl>). The raw blockchain data tables contain the table of all transactions and also the table of all log events. Log events are usually emitted by smart contracts to make the information about their inner state available to potential consumers. For example, most AMM pool smart contracts issue log events such as *swap* (when a swap operation occurs), *mint* (when liquidity is deposited into the pool), *burn* when liquidity is withdrawn from the pool, and more. The source tables created by Blockchain ETL contain the data in binary format, and meticulous research was required to create mappings between binary representations of data and specific events. Using the mappings, a multi-stage SQL transformation process was created. In the process, first the swap and liquidity even data was extracted per protocol, then an overall table of swaps was created. In addition, higher-level tables were also created, uniting several swaps into a logical transaction. For example,

a multi-leg swap consists of several swapping events, and it is useful to be able to represent such multi-leg swap in a single trade.

We also employ data on crypto asset daily prices obtained from www.coinmarketcom.com, data on gas prices on Ethereum and Polygon from respective blockchain explorers, www.etherscan.io and www.polygonscan.com. The unit of our analysis is pool-week. The reason is that our procedure for estimating capital allocation efficiency of a given pool within a given time frame requires a non-trivial number of trades in the pool during that period. For the same reason, we eliminate pool-week observations with fewer than 30 trades in the pool.

Summary statistics are presented in Table 1. Panel A presents the statistics for Ethereum, whereas Panel B presents those for Polygon.

[TABLE 1 HERE]

There is a total of nearly 50,000 pool-week observations on Ethereum, corresponding to roughly 400 pool observations per week, as evident from Figure 5a.

[FIGURE 5 HERE]

The mean \$US value of liquidity of a V3 pool on Ethereum — typically referred to as "Total value locked (TVL)" is nearly \$4 million, whereas the median is an order of magnitude smaller. The TVL of the largest pool on Ethereum exceeds \$1 billion. The combined TVL of Ethereum pools peaked at around \$2.7 billion in June 2022 and has dropped by about 50% since, as evident from Figure 5b.

The mean (median) weekly transaction volume in a V3 pool on Ethereum is \$4 million (\$250 thousand). Figure 5c demonstrates that total weekly transaction volume in V3 pools reached its peak of \$3.5B in January 2022 and is down to roughly \$1 billion over the last year of the sample. After imposing the restriction of at least 30 weekly transactions in a pool, there are over 300 transactions in a pool per week on average, whereas the median exceeds 100. The number of weekly transactions in all V3 pools combined fluctuates between 100,000 to 200,000 throughout the sample, as can be seen in Figure 5d.

Volume-to-TVL — a measure of pool utilization — has a wide range. Liquidity in some pools is not utilized at all, whereas the weekly volume in the busiest pool exceeds 130 times its TVL. The mean (median) weekly pool utilization is slightly below two (one). The mean pool-week level transaction size is also quite widely distributed. The dollar value of transactions in V3 pools is sizable: the mean transaction size exceeds \$7,000, and in a median transaction \$2,000 of crypto assets change hands. Relative to pool TVL, however, most transactions are quite small, consistent with price impact increasing nearly linearly with relative transaction size — the mean (median) transaction corresponding to 0.43% (0.09%) of pool TVL, however the largest transaction saw the value of asset being sold exceeding the pool’s TVL (and resulting in price impact of over 50%).

The volatility of exchange rate of the two assets in the pool, measured as the daily standard deviation of log daily returns of the two assets’ exchange rate over a four-week period preceding the week of the observation ranges between nearly 0 (for pools of two ”stable coins”, i.e. representations of \$U.S. on the blockchain)¹⁶ to over 40%, with the mean (median) of 6% (5%) or roughly 100% annualized standard deviation. We also measure the state of the market, proxied by the return on Ether — the largest and the most important asset in the decentralized finance ecosystem, during the four weeks preceding the observation week. Ether’s 4-week standard deviation is nearly 20%, or around 70% annually.

Finally, we measure the cost of ”gas” on the blockchain that is paid to blockchain validators for updating and maintaining the blockchain’s history. The overall cost of any transaction on the blockchain is the product of the pre-determined number of units of gas (which is generally a function of the complexity of the transaction) and time-varying, demand-for-transaction-driven, price of one unit of gas, typically expressed in Gwei, which is one-billionth of an Ether. Gas price and, as a result, the cost of transacting on the blockchain, tends to be higher when there is higher congestion on the blockchain and when Ether price is higher. For convenience, the price of one unit of gas in Table 1 is expressed in 10^{-9} \$US. The variation in gas prices over time is large: the ratio of the highest and lowest gas prices within the 2.5-year period exceeds 40.

¹⁶See Cong and Xiao (2021) for a taxonomy of crypto assets.

Moving to Polygon (in Panel B of Table 1), there are over 24,000 pool-week observations, corresponding to roughly 250 pools per week; the activity in V3 pools on Polygon started about 8 months later than that on Ethereum, as can be seen in Figure 5a. Polygon pools tend to be significantly smaller than those on Ethereum – by a factor of 5 (10) for the mean (median) pool TVL; the same is true for transaction volume. Interestingly, the mean (median) number of transactions in Polygon pools is over 3 (1.5) times larger than the corresponding figures in Ethereum pools (see also Figure 5d), despite the order-of-magnitude-lower trading volume, consistent with lower economies of scale on Polygon. Pool utilization is slightly higher on Polygon, whereas average and median absolute transaction sizes are 15-30 times lower on Polygon than on Ethereum; on the other hand, transaction sizes computed as a fraction of pool TVL have similar distributions on the two blockchains. Polygon has a larger fraction of low-volatility pools, as evident from lower mean and median volatility of exchange rate of pool assets. Finally, as discussed above, gas prices are vastly different on the two blockchains: the mean (median) gas price — and by implication the overall cost of transaction — is around 500 times lower on Polygon than on Ethereum. The latter finding is consistent with Caparros et al. (2023) who report orders-of-magnitude gas price differences between Ethereum and Polygon.

4. Capital allocation efficiency measurement

We construct our measure of capital allocation efficiency utilizing the definition in Section 2.3, which is based on a comparison between the TVL of a concentrated liquidity (V3) pool and the TVL of a hypothetical non-concentrated liquidity (V2) pool that would result in the same quantity of asset bought by the trader (amount out) for a given quantity of asset sold (amount in) as the V3 pool. This theoretical definition can be satisfied precisely for each individual trade. However, as pool composition and TVL during the week is dynamic — both because each trade changes the amounts of the two assets in the pool and because asset prices fluctuate throughout the week, our procedure for measuring pool-week-level capital allocation efficiency consists of the following steps.

We first find the amounts of the two assets in the hypothetical V2 pool that minimize the

sum of relative deviations between the observed amount out in each trade and the hypothetical amount out that would have been received for that trade in the hypothetical V2 pool, scaled by the trade's price impact. Mathematically, this amounts to minimizing the following function for each V3 pool-week:

$$\min_{X^{V2}, Y^{V2}} \sum_i \left| \frac{\left(\frac{X^{V2} \times Y^{V2}}{IN^{V2} + \Delta IN_i (1-f)} - \Delta OUT_i \right) \times p_{OUT,i}}{\Delta IN_i \times p_{IN,i} - \Delta OUT_i \times p_{OUT,i}} \right|, \quad (4)$$

where X^{V2} and Y^{V2} are the amounts of assets X and Y in a hypothetical nonconcentrated-liquidity (V2) pool; IN^{V2} and OUT^{V2} stand for X^{V2} and Y^{V2} (Y^{V2} and X^{V2}) when X (Y) is being sold by a trader into the pool for Y (X); ΔOUT_i and ΔIN_i are (real) amounts in and out, respectively in trade i that occurred during a given week in a given pool (in trades in which asset in is X (Y), ΔOUT_i and ΔIN_i take the values of ΔY_i and ΔX_i (ΔX_i and ΔY_i), respectively); and $p_{OUT,i}$ and $p_{IN,i}$ are weekly average prices of assets that serve as asset out and asset in in trade i .

The expression $\frac{X^{V2} \times Y^{V2}}{IN^{V2} + \Delta IN_i (1-f)}$ in the numerator of (4) is the hypothetical amount out in trade i in the V2 pool; the difference between this hypothetical amount out and real amount out, multiplied by the price of asset out, is the (raw, non-normalized) error in the estimate of amount out in trade i . We divide this error by the dollar price impact of a trade, which is the denominator of (4) to give weight to both small and large trades in the minimization problem.

We then compute the TVL of the hypothetical V2 pool by multiplying the matched liquidity of the two assets found by minimizing (4) by their respective average weekly prices. Finally, we divide the TVL of the hypothetical V2 pool by the weekly average TVL of the V3 pool to obtain the pool-week-level capital efficiency estimate:

$$CapEff = \frac{X^{V2} \times p_X + Y^{V2} \times p_Y}{\bar{X} \times p_X + \bar{Y} \times p_Y}, \quad (5)$$

where \bar{X} and \bar{Y} are the average amounts of assets X and Y in the concentrated-liquidity (V3) pool throughout the week.

Table 2 presents summary statistics related to capital efficiency estimation for Ethereum and

Polygon blockchains.

[TABLE 2 HERE]

The first row in both panels describes the estimation error at the pool-week level. The mean (median) estimation error is 2% (1.4%) for Ethereum pools and 1.2% (0.7%) for Polygon pools, suggesting that our procedure for estimating the liquidity in hypothetical V2 pools that would match real trading outcomes in V3 pools is quite precise.¹⁷

The second row in each panel presents distribution of estimated capital allocation efficiency — the focus of our empirical investigation. There is wide variation in pool-week-level capital efficiency measure. On Ethereum, the mean capital efficiency is 215, whereas the median is 6. This skewness is driven by several (mostly stable coin) pools with very high capital allocation efficiency values, exceeding 5,000 on several occasions. While there are pools with inefficient allocation of capital (i.e. capital efficiency measure lower than one), capital allocation efficiency of many pools belonging to even the lowest quartile is positive. The distribution of capital allocation efficiency measure on Polygon is quite similar. Given the large skewness of the capital allocation efficiency measure, in the empirical analysis we use its natural logarithm, the distribution of which appears in the third row in both panels. Figure 6 describes the evolution of mean log capital efficiency over time.

[FIGURE 6 HERE]

Capital allocation efficiency on both blockchains is quite persistent. Interestingly, it tends to be higher on Polygon, which may be consistent with the prevalence of lower-volatility pools there — a conjecture we examine empirically below.

In Table 3, we examine the determinants of the precision of our estimation of TVLs of V2-equivalent pools. We estimate, separately for Ethereum and Polygon, linear regressions in which the dependent variable is the pool-week level absolute estimation error (computed as in (4) using the liquidity of the pool’s assets obtained as a result of solving (4)). Even columns include pool and week fixed effects.

¹⁷We have eliminated approximately 0.5% of pool-weeks from the analysis in which the average estimation error exceeds 10%.

[TABLE 3 HERE]

Estimation precision is increasing in the number of weekly transactions in a pool: A one-standard-deviation increase in log weekly transactions (8.2) reduces the estimation error by 0.20% (0.14%) for the case of Ethereum (Polygon), corresponding to 0.1 standard deviations of estimation error. In addition, estimation error is negatively related to the volume of transactions in the pool. The intuition is that in the presence of price variability throughout a week, price impact of a trade is estimated more precisely when trades are larger. Finally, estimation precision is significantly negatively related to the standard deviation of the exchange rate of the pool’s assets: Increasing exchange rate volatility by one standard deviation raises the estimation error by 0.3 (0.38) standard deviations in Ethereum (Polygon) pools. The intuition is that the variation throughout the week in the real V3 pool’s composition is increasing in the volatility of the exchange rate of the pool’s assets, making the assumption of a constant TVL of the hypothetical V2 pool less grounded in reality and distorting the estimation precision.

In light of the variability in the precision of capital allocation efficiency estimation, we examine robustness of empirical tests discussed below, by excluding pool-weeks with estimation errors exceeding 5% and/or 2%. All the qualitative results are robust to estimation within subsamples of pool-weeks in which capital allocation efficiency is estimated with the highest degree of precision.

5. Arbitrage-based measure of market efficiency and other measures of bot trading

5.1. Market-efficiency and efficiency-restoring arbitrage

The central question of our research is the relation between market efficiency and capital allocation efficiency. After having estimated the latter, we turn to estimating market efficiency at the blockchain-pool-week level. According to Fama (1970), the temporary nature of deviations of prices from efficient benchmarks are a result of arbitrageurs’ efficiency-restoring trades. Thus, it is natural to construct a proxy for market efficiency that is based on the prevalence of arbitrage

transactions.

Blockchain-based "on-chain" trading presents a perfect setting for estimating an arbitrage-based measure of market efficiency for the following reasons. First, all transactions on the blockchain are observable in real time — attempted transactions appear in the so-called "memory pool" as soon as they are submitted, whereas completed transactions appear in the block in which they were executed. Thus, it is straightforward to uncover arbitrage opportunities as soon as they appear, and arbitrage-searching "Maximum Value Extraction" bots have become an industry in its own right.

Second, possibilities for arbitrage on DEXes are endless, since assets are connected through a variety of pools on various protocols, and all protocols deployed on the same blockchain are interoperable, i.e. an arbitrage transaction may consist of several legs across several protocols. An example of such transaction is presented in Figure 7.

[FIGURE 7 HERE]

In this example, an arbitrage is performed using three different pools on two separate protocols. The arbitrageur first swaps 413.63 rETH to 408.41 wstETH on Curve, then swaps the received 408.41 wstETH to 462 ETH on Uniswap V3, and then swaps 446.13 ETH to 413.63 rETH on another Uniswap V3 pool. The arbitrageur returns the originally swapped rETH and is left with a handsome profit of 15.9 ETH (around \$30,000 at that moment).¹⁸

Third, there is a single limit to arbitrage — the gas costs that needs to be paid to blockchain validators for including a transaction on the blockchain. These gas costs are typically quite small for transactions of moderate complexity — ranging from a fraction of a cent on Polygon to single-digit dollars on Ethereum. For example, the gas costs associated with the transaction in the example above amounted to 0.0052 ETH (approximately \$10.8).

Last but not least, unlike arbitrage in traditional markets, arbitrage on the blockchain does not require initial capital and is nearly risk-free. The reasons are a) the principle of "atomicity" in blockchain transactions and b) the availability of "flash loans". Atomicity implies that if a

¹⁸This transaction is available at <https://t.ly/XCPtr>.

transaction contains several legs, one of which cannot be executed, the transaction fails in its entirety. A flash loan is a loan that is taken and returned (with interest) within the same transaction. Thus, many arbitrage transactions consist of taking a flash loan in an asset, performing the legs of arbitrage, and returning the loan, all within one transaction, removing the need in initial capital to execute arbitrage. If the flash loan cannot be returned, e.g., because the proposed arbitrage transaction was unprofitable to begin with or because another trader has already taken advantage of the arbitrage opportunity, the transaction (including the flash loan) fails (i.e. is cancelled), making both flash loans and attempted arbitrage transactions effectively risk-free, as only successful arbitrage transactions, in which flash loans are returned, can go through. One exception to this discussion is the gas costs that need to be paid for any attempted transaction regardless of its eventual success (or of its profitability relative to the costs of gas if the transaction is successful), which is the only material risk involved in performing arbitrage on the blockchain.¹⁹

To identify arbitrage opportunities, we search in every block on Ethereum and Polygon blockchains for profitable arbitrage transactions. Ethereum blocks appear roughly every 12 seconds, and each block typically contains low-single-digit hundreds of transactions. Polygon blocks have a two-second frequency, typically having several dozen transactions in a block.

A multi-leg trade is classified as a profitable, market efficiency-restoring arbitrage in a given pool if the following conditions are satisfied:

- a) the net amount out (i.e. the difference between amount out and amount in across all legs of the transaction) of every asset involved in the transaction is non-negative;
- b) the net amount out of at least one asset involved in the transaction is positive;
- c) all legs of the trade occur within a single transaction;
- d) at least one of the legs of the transaction involves a trade in the pool under examination.

Note that any transaction satisfying the first two criteria is necessarily bringing exchange rates in pools involved closer to each other (by buying an asset in a pool in which it is relatively cheap thereby raising its price and selling the asset in a pool in which it is relatively expensive, reducing

¹⁹See Hansson (2022) for a more detailed discussion of arbitrage on the blockchain.

its price as a result). The third condition is a definition of a bundled transaction needed to ensure arbitrage is risk-free.²⁰

Our measure of market efficiency at the pool-week level is the proportion of efficiency-restoring arbitrage transactions out of all trades in the pool during the week. We measure this proportion based on both the volume of trades and transaction count. The first two rows in each of the two Panels in Table 4 present the distribution of the two market efficiency measures in pools deployed on Ethereum and Polygon.

[TABLE 4 HERE]

The median (mean) proportions of arbitrage transactions based on trading volume and transaction count in Ethereum pools is 7-8% (15-16%). Some pools do not see any arbitrage activity, whereas in some other pools all transactions are classified as arbitrage. Arbitrage is more prevalent on Polygon, with the median (mean) fraction of arbitrage transactions being 14-25% (24-33%). The reason is lower limits to arbitrage due to significantly cheaper gas costs on Polygon. Interestingly, arbitrage transactions on Polygon tend to be smaller than non-arbitrage ones, as evident from comparing the first two rows in Panel B. This, again, is consistent with even small arbitrage transactions on Polygon being profitable due to low gas costs. The evolution of proportion of (trading-volume-based) arbitrage is depicted in Figure 8.

[FIGURE 8 HERE]

5.2. *Measures of bot activity unrelated to market efficiency*

It is possible that the proportion of market-efficiency-restoring arbitrage transactions as defined above is correlated with the general trading bot activity in the pool. There are two additional common types of bot trades that we observe: non-efficiency-restoring arbitrage and non-arbitrage bot transactions. To ensure that the relation between market efficiency (proxied by the proportion of efficiency-restoring arbitrage transactions) and capital allocation efficiency is not spurious, it is important to control for arbitrage transactions that do not enhance market efficiency. A prevalent

²⁰Similar criteria for identifying arbitrage transactions are used in Hansson (2022).

example of such transactions is "sandwich arbitrage", which is essentially a combination of front-running and back-running another trader's transaction. Sandwich arbitrage is possible on the blockchain because most transactions submitted to the blockchain can be observed in the public memory pool even before they are executed.²¹

There is a crucial difference between efficiency-restoring arbitrage transactions, which involve multiple pools, and sandwich arbitrage transactions, which occur in a single pool. The latter transactions represent a pure value transfer from the initiator of a transaction to the bot sandwiching it. Sandwich arbitrage transactions do not impact post-trade prices and cannot improve market efficiency.

A transaction is classified as a sandwich arbitrage in a given pool if it satisfies the following conditions:

- a) it has two legs involving the same pair of assets with opposite signs;
- b) the two legs are located in non-consecutive positions in the block;
- c) there is an external transaction involving the same two assets, such that assets in and out are the same as in the first leg of the sandwich arbitrage;
- d) the net amount out of each of the two assets is non-negative;
- e) the net amount out of at least one of the two assets is positive;
- f) both legs of the arbitrage are trades in the pool examined.

Figure 9 depicts an example of a sandwich arbitrage transaction. In this transaction, the attacker

²¹Originally on Ethereum, block validators were examining the transactions sent by the users, and then reordering them or inserting their own transactions to extract maximum value from each block (thus MEV - Maximum Extractable Value). This required validators to run complex software, thus increasing the cost of running a profitable validator and leading to higher centralization of the network. To solve the potential centralization issue, a design feature called Proposer/Builder Separation (PBS) was introduced to Ethereum. The main idea is that the task of building the most profitable block is separated from proposing the new block to the network. The initial version of PBS was implemented by Flashbots and is known as Mev-boost. PBS operates in the following way: Specialized players called Block Builders observe the memory pool and assemble the most profitable blocks. The Builders work in conjunction with another type of specialized players called Searchers, whose goal is to construct most profitable *bundles*, e.g. a efficiency-restoring arbitrage or a sandwich arbitrage, and make sure to be the first to execute it. As a result, market-efficiency-restoring transactions are often placed in beginnings of blocks, whereas the two legs of sandwich arbitrage can typically be found right before and right after a relatively large transaction. Builders perform an auction between the Searchers, and the highest bidders get their bundles included into the block constructed by the Builder. Builders then submit their blocks to Validators (also known as Proposers), who in turn choose the most profitable block out of all the blocks proposed by competing Builders. Usually, around 80% of all the arbitrage revenue is accrued by the Validators, with the remainder going to Searchers and Builders. See also Capponi et al. (2021) for a discussion of the role of block builders in extracting arbitrage opportunities.

identified that a large \$5.2M DAI-ETH swap transaction was submitted to the mempool (itself a multi-leg trade using USDC as transit asset), and found an opportunity to sandwich one of the legs of this trade, specifically the \$4.96M USDC-ETH swap. The arbitrageur constructed a bundle of three transactions:

- 1) A new transaction in which 5.266 million USDC were sold to the pool for 2,677 ETH — frontrunning the original transaction;
- 2) The original organic transaction;
- 3) A new transaction in which 2.677 ETH obtained in the frontrunning leg of the transaction were sold to the pool for 5.289 million USDC — backrunning of the original transaction, making a profit of \$22,900.

[FIGURE 9 HERE]

We define non-market-efficiency-related arbitrage activity at the pool-week level as the proportion of sandwich arbitrage (measured using either trading volume or trade count) in a pool during a given week. The third and fourth rows of Table 4 present the distribution of the proportion of sandwich arbitrage transactions on the two blockchains. Sandwich arbitrage is much more prevalent on Ethereum than on Polygon – the median fraction of sandwich transactions in a pool-week is 2-7% on Ethereum, compared with 0.2% on Polygon. This is likely due to the fact that the relations between bot operators and block builders are more developed on Ethereum due to larger transaction sizes, leading to more profitable arbitrage opportunities.

Bot activity in the pools is not limited to arbitrage. Trading bots often execute non-arbitrage transactions, likely related to high-frequency automated trading strategies. It is important to control for non-arbitrage bot trading in an analysis of capital allocation efficiency. The reason is that bot transactions are more likely to represent informed trading than transactions by noise traders. In AMM pools, the counterparty to every trade is liquidity providers in the pool. Prevalence of non-arbitrage bot trading in the pool against concentrated liquidity positions increases adverse selection that liquidity providers face, reducing the expected profitability of concentrated liquidity pools and, as a result, leading to lower capital allocation efficiency.

We define non-arbitrage activity as transactions by wallets identified in Ethereum and Polygon block explorers as bots that are not classified as efficiency-restoring or sandwich transactions. As evident from the last two rows of the two panels in Table 4, the median proportion of non-arbitrage bot trading on both blockchains is 5-6%, whereas the mean is 13-16%.

6. Determinants of capital allocation efficiency

In this section we examine the relation between our arbitrage-based blockchain-pool-week-level measure of market efficiency and the estimated efficiency of concentrated liquidity allocation in the pool. We begin by estimating panel regressions for the full samples of pools on Ethereum and Polygon and proceed by utilizing a shock to market efficiency on the Ethereum blockchain and performing a difference-in-differences analysis using a matched sample of pools on Ethereum and Polygon.

6.1. *Baseline analysis*

Table 5 reports results of regressing our capital allocation efficiency estimate on the measure of market efficiency and on additional variables that can affect the willingness and ability of liquidity providers to establish and maintain concentrated liquidity positions.

[TABLE 5 HERE]

In Panels A and B, the regressions are estimated using samples of Ethereum and Polygon pools, respectively. We use both trading-volume-based and transaction-count-based measures of proportion of efficiency-restoring arbitrage transactions.

The most important finding in Table 5 is that there is a positive relation between market efficiency and capital allocation efficiency. The coefficients on the proportion of efficiency-restoring arbitrage volume/transactions out of totals in the pool are positive and significant in seven specifications out of eight. In specifications without liquidity pool fixed effects (in odd columns), the relation is economically significant: A one-standard-deviation increase in the measure of market

efficiency is associated with 0.05-0.07 standard deviation increase in capital allocation efficiency in Ethereum pools and with 0.11-0.17 increase in capital allocation efficiency in Polygon pools.

Notably, these results are obtained while controlling for non-market-efficiency-restoring (sandwich) arbitrage, which proxies for the general propensity of bot activity in a pool. Bot activity is positively associated with capital allocation efficiency, the relation being significant in all specifications on Ethereum and in regressions without pool fixed effects on Polygon. Controlling for overall bot activity, capital allocation efficiency is significantly negatively related to the proportion of non-arbitrage bot transactions in the pool, which proxies for the fraction of informed trading. This is consistent with higher adverse selection faced by providers of concentrated liquidity.

Capital allocation efficiency is higher in pools with high liquidity utilization, as proxied by the volume-to-TVL ratio. It pays off to actively manage concentrated liquidity (or pay for automated on-chain concentrated liquidity management using services such as Gamma (www.gamma.xyz) and Algebra (<https://algebra.finance>)) when liquidity providers are competing for a share in a larger pie (i.e. larger pool fees).

Capital efficiency is significantly negatively related to relative transaction size, i.e. the ratio of average dollar value of transactions in the pool to pool TVL. The reason is that larger transactions are more likely to be executed outside of the capital efficiency bounds, leading to inefficient concentration of liquidity and subpar capital allocation. Capital allocation efficiency is also negatively related to the volatility of exchange rate of the pool's assets. The intuition is that higher exchange rate volatility leads to higher likelihood of the exchange rate leaving pre-determined concentrated liquidity provision bounds, leading to idle liquidity and low capital allocation efficiency.

Capital allocation efficiency on Ethereum is higher when the market has performed well, as proxied by the return on Ether over the previous four weeks, suggesting that liquidity providers consider concentrating liquidity a more worthwhile proposition in bull markets. This relation is insignificant on Polygon. Finally, on Ethereum, capital allocation efficiency is negatively related to gas prices. The reason is that concentrated liquidity positions need to be rebalanced often in response to fluctuations in exchange rates of pools' assets. Such rebalancing is costlier when

gas is expensive, reducing the profitability and attractiveness of concentrated liquidity positions. Interestingly, gas prices are not negatively related to capital allocation efficiency in Polygon pools, likely due to the fact that gas is cheap on Polygon and gas cost minimization considerations do not play a significant role in liquidity providers' strategies in pools deployed on Polygon.

To summarize, providers of concentrated liquidity seem to respond to pool and market characteristics that affects relative costs and benefits of leveraged liquidity provision. Controlling for several factors influencing pool-level capital allocation efficiency, it is found to be strongly positively related to arbitrage-based measure of pool-level market efficiency.

6.2. *Difference-in-differences estimation*

The results in the previous subsection demonstrate a positive association between measures of market efficiency and estimated capital allocation efficiency. In what follows, we attempt to uncover a causal association between the two by examining a response of capital allocation efficiency to a shock in market efficiency. The setting we utilize is the so-called "Shapella (Shanghai-Capella) upgrade" of Ethereum blockchain.

Prior to September 15, 2022, Ethereum blockchain used the "proof-of-work (PoW)" consensus mechanism, which was introduced and popularized by the Bitcoin blockchain,²² and in which block validators spend computer resources — a process commonly referred to as "mining" — in a battle for the right to add a new block to the blockchain.²³ On September 15, 2022, Ethereum blockchain transitioned to "proof-of-stake (PoS)" consensus, in which the likelihood of obtaining the right to add the next block is a function of financial resources spent by block validators in the form of "staked" (i.e. deposited) native currency of a blockchain — Ether in the case of Ethereum blockchain.²⁴ This transition of Ethereum from proof-of-work to proof-of-stake consensus is commonly referred to as "the Merge" (between the original Ethereum blockchain and a separate proof-of-stake blockchain called the Beacon Chain).

²²See Nakamoto (2008) for an introduction to proof-of-work consensus.

²³See Capponi et al. (2023) for a game theoretical analysis of mining.

²⁴See John et al. (2022) for an overview of proof-of-work and proof-of-stake consensus mechanisms and Saleh (2021) for an analysis of their differences.

As discussed above, performing arbitrage on a blockchain relies on an arbitrage bot’s ability to insert transactions in particular places in a block — typically either immediately following transactions in the same block or in the very beginning of the next block. This ability to precisely place transactions is possible due to established relations between arbitrageurs and ”block builders”, who assemble transactions into blocks and relay these proposed blocks to validators, who, in turn, append these blocks to the blockchain in case they win the right to do so. Importantly, block builders have complete discretion as to which transactions to include in a block and in which order, leading to tight economic incentives driving the relations between MEV bots and block builders on one hand and block builders and validators on the other hand.

While the Merge changed the identity of important blockchain validators — from ”mining pools”²⁵ to Ether ”staking pools”, such as Lido (<https://lido.fi>) and Coinbase Earn (<https://www.coinbase.com/earn>), this transition was gradual and expected, allowing block builders to establish relations with post-Merge validators prior to the Merge. Thus, the Merge is unlikely to have disrupted arbitrage opportunities on the blockchain.

Shapella upgrade of Ethereum blockchain, which occurred on April 4, 2023, allowed ”unstaking” (withdrawing) of staked Ether from the blockchain. This was widely expected to have significant and hardly predictable impact on the market shares of Ether staking pools and the identity of important validators — an expectation that was borne out in reality. This disruption of the market for blockchain validation had significant effect on the ability of block builders to have their proposed blocks accepted by validators, resulting in reduced possibilities to effectively perform arbitrage transaction. Thus, Shapella upgrade presented a negative shock to arbitrage-based market efficiency. Notably, while the timing and the overall (directional) effect of such shock were largely anticipated, there was little that could be done by block builders (and, by implications, arbitrageurs) to preempt it.

Shapella upgrade corresponds to week 102 in our sample. Notably, Polygon blockchain has utilized proof-of-stake consensus from its inception. Thus, liquidity pools deployed on Polygon

²⁵See Cong et al. (2021) for an analysis of mining pools.

can potentially serve as a control sample for the difference-in-differences analysis of the change in capital efficiency in Ethereum pools around Shapella upgrade. Consistent with this conjecture, the graphs of the proportion of efficiency-restoring arbitrage on Ethereum and Polygon diverge around week 102, as evident from Figure 8. The correlation between this proportion on the two blockchains is quite high (around 0.6) up to week 102, at which point it becomes negative (approximately -0.2 starting week 103, and -0.4 during weeks 103-114).

The subsamples relevant for the analysis are defined as follows. For each (treated) pool-week on Ethereum we attempt to find matching pool-week(s) on Polygon. A matching pool is defined as a pool that has liquidity in the same two assets as the focal pool.²⁶ Each Ethereum pool can have several matching pools on Polygon and multiple Ethereum pools can have the same matching pool on Polygon, since on both blockchains there are often several pools with liquidity in the same asset; these pools can be deployed on the same protocol or on multiple interoperable protocols.

We examine two timeframes. In the first, the pre-event subsample includes all treated and control pools during weeks 34-101 (i.e. from the first appearance of concentrated liquidity pools on Polygon to the week preceding the week of Shapella upgrade) and weeks 103-125 (from the week following Shapella upgrade until the end of the sample period). There are 5,316 Ethereum pool-weeks and 6,208 Polygon pool-weeks in this subsample. In the second subsample, we limit the analysis to 12 weeks before and after the week of Shapella upgrade (i.e. weeks 90-101 and 103-114). The reason for examining the shorter window is that the effect of disruption in arbitrage activity is likely temporary, as block builders eventually establish relations with new block validators. In addition, as is evident from Figure 6, the parallel trends assumption seems to be better satisfied throughout the shorter pre-event window. This subsample contains 1,752 (2,211) pool-weeks on Ethereum (Polygon).

We proceed to estimating regressions in which the dependent variable is our measure of capital allocation efficiency, whereas the main independent variables are: a) Ethereum indicator, b) Post-

²⁶In reality, most assets on Polygon are "bridged" versions of assets on Ethereum. For all practical purposes, these bridged assets are identical to the assets on Ethereum (part of which are themselves bridged from other blockchains, such as Bitcoin).

Shapella-upgrade indicator, and c) the interaction between the two. If the relation between market efficiency and capital allocation efficiency is causal, we expect a negative shock to arbitrage-fueled market efficiency to adversely affect capital allocation efficiency.

The results of the difference-in-differences estimation are presented in Table 6.

[TABLE 6 HERE]

The coefficients on the interaction between treated (Ethereum) and post (Shapella upgrade) indicators are positive and significant in all specifications and for both sample periods. In the longer sample period, the relative reduction in log capital allocation efficiency on Ethereum (-0.68-0.75) corresponds to 0.35-0.4 standard deviations of capital allocation efficiency measure. In the shorter, 24-week sample, the relative reduction in capital allocation efficiency on Ethereum corresponds to 0.27-0.30 standard deviations. The control variables tend to have the same sizes and comparable magnitudes to those in Table 5 with the exception of the coefficient on exchange rate volatility that turns positive.²⁷

Overall, the results in this section indicate that the relation between our pool-level arbitrage-based measure of market efficiency and the pool’s capital allocation efficiency is positive, economically sizable, and likely causal. This evidence is consistent with the capital allocation efficiency channel behind the positive relation between market efficiency and economic outcomes.

7. Conclusions

We examine the relation between market efficiency and capital allocation efficiency in the context of concentrated liquidity pools on decentralized exchanges facilitating trading in crypto assets. We develop novel, highly granular measures of both capital allocation efficiency and market efficiency on two leading public blockchains — Ethereum and Polygon. Our pool-week-level capital allocation efficiency measure is based on a comparison of the amount of concentrated liquidity with the amount of non-concentrated liquidity that would have been required to achieve similar trading

²⁷The reason is sizable correlation between exchange rate volatility and relative transaction size. Removing the latter from the regressions tends to flip the sign of the exchange rate volatility coefficients.

outcomes. We find that concentrated liquidity leads to significant improvements in trading efficiency (i.e. leads to efficient capital allocation) in most cases. However, there is a sizable fraction of observations in which concentrated liquidity results in subpar capital allocation. This highlights the need for developing and studying optimal methods for provision of liquidity — a question that is a subject of ongoing research.

Our market efficiency measure is based on identifying, among all blockchain transactions, arbitrage trades that restore market efficiency, and computing the fraction of such arbitrage transactions out of all weekly trades in every liquidity pool. We also identify non-efficiency-restoring arbitrage transactions and overall bot activity in each liquidity pool and control for these transactions while examining the effect of market efficiency on the efficiency of capital allocation.

Our main finding, obtained using both large samples of pools on Ethereum and Polygon and a smaller sample of matched pools that we use in conjunction with a shock to market efficiency on Ethereum, is that market efficiency seems to positively influence capital allocation efficiency. We choose the setting for analysing our research question — decentralized finance — as its characteristics, such as blockchain transparency and ability to facilitate smart-contract-governed trades that do not require financial intermediation, allow us to construct the unique measures that we use in the empirical analysis. Notably, there are no particular features of this setting that make us believe that our findings would not have external validity in other financial markets. In other words, an implication of our study that developed financial markets (at least in the sense of market efficiency) can contribute to economic growth through efficient capital allocation.

References

- Acemoglu, D., S. Johnson, and J. Robinson (2005). Institutions as a fundamental cause of long-run growth. *Handbook of Economic Growth 1*, 385–472.
- Adams, H., N. Zinsmeister, and D. Robinson (2020). Uniswap v2 core. <https://uniswap.org/whitepaper.pdf>.
- Adams, H., N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson (2021). Uniswap v3 whitepaper. <https://uniswap.org/whitepaper-v3.pdf>.
- Amiram, D., E. Lyandres, and D. Rabetti (2023). Competition and product quality: Fake trading on crypto exchanges. Tel Aviv University working paper.
- Angeris, G., T. Chitra, and A. Evans (2023). When does the tail wag the dog? curvature and market making. Bain Capital working paper.
- Aoyagi, J. (2020). Liquidity provision by automated market makers. University of Hong Kong working paper.
- Barbon, A. and A. Ranaldo (2023). On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. University of Hong Kong working paper.
- Beck, T. and R. Levine (2002). Industry growth and capital allocation:: does having a market-or bank-based system matter? *Journal of Financial Economics 64*(2), 147–180.
- Boyd, J. and E. Prescott (1986). Financial intermediary-coalitions. *Journal of Economic Theory 38*(2), 211–232.
- Caparros, B., A. Chaudhary, and O. Klein (2023). Blockchain scaling and liquidity concentration on decentralized exchanges. Warwick Business School working paper.
- Capponi, A. and R. Jia (2023). The adoption of blockchain-based decentralized exchanges. Columbia University working paper.
- Capponi, A., R. Jia, and Y. Wang (2021). The evolution of blockchain: from public to private mempools. Columbia University working paper.
- Capponi, A., S. Olafsson, and H. Alsbah (2023). Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Management Science forthcoming*.
- Cong, L. and Z. He (2019). Blockchain disruption and smart contracts. *Review of Financial Studies 32*(5), 1754–1797.
- Cong, L., Z. He, and J. Li (2021). Decentralized mining in centralized pools. *Review of Financial Studies 34*(3), 1191–1235.
- Cong, L., X. Li, K. Tang, and Y. Yang (2023). Crypto wash trading. *Management Science forthcoming*.
- Cong, L. and Y. Xiao (2021). Categories and functions of crypto-tokens. *Book Chapter in Palgrave Handbook of FinTech and Blockchain, Palgrave MacMillan*, 267–284.

- Demirguc-Kunt, A. and V. Maksimovic (1998). Law, finance, and firm growth. *Journal of Finance* 53(6), 2107–2137.
- Deng, J., H. Zong, and Y. Wang (2023). Static replication of impermanent loss for concentrated liquidity provision in decentralised markets. *Operations Research Letters* 51(3), 206–211.
- Diamond, D. (1984). Financial intermediation and delegated monitoring. *Review of Economic Studies* 51(3), 393–414.
- Durnev, A., K. Li, R. Morch, and B. Yeung (2004). Capital markets and capital allocation: Implications for economies in transition. *Economics of Transition* 12(4), 593–634.
- Egorov, M. (2019). Stableswap - efficient mechanism for stablecoin liquidity. <https://classic.curve.fi/files/stableswap-paper.pdf>.
- Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25(2), 383–417.
- Foley, S., P. O’Neill, and T. Putnins (2023). Can markets be fully automated? evidence from an “automated market maker. UC Berkeley working paper.
- Frtisch, R., S. Kaser, and R. Wattenhofer (2023). The economics of automated market makers. ETH Zurich working paper.
- Hansson, M. (2022). Arbitrage in crypto markets: An analysis of primary ethereum blockchain data. Stockholm University working paper.
- Hasbrouck, J., T. Rivera, and F. Saleh (2022). The need for fees at a dex: How increases in fees can increase dex trading volume. New York University working paper.
- Heimbach, L., E. Schertenleib, and R. Wattenhofer (2022). Risks and returns of uniswap v3 liquidity providers. ETH Zurich University working paper.
- Hertzog, E., G. Benartzi, and G. Benartzi (2017). Continuous liquidity and asynchronous price discovery for tokens through their smart contracts; aka “smart tokens”. <https://whitepaper.io/document/52/bancor-whitepaper>.
- Jayaratne, J. and P. Strahan (1996). The finance-growth nexus: Evidence from bank branch deregulation. *Quarterly Journal of Economics* 111(3), 639–670.
- Jensen, M. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review* 76(2), 323–329.
- John, K., L. Kogan, and F. Saleh (2022). Smart contracts and decentralized finance. Massachusetts Institute of Technology working paper.
- John, K., M. O’Hara, and F. Saleh (2022). Bitcoin and beyond. *Annual Review of Financial Economics* 14, 95–115.

- King, R. and R. Levine (1993). Finance and growth: Schumpeter might be right. *Quarterly Journal of Economics* 108(3), 717–737.
- Lehar, A. and C. Parlour (2023). Decentralized exchange: The uniswap automated market maker. UC Berkeley working paper.
- Lehar, A., C. Parlour, and M. Zoican (2023). Liquidity fragmentation on decentralized exchanges. UC Berkeley working paper.
- Levine, R. (1998). The legal environment, banks, and long-run economic growth. *Journal of Money, Credit and Banking* 30(3), 596–613.
- Levine, R., N. Loyaza, and T. Beck (2000). Financial intermediation and growth: Causality and causes. *Journal of Monetary Economics* 46(1), 31–77.
- Malamud, S. and M. Rostek (2017). Decentralized exchange. *American Economic Review* 107(11), 3320–3362.
- Martinelli, F. and N. Mushegian (2019). A non-custodial portfolio manager, liquidity provider, and price sensor. <https://balancer.fi/whitepaper.pdf>.
- Milionis, J., C. Moallemi, T. Roughgarden, and A. Zhang (2022). Automated market making and loss-versus-rebalancing. Columbia University working paper.
- Nakamoto, S. (2008). A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>.
- Park, A. (2023). The conceptual flaws of decentralized automated market making. *Management Science* forthcoming.
- Rajan, R. and L. Zingales (1998). Financial development and growth. *American Economic Review* 88(3), 559–586.
- Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *Review of Financial Studies* 34(3), 1156–1190.
- Schaffner, T. (2021). Scaling public blockchains: A comprehensive analysis of optimistic and zero-knowledge rollups. University of Basel working paper.
- Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of Money, Credit and Banking* 1(1), 15–29.
- Wurgler, J. (2000). Financial markets and the allocation of capital. *Journal of Financial Economics* 58(1-2), 187–214.
- Yadav, Y. (2022). (crypto) exchanges as regulators? Vanderbilt University working paper.

Figure 1: TIME-SERIES EVOLUTION OF COMBINED TRADING VOLUME ON DEXES

This figure presents monthly trading volume on Decentralized Exchanges (DEXes) on all blockchains combined from January 2020 to November 2023. The source of data is www.defillama.com.





Figure 2: LIQUIDITY PROVISION IN CONSTANT PRODUCT POOLS

Figure 2a presents an example of liquidity provision into a non-concentrated liquidity pool. Red and blue dots represent liquidity positions of two liquidity providers. Green dot is combined liquidity in the pool. Green curve is the constant product trading curve around the pool's liquidity.

Figure 2b presents an example of trading in the pool. Green dot represents the post-trade liquidity in the pool. Black dot adjacent to is the pre-trade liquidity in the pool. Red and blue dots represent the two liquidity providers' post-trade positions, whereas black dots adjacent to them represent their pre-trade positions.

(a) Liquidity provision and trading curve



(b) Trading in the pool





Figure 3: LIQUIDITY PROVISION IN CONCENTRATED LIQUIDITY POOLS


Figure 3a presents an example of liquidity provision into a concentrated liquidity pool. Red dot represents the first liquidity provider's concentrated liquidity position. The upper red curve represents trading curve resulting from the first LP's position, and the rays coming from the lower left corner represent the concentrated liquidity bounds. Blue dot represents the non-concentrated position of the second liquidity provider. Black dot represents combined real liquidity in the pool. Green dot represents combined virtual liquidity in the pool. Green curve represents the trading curve in the pool.

Figures 3b and 3c present examples of trading in the pool. Green dot represents the post-trade liquidity in the pool. Black dot adjacent to is is the pre-trade liquidity in the pool. Red and blue dots represent the two liquidity provider's post-trade positions, whereas black dot adjacent to them represent their pre-trade positions. (In Figure 3c blue dot coincides with green dot and red dot coincide with black dot.)

(a) Concentrated liquidity provision and trading curve



(b) Trading in the concentrated liquidity pool



(c) Inefficient concentrated liquidity





Figure 4: FRACTION OF TRADING VOLUME IN V3 POOLS OVER TIME

Figure 4a presents the fraction of trading volume on main DEXes deployed on Ethereum blockchain from March 2020 until November 2023. The DEXes analyzed are: [Please provide the list.](#)

Figure 4b presents the fraction of trading volume on main DEXes deployed on Polygon blockchain from March 2020 until November 2023. The DEXes analyzed are: [Please provide the list.](#)

(a) Distribution of trading volume on Ethereum



(b) Distribution of trading volume on Polygon

Figure 5: EVOLUTION OF V3 POOLS

Figure 5a presents the weekly number of concentrated liquidity (V3) pools with non-zero liquidity on Ethereum (solid blue curve) and Polygon (dashed red curve).

Figure 5b presents the weekly total total value locked (TVL) of all V3 pools on Ethereum (solid blue) and Polygon (dashed red) in \$U.S. Pool-level weekly TVL is the sum of the average values of the two assets in the pool throughout the week. Total weekly TVL is the sum of pool-level weekly TVLs across all V3 pools on a given blockchain.

Figure 5c presents the total weekly volume of transactions in all V3 pools on Ethereum (solid blue) and Polygon (dashed red) in \$U.S.

Figure 5d presents the total weekly number of transactions in all V3 pools on Ethereum (solid blue) and Polygon (dashed red) in \$U.S.

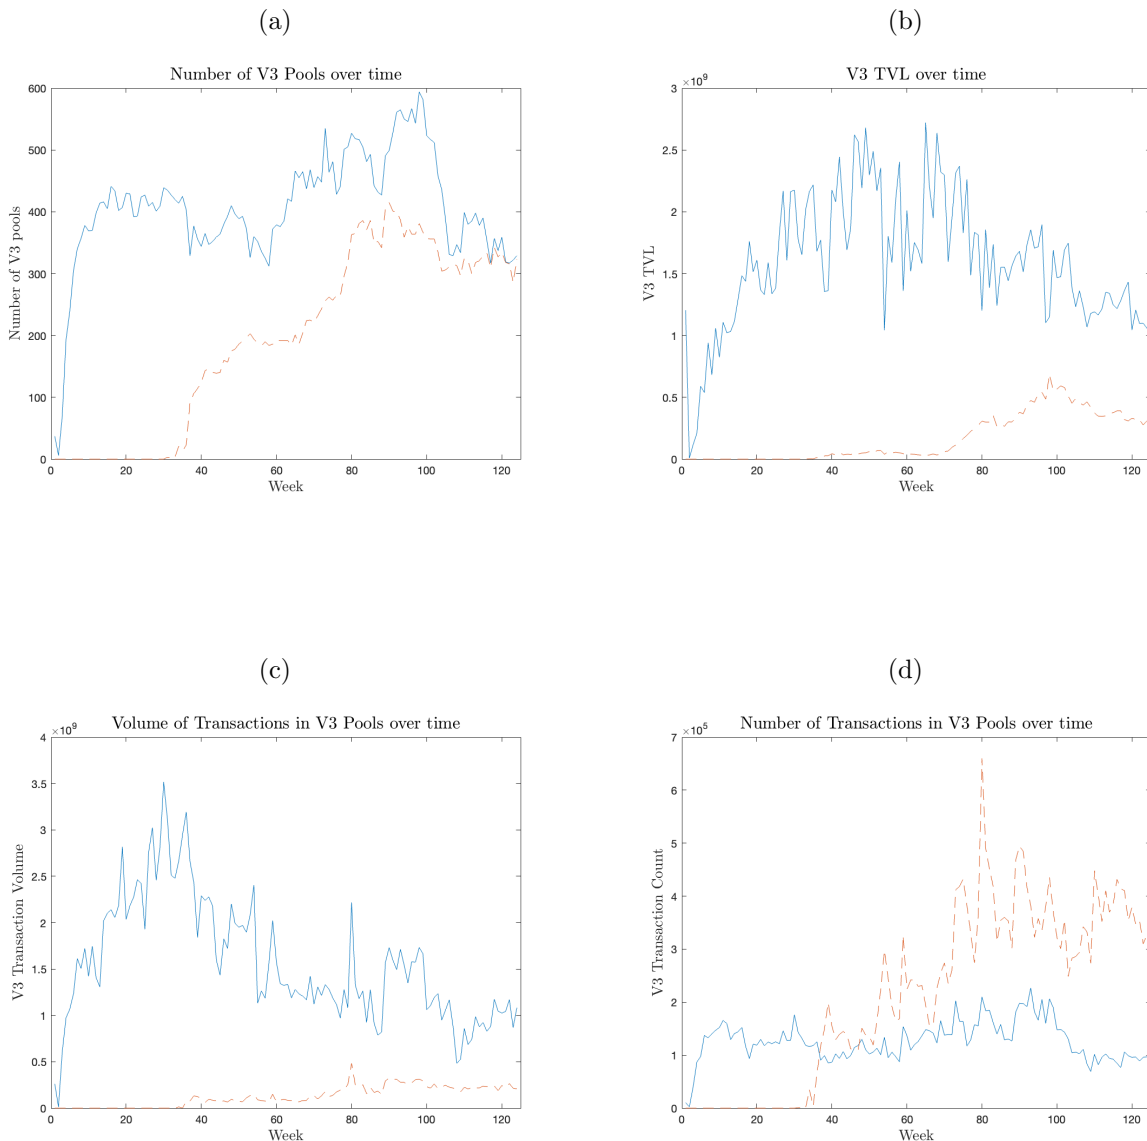


Figure 6: EVOLUTION OF CAPITAL EFFICIENCY

This figure presents weekly mean log capital efficiency of V3 pools deployed on Ethereum (solid blue) and Polygon (dashed red), for 125 weeks starting from May 5, 2023 and ending September 17, 2023. The estimation of capital efficiency is described in Section 4.






Figure 7: EXAMPLE OF AN EFFICIENCY-RESTORING ARBITRAGE TRANSACTION

This figure presents an example of an efficiency-restoring arbitrage transaction. The discussion of this transaction is found in Section 5.1.

-
- ▶ Swap 413.62512491381406395  rETH For 408.406665350687813761  wstETH On  Curve Finance
 - ▶ Swap 408.406665350687813761  wstETH For 462.044546737133279756 Ether On  Uniswap V3
 - ▶ Swap 446.142310925392871424 Ether For 413.62512491381406395  rETH On  Uniswap V3
-

[0x000032B9FF9E8143F6a8a28fB0e67d9739e22dd0](#) 

 [0x5e51328C0583094b76f28Cfd532AbC3d454fCfEA](#)  

↳ Transfer 15.895355449566982144 ETH From [Wrapped Ether](#) To [0x5e5132...454fCfEA](#)

↳ Transfer 15.895355449566982144 ETH From [0x5e5132...454fCfEA](#) To [eth-builder](#)

Figure 8: EVOLUTION OF EFFICIENCY-RESTORING ARBITRAGE

This figure presents weekly mean proportion of efficiency-restoring arbitrage transactions in V3 pools deployed on Ethereum (solid blue) and Polygon (dashed red), for 125 weeks starting from May 5, 2023 and ending September 17, 2023. The procedure for identifying relevant transactions is described in Section 5.1.

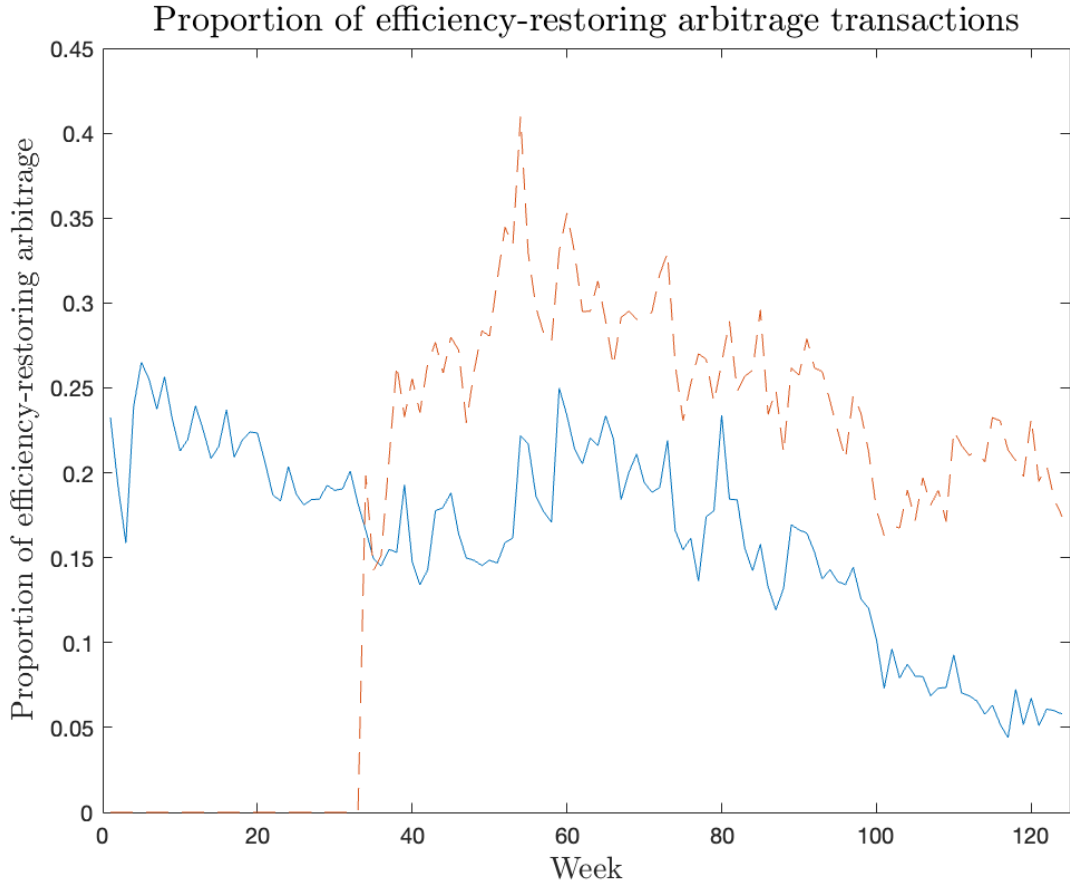


















Figure 9: EXAMPLE OF A SANDWICH ARBITRAGE TRANSACTION

This figure presents an example of a sandwich arbitrage transaction. The discussion of this transaction is found in Section 5.2.

- Frontrun**
 - ▶ Enable Ether as Collateral on  Aave Protocol V2
 - ▶ Supply 3,266.633166154861916194 Ether To  Aave Protocol V2
 - ▶ Borrow 5,266,312.3968  USDC From  Aave Protocol V2
 - ▶ Swap 5,266,312.3968  USDC For 2,677.244470663396227327 Ether On  Uniswap V3

- Victim Trade**
 - ▶ Swap 4,963,954.085663896321955049  DAI For 4,962,984.570046  USDC On  Uniswap V3
 - ▶ Swap 4,962,984.570046  USDC For 2,509.160889709629881976 Ether On  Uniswap V3
 - ▶ Swap 261,260.741350731385366055  DAI For 261,209.671706  USDC On  Uniswap V3
 - ▶ Swap 261,209.671706  USDC For 132.41071713520100745 Ether On  Uniswap V3









- Backrun**
 - ▶ Repay 5,266,312.3968  USDC To  Aave Protocol V2
 - ▶ Disable Ether as Collateral on  Aave Protocol V2
 - ▶ Withdraw 3,266.633166154861916194 Ether From  Aave Protocol V2
 - ▶ Swap 2,677.244470663396227327 Ether For 5,289,206.844108  USDC On  Uniswap V3
 - ▶ Swap 22,894.447308  USDC For 11.598027069242968987 Ether On  Uniswap V2

Table 1: SUMMARY STATISTICS OF CONCENTRATED LIQUIDITY (V3) POOLS

This table presents summary statistics of concentrated liquidity (V3) pools deployed on Ethereum blockchain in Panel A and Polygon blockchain in Panel B. The sample period is May 6, 2021 to September 17, 2023 (125 weeks). TVL stands for "total value locked", i.e. the amount of liquidity in the pool, computed as weekly average of the \$US value of the two assets in the pool. Transaction volume refers to the \$US value of trades in a pool throughout a week. Transaction count is the number of trades in a pool throughout a week. Volume-to-TVL is the ratio of weekly transaction volume to average weekly TVL. Average transaction size is the weekly average of \$US values of transactions in a pool. Average transaction to TVL is the ratio of average weekly transaction size in a pool to average TVL of the pool throughout the week. St. dev. Exchange rate is the standard deviation of the daily change in the logarithm of the ratio of prices of two assets in a pool, computed over 28 days prior to the observation week. ETH return is the difference between the logarithm of ETH price one day prior to the start of the week of an observation to the logarithm of ETH price 29 days prior to the observation week start. Gas price is the price of one unit of gas in 10^{-9} \$US.

	Min	Q1	Median	Q3	Max	Mean	St. dev.	Num obs.
Panel A: Ethereum								
TVL	84	95,091	357,320	1,358,238	1,115,790,411	3,928,214	23,916,018	49,451
Transaction volume	52	65,536	241,819	1,137,434	235,944,483	3,864,930	16,838,902	49,451
Transaction count	30	58	114	265	23,616	313.16	832.56	49,451
Volume-to-TVL	0.000	0.345	0.852	1.991	131.949	1.858	3.475	49,451
Avg transaction size	1.41	812.10	1998.82	5507.64	778063.84	7177.53	19433.81	49,451
Avg transaction to TVL	0.0000	0.0001	0.0009	0.0033	1.0558	0.0043	0.0136	49,451
St. dev. Exchange rate	0.03%	2.91%	4.84%	7.65%	41.04%	6.09%	5.39%	49,451
ETH return	-61.63%	-11.38%	-1.17%	12.90%	43.76%	-0.72%	19.51%	49,451
Gas price (10^{-9} \$US)	18.46	38.74	58.69	204.41	755.40	144.50	166.09	49,451
Panel B: Poygon								
TVL	0	1,990	15,530	148,237	225,180,047	878,000	7,248,649	24,070
Transaction volume	0	1,370	11,917	115,591	35,737,615	649,894	2,845,637	24,070
Transaction count	30	72	182	697	102,158	1,096.03	3,420.00	24,070
Volume-to-TVL	0.000	0.326	0.943	2.450	207.183	2.554	6.178	24,070
Avg transaction size	0.00	13.37	59.86	246.35	11898.57	242.27	517.99	24,070
Avg transaction to TVL	0.0000	0.0001	0.0005	0.0019	2.8662	0.0032	0.0221	24,070
St. dev. Exchange rate	0.03%	1.87%	3.15%	5.01%	22.25%	3.93%	3.41%	24,070
ETH return	-59.47%	-11.45%	-1.58%	6.00%	43.25%	-1.50%	16.93%	24,070
Gas price (10^{-9} \$US)	0.04	0.10	0.14	0.24	2.12	0.25	0.34	24,070

Table 2: SUMMARY STATISTICS OF CAPITAL ALLOCATION EFFICIENCY OF CONCENTRATED (V3) POOLS

This table presents summary statistics of measures of capital allocation efficiency in concentrated liquidity (V3) pools deployed on Ethereum blockchain in Panel A and Polygon blockchain in Panel B. The sample period is May 6, 2021 to September 17, 2023 (125 weeks). Mean (estimation) error refers to average, across all trades in a given pool in a given week, of absolute value of the distance between the dollar value of the actual amount out and amount out that would have been obtained in a hypothetical non-concentrated liquidity pool, normalized by the dollar price impact of the trade, as described in (4). Capital allocation efficiency is the ratio of TVL of a hypothetical V2 pool obtained as a result of solving the minimization problem in (4) to the average weekly TVL of the V3 pool, as in (5). \log (Capital allocation efficiency) refers to natural logarithm.

	Min	Q1	Median	Q3	Max	Mean	St. dev.	Num obs.
	Panel A: Ethereum							
Mean error	0.011%	0.738%	1.388%	2.637%	9.998%	2.046%	1.931%	49,451
Capital allocation efficiency	0.006	1.664	6.283	62.922	6,149.490	215.600	754.210	49,451
\log (Capital alloc. efficiency)	-5.160	0.509	1.838	4.142	8.724	2.316	2.579	49,451
	Panel B: Polygon							
Mean error	0.004%	0.391%	0.708%	1.334%	9.993%	1.180%	1.449%	24,070
Capital allocation efficiency	0.001	1.466	6.884	76.928	38,412.360	947.120	4,357.230	24,070
\log (Capital alloc. efficiency)	-6.792	0.383	1.929	4.343	10.556	2.420	3.120	24,070

Table 3: DETERMINANTS OF PRECISION OF CAPITAL ALLOCATION EFFICIENCY ESTIMATES

This table reports estimates, at the pool-week level, of linear regressions in which the dependent variable is mean estimation error, as defined in Table 2. The independent variables are defined in Table 1. The regressions are estimated on the pools deployed on Ethereum blockchain in Panel A and Polygon blockchain in Panel B. In even columns, regressions are augmented by pool and week fixed effects. Standard errors, reported in parentheses, are clustered at the pool and week levels.

	Panel A: Ethereum		Panel B: Polygon	
Intercept	2.804*** (0.111)	5.479*** (0.882)	1.445*** (0.075)	1.099** (0.450)
log (Transaction count)	-0.148*** (0.023)	-0.017 (0.019)	-0.070*** (0.020)	0.007 (0.016)
log (Transaction volume)	-0.070*** (0.013)	-0.059*** (0.011)	-0.028** (0.011)	-0.007 (0.008)
St. dev. exchange rate	10.887*** (0.359)	2.081*** (0.190)	16.101*** (0.631)	4.616*** (0.331)
Pool fixed effects	No	Yes	No	Yes
Week fixed effects	No	Yes	No	Yes
R squared	0.240	0.442	0.252	0.490
Number obs.	49,451	49,451	24,070	24,070

Table 4: SUMMARY STATISTICS OF BOT TRADING

This table presents summary statistics of measures of various types of bot trading activities in concentrated liquidity (V3) pools deployed on Ethereum blockchain in Panel A and Polygon blockchain in Panel B. The sample period is May 6, 2021 to September 17, 2023 (125 weeks). % of efficiency-restoring arbitrage is the fraction (volume-based in odd rows and number-of-transactions-based in even rows) of trades classified as efficiency-restoring arbitrage transactions out of all transactions in a pool during a week. The procedure for identifying efficiency-restoring arbitrage transactions is described in Section 5.1. % sandwich arbitrage is the fraction of trades classified as sandwich transactions out of all transactions in a pool during a week. The procedure for identifying sandwich arbitrage transactions is described in Section 5.2. % other bot trading is the fraction of trades initiated by MEV bots and not classified as efficiency-restoring arbitrage or sandwich arbitrage out of all transactions in a pool during a week.

	Min	Q1	Median	Q3	Max	Mean	St. dev.	Num obs.
Panel A: Ethereum								
% efficiency-restoring arbitrage								
Volume-based	0%	1.35%	7.29%	22.11%	100%	16.09%	21.24%	49,451
Count-based	0%	2.13%	7.52%	19.66%	100%	14.98%	19.59%	49,451
% sandwich arbitrage								
Volume-based	0%	0%	1.47%	7.44%	99.67%	5.95%	10.72%	49,451
Count-based	0%	0%	0.67%	2.30%	74.07%	1.64%	2.54%	49,451
% other bot trading								
Volume-based	0%	0.33%	4.96%	20.63%	100%	13.88%	19.08%	49,451
Count-based	0%	0.72%	5.26%	23.74%	100%	15.79%	21.36%	49,451
Panel B: Polygon								
% efficiency-restoring arbitrage								
Volume-based	0%	5.47%	14.49%	36.04%	100%	24.46%	25.63%	24,070
Count-based	0%	9.67%	24.75%	52.42%	100%	33.07%	27.91%	24,070
% sandwich arbitrage								
Volume-based	0%	0%	0%	0.20%	90.14%	1.17%	4.53%	24,070
Count-based	0%	0%	0%	0.18%	25.00%	0.39%	1.14%	24,070
% other bot trading								
Volume-based	0%	0.80%	5.77%	19.33%	100%	13.25%	17.52%	24,070
Count-based	0%	1.90%	6.53%	18.75%	100%	14.14%	18.36%	24,070

Table 5: DETERMINANTS OF CAPITAL ALLOCATION EFFICIENCY

This table reports estimates, at the pool-week level, of linear regressions in which the dependent variable is capital allocation efficiency, as defined in Table 2. The independent variables are defined in Tables 1 and 4. The measure of efficiency-restoring arbitrage is based on the fraction of volume of arbitrage transactions out of total transaction volume in a pool during a week (columns 1-2 and 5-6) and on the fraction of the number of arbitrage transactions out of total number of transactions in the pool during the week (columns 3-4 and 7-8). The regressions are estimated on the pools deployed on Ethereum blockchain in Panel A and Polygon blockchain in Panel B. In even columns, regressions are augmented by pool and week fixed effects. Standard errors, reported in parentheses, are clustered at the pool and week levels.

Arb measure:	Panel A: Ethereum				Panel B: Polygon			
	% arb volume		% arb transactions		% arb volume		% arb transactions	
Intercept	0.713*** (0.099)		0.891*** (0.162)		2.860*** (0.502)		3.378*** (0.496)	
% arbitrage	0.870*** (0.049)	0.637*** (0.081)	0.560*** (0.117)	0.334*** (0.107)	1.394*** (0.075)	0.151 (0.115)	1.907*** (0.170)	0.598*** (0.133)
% sandwich	0.558*** (0.096)	0.383*** (0.109)	0.279*** (0.786)	0.820* (0.497)	1.946*** (0.413)	0.777 (0.417)	0.694 (2.224)	3.948** (1.807)
% non-arb bot	-0.875*** (0.056)	-0.214** (0.084)	-0.795*** (0.135)	-0.396*** (0.091)	-2.130*** (0.112)	-0.780*** (0.130)	-1.876*** (0.370)	-0.920*** (0.146)
log volume-to-TVL	0.576*** (0.008)	0.439*** (0.011)	0.572*** (0.021)	0.439*** (0.011)	0.503*** (0.010)	0.429*** (0.016)	0.517*** (0.035)	0.433*** (0.016)
Rel. trans. size	-72.928*** (0.745)	-73.979*** (0.776)	-73.236*** (8.736)	-73.981*** (0.777)	-26.161*** (0.843)	-23.483*** (0.775)	-25.941 (16.297)	-23.476*** (0.774)
St. dev. exch. rate	-0.849*** (0.194)	-1.668*** (0.264)	-1.026** (0.452)	-1.717*** (0.264)	-2.185*** (0.561)	-4.122*** (0.770)	-3.851*** (1.359)	-4.210*** (0.769)
Lagged ETH return	0.139*** (0.052)		0.151*** (0.056)		0.095 (0.120)		0.098 (0.126)	
Log gas price (\$US)	-0.218*** (0.010)		-0.205*** (0.018)		0.015 (0.550)		0.062** (0.031)	
Pool fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
Week fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
R squared	0.250	0.391	0.247	0.390	0.144	0.404	0.161	0.405
Number obs.	49,451	49,451	49,451	49,451	24,070	24,070	24,070	24,070

Table 6: DETERMINANTS OF CAPITAL ALLOCATION EFFICIENCY: SHAPELLA UPGRADE

This table reports estimates, at the pool-week level, of linear regressions in which the dependent variable is capital allocation efficiency, as defined in Table 2. Post Shapella is an indicator variable equaling one for weeks 103-125 and equaling zero otherwise. Ethereum is an indicator variable equaling one (zero) for pools deployed on Ethereum (Polygon) blockchain. Post Shapella \times Ethereum is the product of the two indicators. Log gas price (\$US): Ethereum (Polygon) equals the natural logarithm of gas price for Ethereum (Polygon) pools and equals zero for Polygon (Ethereum) pools. The rest of the independent variables are defined in Table 1. The sample includes pools on Ethereum and Polygon blockchains that can be matched, i.e. a pool on Ethereum (Polygon) blockchain is included only if there is a pool with the same assets on Polygon (Ethereum) blockchain. The regressions are estimated using the full sample period in Panel A and using data from weeks 90-101 and 103-114 in Panel B. In even columns, regressions are augmented by pool fixed effects. Standard errors, reported in parentheses, are clustered at the pool level.

Sample period	Full sample		12 weeks pre and post	
Intercept	2.804*** (0.933)	1.649*** (1.780)	3.662*** (1.288)	2.802 (1.872)
Post Shapella	-0.361*** (0.124)	-1.286** (0.532)	-0.352** (0.154)	-0.229** (0.106)
Ethereum	-1.687 (1.047)	-1.613 (1.018)	-3.434* (1.988)	-0.435 (1.906)
Post Shapella \times Ethereum	-0.685*** (0.170)	-0.751*** (0.099)	-0.387* (0.212)	-0.424*** (0.158)
log volume-to-TVL	0.454*** (0.049)	0.440*** (0.019)	0.503*** (0.055)	0.350*** (0.037)
Rel. trans. size	-85.290*** (9.405)	-73.516*** (1.894)	-89.841*** (10.735)	-76.840*** (3.389)
St. dev. exch. rate	7.029*** (2.388)	-1.114 (-1.215)	6.670** (3.192)	3.756* (2.111)
Lagged ETH return	0.242* (0.136)	-1.900 (1.839)	-0.442 (0.352)	-0.295 (0.340)
Log gas price (\$US): Ethereum	-0.097 (0.061)	-0.298*** (0.067)	-0.216 (0.156)	0.022 0.149
Log gas price (\$US): Polygon	0.027 (0.060)	-0.091 (0.063)	0.082 (0.085)	0.096 (0.077)
Pool fixed effects	No	Yes	No	Yes
R squared	0.197	0.475	0.206	0.543
Number obs.	11,524	11,524	3,963	3,963