



BIS Working Papers

No 1227

Decentralised dealers? Examining liquidity provision in decentralised exchanges


By Matteo Aquilina, Sean Foley, Leonardo Gambacorta
and William Kregel

Monetary and Economic Department

November 2024

JEL classification: D47, G14, G23.

Keywords: Market design, market making, liquidity,
automated market maker, decentralized finance



BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

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

© *Bank for International Settlements 2024. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

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

Decentralised dealers?

Examining liquidity provision in Decentralised Exchanges*

Matteo Aquilina^{1,2} Sean Foley² Leonardo Gambacorta^{1,3} William Krekel²

¹Bank for International Settlements, Switzerland

²Macquarie University, Australia

³CEPR

Abstract

Decentralised exchanges allow participants to buy and sell assets without the need for intermediaries, in theory *democratising* liquidity provision. However, using data from the largest decentralised exchange, we show that liquidity provision – rather than being the purview of a diffused set of market participants – is instead confined predominantly to a small group of sophisticated ones. These participants submit orders that mimic bids and asks and are able to extract significantly higher profits (both in absolute and relative terms) compared to their unsophisticated counterparts. They also exhibit considerable skill, extracting higher profits during periods of high volatility by capturing a higher share of trading without incurring additional adverse selection.

Keywords: Market Design, Market Making, Liquidity, Automated Market Maker, Decentralized Finance

JEL Codes: D47, G14, G23

*We thank Luke Johnson for valuable assistance with the underlying data. We thank Inaki Aldasoro, Douglas Araujo, Rodney Garratt, Wenqian Huang, Andreas Schrimpf, Hyun Song Shin, Vladyslav Sushko, participants to the 7th Cryptocurrency Research Conference in Dubai and a Bank for International Settlements (BIS) seminar for comments and suggestions. We are especially thankful to an anonymous referee for insightful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect those of the BIS. Krekel thanks the Australian Government Research Training Program Scholarship and the BIS, which funded his period of study during this research.

“A decentralized exchange [...] is a peer-to-peer marketplace where transactions occur directly between crypto traders. DEXs fulfill one of crypto’s core possibilities: fostering financial transactions that aren’t officiated by banks, brokers, payment processors, or any other kind of intermediary.”

Coinbase - A leading crypto exchange

1 Introduction

Decentralized finance (DeFi) refers to a set of new applications in the crypto-asset space designed to disintermediate the functions fulfilled by traditional financial service intermediaries. In DeFi, computer code (*smart contracts* in DeFi terminology) is deployed on the blockchain and automatically performs certain actions when some pre-determined conditions are met.¹ DeFi applications (DApps) were enabled by the launch of the Ethereum blockchain in 2015, but the DeFi ecosystem is even newer, only gaining noticeable traction from 2020.

DApps have been built to mimic and replicate many of the functions carried out by intermediaries in traditional finance (Aquilina et al., 2024). Some applications allow users to borrow and lend their cryptocurrencies, others focus on asset management or insurance, another category allows participants to develop derivatives and synthetic assets. However, one of the innovative applications that arguably exhibits the most potential is the development of decentralised exchanges (DEXs). Currently, DEXs are responsible for more than 10bn US dollars of transactions of digital assets every day, with this figure continuing to grow as these markets gain further adoption.

In contrast to what happens on a traditional stock exchange, trades on a DEX - typically referred to as *swaps* - take place directly between two counterparts without the need for the custody of assets to be passed on to a broker, a member of the exchange or a clearinghouse. As a result, there are no traditional market makers that carry inventory to intermediate between buyers and sellers. Previous studies have highlighted that DEXs have the potential to be used in a number of settings, either complementing or competing with more traditional exchanges based on central limit order books (Capponi & Jia, 2021; Foley, O’Neill, & Putnins, 2023).

In DEXs, market participants are incentivised to provide liquidity through fees that are paid to liquidity providers (LPs) who commit their assets to a liquidity pool. Access to these pools is possible for all market participants and is not restricted to any subset of agents - unlike the high barriers to entry that exist in many traditional financial markets. Hence liquidity provision is, at least in theory, *democratised*. Anyone

¹For a description of the different components of DeFi, see Schär (2021).

with access to the relevant blockchain can commit their assets to a liquidity pool and earn a share of the relevant liquidity fees. DeFi proponents argue that this is an inherent advantage as it eliminates – or at least substantially reduces – the need for intermediation in financial markets. As highlighted in the quote at the beginning of this paper, removing intermediaries is a core stated objective of DeFi.

However, there may be economic forces at play that favour the emergence of sophisticated liquidity providers, even in a market design where access itself is unrestricted. Differences in skills, economies of scale and specialisation advantages may favour a relatively small numbers of participants in their intermediation activities and make it difficult for unsophisticated participants to compete.²

In traditional financial services different ways of intermediating trades characterise different asset classes: from the mostly dealer-based intermediation in the corporate bond market, to high-frequency liquidity provision in equity and foreign exchanges markets. Notwithstanding these differences, all these set ups rely on a limited number of intermediaries for the provision of liquidity. In over the counter markets, those participants that have the ability to connect willing counterparties to each other have emerged as dealers (Duffie et al., 2005); in markets characterised by electronic trading, liquidity providers tend to compete on technological sophistication and speed (Aquilina et al., 2021; Budish et al., 2015, 2024). It is therefore interesting to analyse whether the innovation brought about by DEXs is capable of overcoming this shortcoming of traditional finance.

In this study we analyse liquidity provision in one of the largest DEXs (Uniswap V3) and provide insights into the behaviour of market participants in terms of liquidity provision, trading and profitability. Notwithstanding the decentralised nature of the system, we show that a subset of participants behave in a manner similar to market makers and dealers (which we label *sophisticated participants*) in traditional finance have emerged in the DeFi space and that it is difficult for unsophisticated (*retail*) participants to compete with them. In particular:

- The vast majority of liquidity (65-85%) is provided by sophisticated participants, posting orders that are tradeable in a price range relatively close to the prevailing market price, mimicking the posting of bid and ask orders in traditional exchanges.
- Retail liquidity providers manage their liquidity positions much less actively: they interact with substantially fewer liquidity pools and adjust their positions less often.
- Retail participants capture a much smaller share of trading fees (10-25%) and — while they suffer less adverse selection — earn substantially lower returns on their invested capital than their sophisticated

²For example, Cong et al. (2023) find significant concentration of activity in the Ethereum ecosystem as a whole.

counterparts.

- Retail participants exhibit substantially lower skill. They are much less profitable in highly volatile periods and do not seem to adapt their liquidity provision to changing market conditions.

Our paper contributes to the rapidly developing literature on decentralised exchanges and decentralised finance more generally. Heimbach et al. (2022) closely align with our study: they analyse the choice of liquidity providers in Uniswap V3 from a computer science perspective. However, they focus only on a limited number of currency pairs and do not disentangle the contribution of different types of investors. Capponi, Ruizhe, and Shiao (2023) analyse price discovery in DEXs and find that high-fee trades reveal more private information, as highly informed traders compete with each others to capture additional fees. Barbon and Ranaldo (2023) compare centralised and decentralised exchanges and find that they have a similar level of transaction costs while Lehar and Parlour (2021) discuss the differences between exchanges utilizing central limit order books (CLOB) versus DEXs based on automated market makers, analysing the behaviour of liquidity providers across the two types of markets. Lehar et al. (2023) analyse the fragmentation of liquidity in different DEX pools. Our contribution to this literature involves an in-depth analysis of the behaviour of sophisticated and retail investors. We examine their characteristics across a diverse range of currency pairs and liquidity pools and document the emergence of *sophisticated* LPs across the entire cross-section.

The remainder of the paper is structured as follows: Section 2 describes the functioning of DEXs and automated market makers, Section 3 describes the source of the data and the methodology used, Section 4 reports our results and Section 5 concludes.

2 An introduction to DEXs

The first (order-book-based) DEXs were launched in 2016, but it was not until 2018 that they gained more traction with the introduction of automated market makers (AMMs), notably through the Uniswap protocol. In line with the general *ethos* of crypto and DeFi, the stated objective of DEXs was to remove the need for intermediaries in the trading process and allow all participants to provide liquidity, thereby *democratising* this aspect of financial markets.³

2.1 How do AMMs work?

The implementation of an AMM is realized through multiple smart contracts residing on a blockchain.⁴ Trading in AMMs occurs via liquidity pools, typically comprising two asset reserves. A liquidity provider (LP) contributes its assets to these pools, earning trading fees in return, distributed proportionally based on the LP's liquidity share in a pool. Traders execute swaps between assets, paying a fixed percentage fee that remunerates LPs for their inventory and adverse selection risks.⁵

In AMMs, no centralized counterparty takes custody of assets and trading is facilitated in a trustless, atomistic manner, eliminating the counterparty default risk that has plagued CEXs such as the now defunct FTX. Other advantages include quasi real-time settlement and consistently available liquidity (Aspris et al., 2021). Part of the reason for AMMs' growing popularity is the potential to earn passive income as a LP. In providing liquidity to trading pools, LPs are able to earn income on their staked assets in the form of trading fees. The total-value-locked (TVL) in the Uniswap protocol in June 2024 amounted to over USD6 billion, representing about 6% of all TVL in decentralized finance.⁶ Transactions in AMMs are settled through the inclusion into a block that is appended to the respective blockchain (in our case Ethereum). Blocks are added in discrete time intervals, called the 'block time'. Broadly speaking, for the time interval of the block time, transactions are considered to have arrived at the same time and are generally ordered by economic principles i.e. agents paying a high enough network fee (called 'gas' in Ethereum) that rewards block validators. Importantly, no market participant is able to jump to the front of the queue, solely based on speed advantages as is the case in traditional exchanges.⁷ Instead, transactions within a block are executed in a batch based on an inclusion process similar to pure price-priority in CLOB markets. Xu et al. (2023) give an

³For instance, the Uniswap protocol is described as follows by its developers: "The Uniswap Protocol is an open-source protocol for providing liquidity and trading ERC20 tokens on Ethereum. It eliminates trusted intermediaries and unnecessary forms of rent extraction, allowing for safe, accessible, and efficient exchange activity." See <https://uniswap.org/faq>

⁴At the time of writing most of the trading volume is routed through AMMs deployed on the Ethereum chain. Other notable blockchains hosting AMMs are the Binance Smart Chain, Avalanche and Solana.

⁵In Uniswap V3, liquidity pools can trade at fee tiers of 1, 5,30 and 100 bps. For further insights into the mechanics of fixed fee liquidity pools, we refer to Foley et al. (2024).

⁶Data from DefiLlama retrieved from <https://defillama.com/dexs>.


⁷The provision of just-in-time liquidity as described below can be seen as the equivalent behaviour in DEXs.

overview and background on different classes of AMMs. We provide a short introduction on how basic AMMs (following the constant product mechanism, such as Uniswap V2), work. For a more theoretical discussion we refer to Angeris and Chitra (2020). Further, we describe the main characteristic that differentiate Uniswap V3 from the previous generation of AMMs, namely the introduction of ‘concentrated liquidity’.

2.1.1 The ‘classic’ constant-product AMM

Constant Product Market Makers (CPMMs), the pioneering AMMs in the cryptocurrency domain, generate prices using a constant product function, commonly denoted as $x * y = k$. Here, x and y denote the quantities of Token X and Y in the pool reserves, and k is their product. Notably, k ’s value remains unchanged unless LPs alter liquidity in either token reserve. This straightforward mathematical relationship guarantees that the pool’s token price is governed by supply and demand. Figure 1 demonstrates this pricing mechanism. The initial pool allocations x_0 and y_0 set the price $P_0 = \frac{y_0}{x_0}$. If a liquidity taker intends to purchase Δx Token X, she will extract Δx Token X from the pool and contribute Δy Token Y. The pool’s updated liquidity state maintains the constant k and is $x_1 = (x_0 - \Delta x_0)$ and $y_1 = (y_0 + \Delta y_0)$, resulting in the new price $P_1 = \frac{y_1}{x_1}$. The asymptotic and convex shape of the pricing curve has two important implications. Firstly, larger trades result in greater price impacts. Secondly, emptying a token reserve is infinitely costly, ensuring perpetual liquidity. Park (2023) provides a comprehensive discussion of CPMMs.

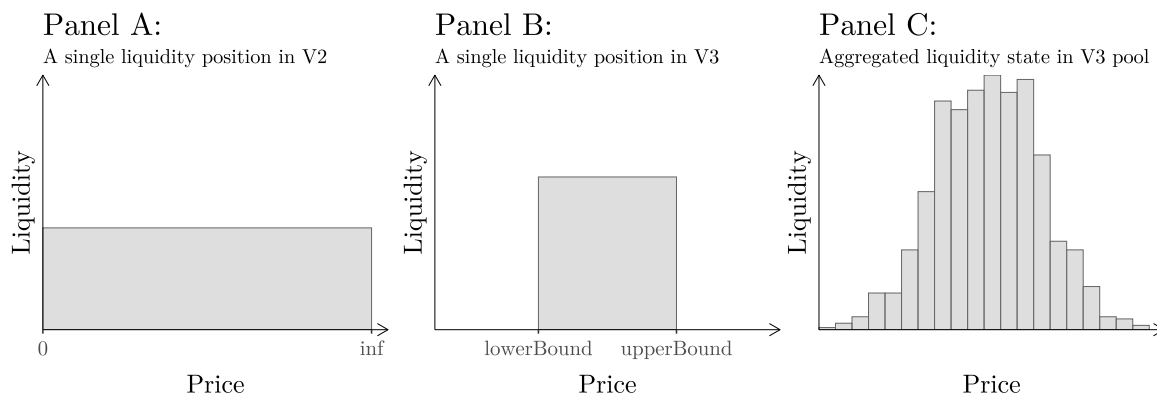
Figure 1: Bonding Curve of a CPMM



2.1.2 The new generation of ‘concentrated liquidity’ AMMs

In a regular CPMM (such as Uniswap V2), LPs distribute their liquidity across the whole range of possible prices (Figure 2, Panel A). With the introduction of Uniswap V3 in May 2021, the AMM landscape underwent a major architectural upgrade when Uniswap V3 introduced the concept of ‘concentrated liquidity’. In this set up, the pricing curve is divided into small, discrete steps known as ‘ticks’, which are converted into an asset’s price using a deterministic formula. When creating a liquidity position, LPs select a specific tick range, defined by a lower and upper limit (Figure 2, Panel B), in which they provide liquidity in. To characterise the width of a single liquidity position, we introduce a metric called *tickrange spread* that is the range of the position scaled by its midpoint: $TickRangeSpread = \frac{upperBound - lowerBound}{0.5 * (upperBound + lowerBound)}$. Notably, ‘swaps’ in a pool are not limited to these tick ranges, unlike the ‘tick size’ in a conventional market. The global liquidity state of a pool is obtained by aggregating individual liquidity positions across all LPs and ticks (Figure 2, Panel C). Concentrating liquidity typically results in greater capital efficiency, as liquidity is deepest where it is most required, at or near the current market price. On the other hand, liquidity is thinner at prices further away from the current market price. Theoretically, this leads to smaller price impacts and slippages for traders when prices are relatively stable. The price range across which assets are deployed is fixed to a liquidity position. A position is considered active if the position’s tick range includes the current market price. Importantly, LPs only earn fees and interact with incoming transactions if a swap takes place within their respective price range. Specifically, LPs earn fees according to a pro-rata mechanism depending on their proportion of provided liquidity at a tick. Within a tick, this type of AMM operates similarly to a CPMM. At the position bounds, the asset composition of a position shifts entirely to the less valuable of the two provided assets, exposing LPs to adverse selection (also known as impermanent or divergence loss in this context). If the price exceeds a position’s lower or upper limit, the position consists entirely of one asset. The position then becomes inactive, does not contribute to the global active liquidity state, and is not eligible for any fee reward. If the price re-enters a position’s range, the position becomes active again.

Figure 2: “Liquidity state in Uniswap V2 and V3” adapted from Adams et al. (2021). This figure exemplifies the liquidity states in Uniswap V2 and V3. Panel A shows the liquidity state in a Uniswap V2 pool that provides liquidity across the whole price range. Panel B shows the liquidity distribution of a single, concentrated liquidity position in Uniswap V3 that distributes its liquidity in a price range between a lower bound and an upper bound. Panel C depicts the aggregated liquidity state of multiple liquidity positions in a Uniswap V3 liquidity pool.



Crucially, LPs only receive fees when a swap takes place within their respective price range. As a result, an LP must find an economic equilibrium between monitoring costs (to keep their position active) and the expensive relocation of a position to a different price range.⁸ We explore the dynamics of concentrated liquidity positions in Section 2.2.1 in more detail. Altogether, LPs now compete for strategic liquidity provision via various allocation parameters, unlike in UniswapV2 when all fees were distributed pro-rata across all positions in the entire possible price range. In summary, LPs now have access to new tools that enhance their flexibility. However, while liquidity provision has become more adaptable, it has also grown more complicated.

2.1.3 Transaction costs

To provide liquidity on a DEX, and indeed for any other transaction on the Ethereum blockchain, users need to pay a fee, the so-called ‘gas’ fee, that compensates block validators for adding a transaction into a block. This cost is generally dependent on two components. First, the *gas* that is used to complete a transaction. This amount mirrors the computational difficulty of a transaction that is larger for more complex and lower for simpler ones. For example, the liquidity provision of a DEX uses multiple input parameters that are processed by smart contracts and is therefore computationally more expensive than a simple ETH transfer between two addresses. Importantly, for the same transaction (and hence computational difficulty), this parameter is time-invariant.

The second component is the *gasPrice*. The *gasPrice* reflects the cost of a unit of *gas* and is variable. For example, during times of high demand, transactions with a too low *gasPrice* might need to wait longer to

⁸We refer to Lehar et al. (2023) and Caparros et al. (2023) for a discussion on the implications of gas costs for liquidity provisioning in Uniswap V3.

be included in a block as block validators sort and process transactions according to their profitability and may only include higher fee-paying transactions. The total transaction cost is then given by the product of the amount of gas used ($gasUsed$) and the price paid for each unit of gas ($gasPrice$): $TX_{cost} = gasUsed \times gasPrice$.

2.2 Liquidity provision strategies in modern AMMs

The flexibility introduced by the new generation of AMMs allows for significantly more sophisticated liquidity provision strategies compared to their earlier counterparts. In this section, we describe the four strategies, which we label respectively *concentrated*, *unconcentrated* (or *V2-like*), *range order* and *just-in-time* (JIT) liquidity strategies.

2.2.1 Concentrated strategy

LPs opting for a concentrated strategy leverage the newly introduced feature of Uniswap V3. and choose a lower and an upper bound for the position in which they wish to provide liquidity in. To earn fees from swaps, the swap price must be situated in this range. Figure 3, Panel A illustrates the mechanics of concentrated liquidity provision using a hypothetical position.

Suppose at time t_0 the market price in a ETH/USDC pool is at 4,000 ETH/USDC.⁹ An LP endowed with equal values of ETH and USDC decides to provide liquidity with a lower bound of 3,000 ETH/USDC and an upper bound of 5,000 ETH/USDC. The position’s tickrange spread is 50% ($TickRangeSpread = \frac{upperBound - lowerBound}{0.5 * (upperBound + lowerBound)} = \frac{5000 - 3000}{0.5 * (5000 + 3000)} = 0.5$). Half of the LP’s assets are posted as USDC below the market price and the other half as ETH above the market price. At t_1 the price has increased to 4,750 ETH/USDC as liquidity demanders bought ETH from the AMM (and LPs sold ETH), driving up the price. At this point in time, the composition of assets of the single liquidity position is no longer evenly distributed. Instead, the position has progressively sold its ETH reserves and is left with a higher proportion of the less-valuable asset, USDC. Past the position’s upper bound (i.e. 5,000 ETH/USDC), the position consists entirely of USDC. Between t_0 and t_1 the swap price remained within the LP’s position bounds and the LP is eligible for fee revenue. The exact fee revenue depends on two factors: the size of the executed swaps (as liquidity demanders pay a fee proportional to the size of their trades) and the amount of liquidity the total fee is shared with. In practice, many other LPs likely provided liquidity at or around the prevailing market price and facilitated trading. At t_2 the market price has moved above the upper position bound. The position is ‘out of range’ and is no longer eligible for any fee revenue.

⁹Decentralized finance applications built on the Ethereum blockchain typically use the ERC20-compliant version called wrapped ETH (wETH), which is a fully fungible representation and can be interchanged at a 1:1 rate.

The LP now has two options: 1) wait for the market price to move back in range or 2) burn the position and mint a new one around the new market price, incurring network costs.¹⁰ It becomes apparent that the choice of the tickrange spread has important implications. A larger tickrange spread results in a position that is more immune to price swings, ensuring it remains in an active, in-range state for longer. In contrast, a narrower tickrange spread that concentrates most of a position’s liquidity to only a few ticks acts as a form of leverage, and entitles the LP to a larger proportion of total fees, given the pro-rata distribution mechanism within a tick.¹¹

2.2.2 Unconcentrated or ‘V2-like’ strategy

While more recent AMMs provide additional options, LPs can still provide liquidity across the whole range of potential prices, by using an infinite price range, i.e. from 0 to ∞ (see Figure 2, Panel A). However, by providing ‘unconcentrated’ liquidity, LPs choose inferior capital efficiency as they are unlikely to capture a substantial amount of trading fees. Yet, LPs might prefer to set an unconcentrated position to avoid incurring any monitoring costs.

2.2.3 Range Order strategy

Range orders mimic traditional limit orders used in stock exchanges and CEXs. In this strategy, LPs provide one-sided liquidity in extremely tight positions that are initially out of range, where the position bounds mirror the targeted limit price. Once the market price crosses the outer bound of the position it is converted into the other asset, equivalent to a limit order being filled. Importantly, the LP needs to withdraw the position after the price has crossed its position’s bound or risks that the position is reverted into the original state, i.e. becomes unfilled. Rather than paying a swap fee for the transaction, the LP instead receives compensation for the liquidity it provided, similar to a maker fee in traditional finance. However, the LP also needs to pay network costs twice (once for the creation and once for the withdrawal of the liquidity position).

Figure 3, Panel B illustrates the mechanism of a range order to sell ETH for USDC at a target (limit) price of about 4,190.5 USD per ETH. Suppose the price between t_0 and t_3 steadily increases. At t_0 the price is 4,188 and the LP mints a position that exclusively consists of ETH in a narrow range between 4,190 and 4,191.¹² Importantly, at t_0 the position is out of range. At time t_1 the price reaches the position’s lower

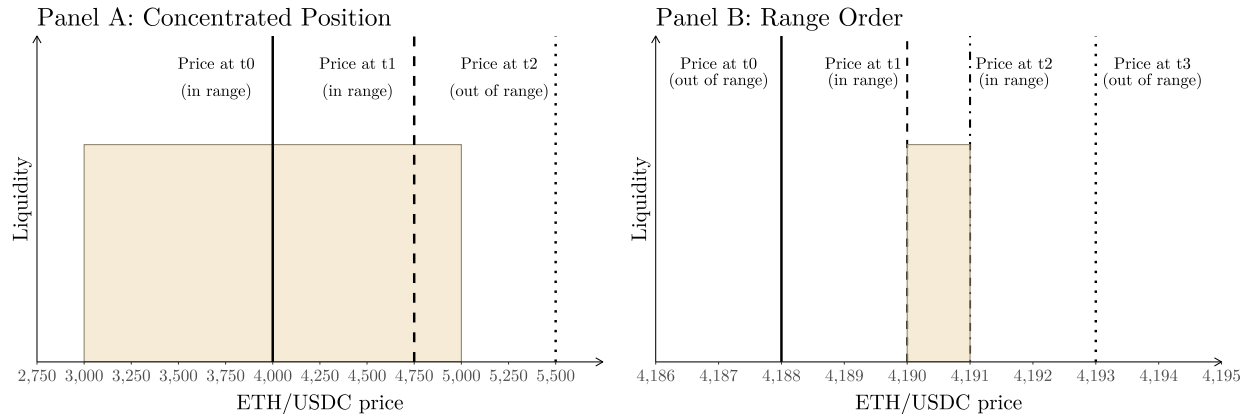
¹⁰We refer to Caparros et al. (2023) for a detailed discussion on the repositioning of LPs in AMMs.

¹¹We refer to Barbon and Ranaldo (2023) for a further discussion and mathematical explanation on how concentrating liquidity equates to a form of leverage.

¹²The minimum tick size in which a liquidity position can be posted is dependent on the liquidity pool’s fee levels, which are approximately as follows: 1) 1bps fee level: 1bps tick size, 2) 5bps fee level: 10bps tick size, 3) 30bps fee level: 60bps tick size, 4) 100bps fee level: 202bps tick size (Adams et al., 2021).

bound of 4,190. The price is now in range, and the position becomes active. Between t_1 and t_2 the position is gradually converted and at t_2 consists of USDC only. Once the prices moves out of range at t_3 , the LP needs to burn the position to avoid any risk of the position being converted back to ETH.

Figure 3: Liquidity Provision strategies in Uniswap V3. This figure illustrates two liquidity provision strategies in Uniswap V3: Panel A shows a concentrated position with a lower bound of 3,000 and an upper bound of 5,000; Panel B shows the set up of a range order, at an approximate limit price of 4,190.5.



2.2.4 Just-in-time strategy

In JIT transactions, sophisticated LPs provide liquidity immediately before large trades and are able to earn most related fees. They do so by observing incoming trades in the blockchain’s mempool and bundling them with the creation (and subsequent removal) of a liquidity position. The JIT LP is able to earn the majority of an incoming orders fee revenue by posting a large amount of liquidity in an extremely narrow range, most often only around the tick in which the swap is expected to execute. Through this mechanism, the JIT LP collects most of the accruing fees given their significant share of liquidity at the corresponding tick.

JIT transactions are generally beneficial for liquidity demanders, as they increase the depth of the order book and reduce slippage, similar to midpoint dark pools in traditional finance markets.¹³ However, JIT liquidity is detrimental to existing LPs who suffer from fee dilution, as the majority of fees related to the attacked trade will be channeled to the JIT LP only. In the extreme, a sufficiently large proportion of JIT liquidity can almost completely ‘crowd-out’ non-JIT liquidity. For a theoretical discussion on JIT liquidity and its impacts on liquidity providers we refer to Capponi, Jia, and Zhu (2023).

¹³See Foley and Putniņš (2016) for further discussions of the merits of dark pools in traditional financial markets.

2.3 Comparison of DEXs vs. CEXs

DEXs are not the only type of exchange that characterises the trading of cryptocurrencies. Centralized exchanges (CEXs) utilizing a CLOB are the other – and still more widely used – structure currently in use to trade crypto-assets. Trading takes place differently in these two types of exchanges and as CEXs are more familiar it is interesting to compare and contrast some of the main characteristics of these two types of exchanges. The basis for the following discussion is a public, permissionless DEX on a public blockchain, the type of DEX that currently facilitates the majority of cryptocurrency trading.¹⁴

To trade on CEXs, agents transfer their funds and (consequently custody) to the exchange. Conversely, users of DEXs retain custody of their assets at all times. This limits their exposure to security incidents like hacks or bankruptcies, such as the now infamous collapse of FTX in November 2022. After depositing assets, traders on CEXs benefit from (almost) real-time trading, liquid order books and advanced functionality. Users of registered CEXs typically need to comply with KYC and AML laws. Therefore, proof of identification is needed during the registration process. In contrast, trading on DEXs generally comes without a formal registration process. On the other hand, DEXs suffer from lower speed, which is typically constrained by the block time of the respective blockchain.¹⁵ Moreover, they only offer a rather basic trading experience in terms of functionality and order types. However, customers of CEXs need to trust the proprietary central entity, whereas DEXs are trustless due to their distributed and open-source nature (Han et al., 2022). Trading fees on CEXs are generally a fixed percentage of the volume and may vary if the order adds liquidity (maker) or takes liquidity (taker) from the order book. Liquidity takers on DEXs additionally pay a flat fee to the network (‘gas costs’ in the case of Ethereum) for settlement. The universe of traded securities on CEXs is constrained by the decisions of each exchange to list specific coins, whereas permissionless DEXs offer nearly endless opportunities, with the potential to almost instantaneously list any asset immediately after it is created on the blockchain. Further, CEXs are prone to potential censorship or simple server outages. In contrast, the distributed nature of DEXs make it nearly impossible to shut down exchange activity.

¹⁴We acknowledge that future implementations and current prototypes such as the one used in the BIS’s Project Mariana partly exhibit different characteristics. The traditional, limit-order-book exchanges contrasted include large existing CEXs such as Binance, Coinbase and Kraken.

¹⁵For a detailed analysis on the speed of CEXs, we refer to Foley, Krekel, et al. (2023).

Table 1: Characteristics of DEXs vs. CEXs. This table summarises key differences between DEXs such as AMMs and more traditional CEX that follow a limit-order-book model.

	Centralized Exchange	Decentralized Exchange
Transaction ledger	Private	Public
Custody of funds	Custodian (often exchange itself)	User
Trade matching	Matching Engine	Smart Contract
KYC/AML	Yes	No
Trustless	No	Yes
Trading universe	Listed by exchange (permissioned)	Listed by anyone (permissionless)
Execution Speed	Fast	Slower (depending on implementation)
Fees	Proportional to size	Proportional to size and network fee ('flat fee')
Functionality	Advanced order types, margin trading	Basic but improving
Censorship	Possible	Not possible

3 Data and methodology

3.1 Data

We parse the raw transaction logs of Uniswap V3 liquidity pools from their creation (Uniswap V3 launched on 5 May 2021) until 1 January 2024 from an Ethereum Archive node. Transaction data for the trading of assets, the creation and removal of LP positions (known as “swaps”, “mints” and “burns”, respectively) include the transaction hash, the block number a transaction is included in, the corresponding UNIX timestamp, the wallet address of interacting agents / smart contract addresses and the corresponding amounts in token currencies. Liquidity position data contains the timestamp and block number of the mint/burn transaction, the wallet address of the owner, the liquidity state (assets of the position) and the ID that makes the respective liquidity position uniquely identifiable, as well as the lower and upper tick bounds of the position. We then convert trading volume to USD using hourly prices sourced from CCData¹⁶ and rank liquidity pools according to their lifetime trading volume. We limit our analysis to the top 250 pools that constitute 96% of total trading volume, as the computational cost of adding all other pools would not be matched by any additional benefits in terms of adding to our understanding of the market.

We use this data to reconstruct the state of each liquidity pool over time. This exercise is similar to reconstructing an orderbook of a CLOB market at all orderbook levels after each update. Within a pool, the non-trivial computational complexity increases with the amount of liquidity positions and the amount of initialized ticks liquidity is provided across. For example, a USDT-USDC pool, where liquidity is concentrated around \$0.98-1.02 will be far less computationally intensive than an ETH-USDC pool with the same number of trades, simply because the latter will require the evaluation across a much larger number of ticks. We replay each swap and assess which liquidity position(s) facilitate the transaction. This reconstruction

¹⁶API access is available via <https://developers.cryptocompare.com/>.

makes it possible to calculate the individual LP’s earned fees from each swap and allows us to identify the dynamic asset composition of each liquidity position. We exclude liquidity positions that are worth less than 1 USD at time of the mint and aggregate the transaction level data to the daily level that is used for further analysis by using the last observation per day for each liquidity position.¹⁷ We also source the amount of *gasUsed* and the *gasPrice* paid for each liquidity-related transaction from the raw transaction logs.

3.2 Identifying ‘sophisticated’ and ‘unsophisticated’ market participants

The main objective of our study is to assess whether liquidity provision can be *democratised* in a DEX or whether it remains in the hands of a small set of sophisticated LPs. We therefore need to establish whether these sophisticated investors exist in the first place and whether they are able to outcompete retail ones. In contrast to traditional finance, where the identity of the counterparties involved in a trade is typically only known to the traders themselves, the regulator and the exchange,¹⁸ the transparent nature of the blockchain allows everyone to identify the addresses involved in transactions. While the same entity may well control multiple wallets (thereby masking the ultimate decision maker involved in trades), DeFi makes the job of researchers that want to track market participants’ activity easier because the activity of wallets is publicly broadcast on the blockchain.

Leveraging the transparent nature of the blockchain, we use two approaches to separate sophisticated from retail investors. The first involves analysing the characteristics of a wallet to infer the category to which it belongs. The second is more direct and exploits the fact that, in some cases, it is possible to link a wallet with a specific sophisticated investor.

Within the first approach we focus on five different characteristics, some of which relate to the amount of capital committed to liquidity provision and others that relate to technological sophistication. We use two metrics to measure capital intensity: 1) The maximum mint value in USD per wallet and then, similar to Cornelli et al. (2024), we use the 95th percentile as our cutoff to classify wallets as ‘sophisticated’. 2) We assign all wallet addresses that have minted a position of at least USD1 million in value as sophisticated. In terms of technological sophistication we identify sophisticated wallets as those that are in the top 5th percentile of: 1) the total number of liquidity positions minted, 2) the number of distinct pools a wallet has provided liquidity in and 3) the total number of liquidity transactions. If a wallet address provides liquidity through multiple positions and different pools simultaneously, it is likely that it would use a more

¹⁷The high computational cost of this task has resulted in very few authors attempting to identify granular fee revenue at the LP level in Uniswap V3. Studies in computer science have come up with a solution, but are either limited to a small number of pools and shorter sample periods (Heimbach et al., 2022) and / or rely on pool aggregate numbers (Loesch et al., 2021).

¹⁸In some cases, such as for equities traded on exchanges the identity is not even known to the counterparties themselves but only to the exchange and potentially to the regulator, if there are reporting requirements in place.

complex form of automation to manage these positions. This suggests that the liquidity provider is more sophisticated.

For our second — direct — approach, we gather labelled wallet address tags from the blockchain explorer Etherscan and the blockchain-intelligence company Arkham Intelligence and flag wallet addresses of known sophisticated market participants accordingly.¹⁹ We source labels for all LP wallet addresses identified in our sample and find labels for 7,153 addresses (7.5% of all observed addresses). After removing labels related to the Ethereum Name Service, we are left with 826 unique labels.²⁰ We manually flag labels as ‘sophisticated’ for investment funds, venture capital firms, asset management protocols, arbitrage/MEV bots or professional liquidity providers (e.g. market makers or high-frequency trading firms). We identify a total of 240 labels as sophisticated, that are mapped to 839 different addresses.

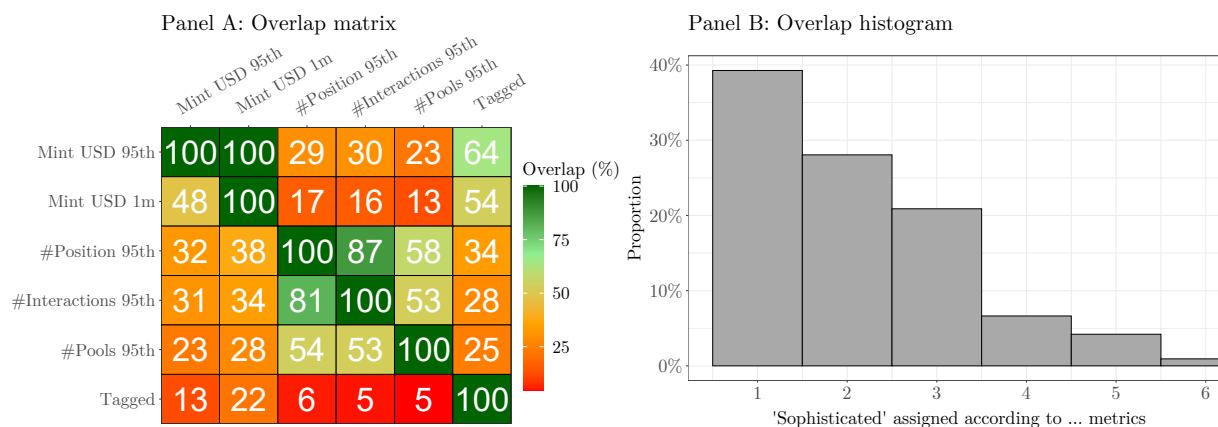
Figure 4 summarises the similarities obtained by using these different definitions. Panel A constructs an ‘overlap matrix’ that shows how the different criteria are related to each other. We find higher overlaps within the group of capital intensive LPs. For example, the cell in row 2 and column 1 highlights that almost half of wallet addresses (48%) that are part of the identification mechanism ‘Mint USD 95th’ (that used the distribution cutoff) are also identified as sophisticated through the identification mechanism ‘Mint USD 1m’ (that used an absolute value as cutoff). Similarly, a higher overlap is present within the group of technologically sophisticated agents. For example, more than half (54%) of wallet address that are identified as sophisticated through the number of unique liquidity positions are also detected by the mechanism focusing on the number of distinct traded pools.

We further construct an ‘overlap histogram’ (Figure 4, Panel B). Here, we consider all wallet addresses that have been labelled ‘sophisticated’ by any of our metrics. Then, we investigate how many metrics would have assigned the respective address to the sophisticated category. For example, just under 40% of sophisticated addresses would have been classified as such through only one metric (first grey bar).

¹⁹Appendix D provides an example on how Etherscan labels wallet addresses.

²⁰The Ethereum Name Service (ENS) is a decentralized system that allows anyone to register and manage human-readable domain names on the Ethereum blockchain, similar to assigning a username to a wallet address.

Figure 4: Classification overlap. This figure shows the overlap of wallet addresses classified as sophisticated across six proposed mechanisms. Panel A shows an overlap matrix in which Mint USD 95th, #Position 95th, #Interactions 95th and #Pools 95th are the groups that use a distribution threshold mechanism that classifies wallet addresses above the 95th percentile as sophisticated and retail otherwise. Mint USD is the maximum mint size of a liquidity position in USD, Mint USD 1m is a binary variable equal to 1 if a wallet address has minted a position greater than USD 1m, #Positions, #Interactions and #Pools use the absolute number of distinct liquidity positions, liquidity transaction or used liquidity pools as their underlying variable. Tagged is a binary variable equal to 1 if a wallet address has been identified as sophisticated through tagged wallet labels of Etherscan or Arkham Intelligence. Panel B shows the overlap histogram that illustrates the percentage of wallet addresses that are classified as sophisticated according to one, two, ..., six metrics.



Appendix B compares the six different sophisticated groups across key liquidity provision characteristics against retail. It becomes evident that the liquidity provision behaviour of sophisticated agents, no matter the specific classification mechanism, is similar and differs drastically from the behaviour of retail investors. For the subsequent analysis we classify wallet addresses as ‘sophisticated’ if they meet at least two identification criteria.²¹ By focusing on the *at least two* measure, we are attempting to balance a trade off between being inclusive in our definition: we do not want to leave out truly sophisticated LPs, while at the same time we do not want to risk diluting the differences by including wallets that are operated by retail investors. In turn, an address is classified as ‘retail’ if it satisfies at most one criterion. Using our sophisticated threshold of satisfying ‘at least two’ criteria, we classify 6,124 wallet addresses as sophisticated, representing about 7% of the total 88,299 addresses. When we directly identified wallets belonging to sophisticated agents using labelled addresses, we found that each agent was linked to 3.5 different wallets on average (240 agents and 839 wallets). If the same ratio applies using the ‘at least two’ criterion, and assuming that unsophisticated investors do not have multiple wallets, sophisticated would therefore represent between 2-3% of all market participants.

²¹At first glance the criterion ‘Mint USD 1m’ might appear redundant as all positions satisfying this threshold are also captured by the criterion ‘Mint USD 95th’. However, given our threshold of satisfying *at least two* criteria, the ‘Mint USD 1m’ criterion ensures that LPs with exceptionally large positions are not categorized as retail, even if they do not meet any other sophisticated criteria.

4 Empirical Analysis

We begin this Section with a discussion of the prevalence of the liquidity provision strategies described in Section 2.2 in our sample. Subsequently, we provide summary statistics of liquidity positions (Section 4.2). We then move on to the main empirical contribution of the paper: a detailed analysis of the behaviour of retail and sophisticated liquidity providers (Section 4.3) in terms of their prevalence (Section 4.3.1), profitability (Section 4.3.2) and their responses during times of market stress (Section 4.3.4).

4.1 Liquidity provision strategies in Uniswap V3

Section 2.2 described four liquidity provision strategies; concentrated, unconcentrated, range order and just-in-time. Using the data in our sample, we assess the prevalence of each of these strategies.

‘Concentrated’ positions can be characterized by the tickrange spread we introduced in Section 2.2.1. Figure 5 shows the histogram of tickrange spreads across liquidity positions of our sample, where the width of each bin equals 2%. Panel A shows the full sample, whereas Panel B and C illustrate the sophisticated and retail sub-samples respectively. We find several distinct patterns.²² First, a large proportion of positions post liquidity in tight ranges, e.g. more than 25% of positions have a tickrange spread of 2% or less. Positions with small tickrange spreads require constant active monitoring.²³ Liquidity positions that are classified as sophisticated show a much higher proportion of very narrow tickrange spreads (37% vs 11%), which is consistent with sophisticated LPs pursuing a more active liquidity management style.

Second, a notable proportion of liquidity is posted at a tickrange spread of 120%. This is due to the Uniswap graphical user interface providing a default tickrange spread for pools in the 30 basis points (bps) and 100bps fee tier at this level. Indeed, many more retail participants deploy the default 120% spread. In contrast, the default tick range is significantly less prevalent in sophisticated liquidity positions, as more sophisticated agents deploy their liquidity positions programmatically (directly interacting with the smart contract through code) and not through a user-interface.

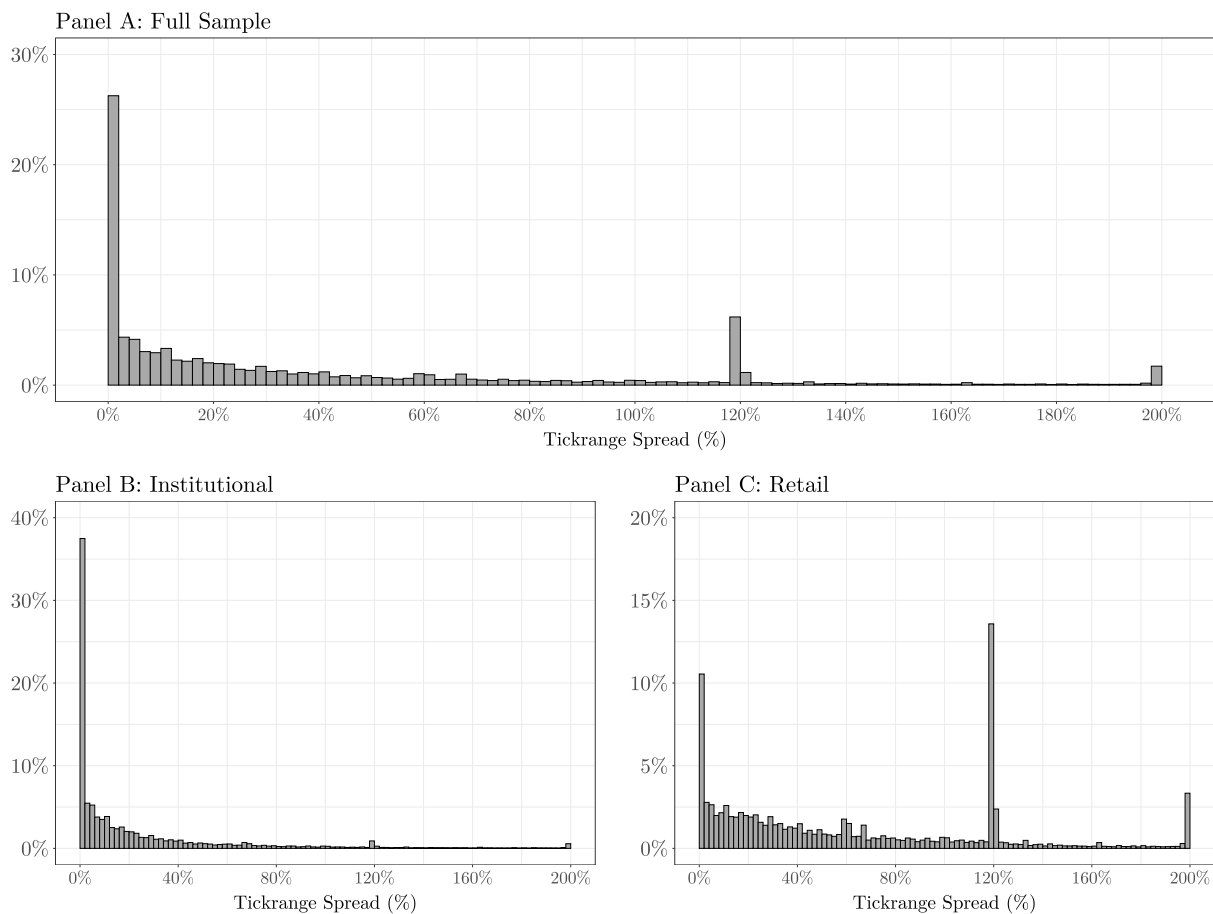
We also find a small but non-trivial proportion of positions mimicking the ‘unconcentrated’ liquidity provision strategy, similar to Uniswap V2. We observe a tickrange spread value of 200% for about 2% of positions. Given the large width of the position, this behaviour is in practice equivalent to providing liquidity at all price levels. We find that these positions are more heavily used by retail participants, which might have a

²²It should be noted that in a truly decentralised and *democratised* system, there should not be significant differences in such characteristics and in the behaviour of wallets among sophisticated and retail investors. The fact that we do observe such differences is already an indicator that intermediaries emerge even in a system that attempts to eliminate them. Lehar et al. (2023) highlight that one reason why large LPs emerge is the fixed cost of managing liquidity.

²³Such positions can also be employed for JIT and range orders strategies.

‘set and forget’ motivation to avoid undertaking any monitoring at all. These LPs minimise the monitoring costs of their positions as such a position will never go out of range, though this comes at the expense of capital efficiency.


Figure 5: Histogram Tickrange Spreads. This figure shows the distribution of the liquidity positions’ widths. Panel A provides an overview of the full sample, whereas Panels B and C focus on liquidity positions of sophisticated and retail LPs respectively.



We further query our dataset for ‘range orders’. To detect this variation of a CLOB limit order, we use the following heuristic to filter liquidity positions. First, we verify that liquidity provision is one-sided, i.e. only one of asset X or Y has been provided at mint. Second, the liquidity is posted out-of-range (above or below the prevailing market price). Third, at time of the position’s burn the liquidity has been fully converted into the other asset. Fourth, at the time of the burn, the position is out-of-range once again but on the other side of the prevailing market price (below if it was above and vice versa). To limit the false positive rate we also filter for a maximum time in the market (between mint and burn) of 48 hours. We find about 6,000 positions following this pattern, equaling slightly more than 1% of all positions, indicating that this limit-order-like type of liquidity provisioning is not particularly prevalent in practice.

Finally, we scan our sample for ‘JIT transactions’. Following the methodology outlined in Wan and Adams (2022) we look for positions that fulfil the following conditions. They have been minted and burned in the same block and the same pool. Within a block, the mint and burn transaction are two block positions apart and the transaction in between the mint and the burn is a swap in the same liquidity pool, i.e. the swap is sandwiched by mint and burn. We find less than 1% of trades result from JIT transactions, consistent with the results of Wan and Adams (2022) and Lehar et al. (2023). Figure 6 shows the number of JIT transactions per day in our sample.

Figure 6: JIT Liquidity Transactions. This figure shows a timeseries of the number of ‘just-in-time’ liquidity transactions per day. The sample period is from the launch of Uniswap V3 (May 2021) to December 31, 2023.



Our analysis also suggests that the activity of JIT liquidity provisioning is highly concentrated. We find almost 40% of identified JIT transactions to have occurred in the most active pool of our sample (ETH/USDC 5bps). This result is unsurprising for several reasons. First, JIT activity is more lucrative in pools with large trade sizes and more opportunities arise with larger, more frequent transactions. Second, JIT transactions often use borrowed funds, favoring highly liquid assets like ETH and USDC with readily available lending options. Third, volatile pools, such as stablecoin/token pairs, exhibit a more fragmented liquidity landscape. This fragmentation allows attackers to strategically inject large amounts of liquidity within the precise price range anticipated for the targeted swap. Furthermore, the analysis identified only a small number of actors responsible for JIT attacks. Notably, only ten wallet addresses account for approximately half of all observed JIT transactions. This behaviour has some commonality with the behaviour of sophisticated liquidity providers in traditional financial markets. There, liquidity providers are involved in speed races

(Budish et al., 2015, 2024) to exploit the continuous nature of trading rather than attempting to front run a transaction that they can predict will take place by looking at the blockchain’s mempool. Aquilina et al. (2021) show that almost a quarter of all traded volume takes place in such races in UK equity markets, a much higher percentage compared to our results. They also show that, similar to our findings in DEXs it is just a handful of firms that are involved in such behaviour.

Overall, our analysis of liquidity provision strategies shows that the overwhelming majority of liquidity is provided using the ‘concentrated’ strategy, i.e. LPs in DEXs mimic the behaviour of market makers in traditional finance by quoting prices close to the prevailing market price. While the range of prices quoted is wider than what typically happens in CLOB markets in traditional finance, most of the liquidity provision takes place using orders that mimic providing bid and ask quotes.

4.2 Descriptive statistics

Table 2, Panel A provides summary statistics for the full sample of 430,799 liquidity positions created by 88,299 distinct wallet addresses. Panel B divides the sample into the asset pair categories ‘stable/stable’ (if an asset pair consists of two stablecoins), ‘stable/token’ (if an asset pair consists of one stablecoin and one non-stable token) and ‘token/token’ (if an asset pair consists of two non-stable tokens). Histograms of selected parameters can be found in Appendix A. Of the analysed positions, just less than half (49%) are in token/token liquidity pools, 47% in stable/token pools and only 4% in stable/stable pairs. The average position size, as measured through its value at mint, is substantial - over 2 million USD. In contrast, the median size is significantly smaller, standing at about 26,000 USD. This disparity between the average and median suggests a pronounced positive skew in the data, which is similarly observable in other parameters such as earned fees and the duration of a position. The average (median) position accrued about 1,758 USD (129 USD) in fees and has a lifetime of 66 days (2.8 days). Again, the data exhibits a large variance with a position’s lifespan ranging from only a few blocks to multiple years. We also find different approaches with regards to the management of the provided liquidity, measured through the number of interactions with a position. The median of three interactions corresponds to a ‘mint’ (in which the position is created), a ‘burn’ (in which the position is removed) and a ‘collect’ (in which accrued fees are transferred to the beneficial owner). Numbers higher than three indicate that the LP has amended a position’s liquidity or prematurely collected accrued fees during its lifetime, i.e. engaging in more active liquidity management. We further find that LPs are adept at keeping their positions ‘active’ with an average position remaining in an active price range for 86% of its lifetime.

Table 2: Summary statistics: Liquidity positions. This table reports summary statistics of key characteristics of liquidity positions. Panel A shows the full sample. Panel B splits the sample across asset pair categories (stable/stable, stable/token, token/token). MintSize (\$) is the mint size of a position in US\$. Fees (\$) is the amount of accrued fees of a position in US\$, winsorized at the 99th percentile. Duration (days) is the time between mint and burn of a position in fractional days. Interactions is the number of interactions with a liquidity position. TickRange (%) is the tickrange spread of a positions in percentage. TimeActive (%) is the proportion of days a position has been active. Stable/Stable, Stable/Token and Token/Token indicate the respective pool classification.

Panel A: Overall sample								
Variable	N	Mean	Std. Dev.	Min	Q5	Median	Q95	Max
MintSize (\$)	430,799	2,167,861	7,187,902	1.00	37.36	25,899	18,449,526	102,499,235
Fees (\$)	430,799	1,758	5,597	0.00	0.00	129.0	8,408	41,847
Duration (days)	430,799	65.61	171.7	0.00	0.00	2.82	498.7	974.2
Interactions	430,799	4.09	16.60	1.00	1.00	3.00	8.00	5,237
TickRange (%)	430,799	39.50	49.38	0.01	0.10	17.65	133.3	200.0
TimeActive (%)	430,799	86.35	27.19	0.00	14.29	100.00	100.00	100.00
Stable/Stable (%)	430,799	3.81	19.15	0.00	0.00	0.00	0.00	100.00
Stable/Token (%)	430,799	46.76	49.90	0.00	0.00	0.00	100.00	100.00
Token/Token (%)	430,799	49.42	50.00	0.00	0.00	0.00	100.00	100.00

Panel B: Asset pair categories								
Variable	N	Mean	Std. Dev.	Min	Q10	Median	Q90	Max
Category: stable/stable								
MintSize (\$)	16,424	1,224,735	4,288,401	1.00	23.86	47,880	7,284,464	100,744,800
Fees (\$)	16,424	1,257	4,892	0.00	0.00	35.15	5,208	41,847
Duration (days)	16,424	96.91	216.2	0.00	0.00	5.99	717.0	974.2
Interactions	16,424	5.12	51.31	1.00	1.00	3.00	10.00	5,237
TickRange (%)	16,424	2.89	19.50	0.01	0.01	0.20	4.40	200.0
TimeActive (%)	16,424	92.71	20.12	0.00	43.52	100.00	100.00	100.00
Category: stable/token								
MintSize (\$)	201,453	3,739,467	9,536,970	1.00	10.60	38,867	27,645,439	102,499,235
Fees (\$)	201,453	1,762	5,850	0.00	0.00	101.1	8,516	41,847
Duration (days)	201,317	72.50	182.3	0.00	0.00	2.14	560.1	974.1
Interactions	201,317	3.90	8.42	1.00	1.00	3.00	7.00	2,306
TickRange (%)	201,453	35.48	46.84	0.01	0.10	13.78	120.4	200.0
TimeActive (%)	201,453	86.93	26.61	0.00	15.43	100.00	100.00	100.00
Category: token/token								
MintSize (\$)	212,922	753,659	3,557,512	1.00	152.0	19,517	2,243,118	81,477,158
Fees (\$)	212,922	1,793	5,398	0.00	0.02	173.1	8,517	41,847
Duration (days)	212,680	56.68	156.2	0.00	0.00	3.10	373.6	974.2
Interactions	212,680	4.19	16.94	1.00	2.00	3.00	9.00	5,080
TickRange (%)	212,922	46.12	51.63	0.01	0.60	24.08	151.5	200.0
TimeActive (%)	212,922	85.31	28.12	0.00	12.66	100.00	100.00	100.00

When comparing liquidity positions across asset pair categories, notable variations emerge, especially between positions in stable/stable pools and stable/token pools or token/token pools. We find the median position size in stable/stable pairs to be 30% larger compared to stable/token pairs and over two and a half times larger than pools in token/token pairs. Positions in stable/stable pools also tend to have a longer lifespan, with the

average time between mint and burn being 97 days, or roughly 30% longer than positions in stable/token (73 days) and 70% longer than token/token (57 days) pools. Additionally, liquidity positions in stable/stable pools remain active for a greater proportion of their lifetime (93%) compared to positions in stable/token (87%) or token/token pairs (54%). We also find positions in stable/stable pairs to be considerably narrower than in non stable/stable pools. Specifically, we find the average position’s tickrange spread to be 2.9% in stable/stable pairs, significantly tighter than in stable/token (35%) and token/token pairs (46%). The observed differences are unsurprising, considering that stable/stable pools exhibit lower volatility - with both tokens typically remaining very close to their ‘peg’ of 1 USD. Therefore, 1) LPs can leverage the feature of ‘concentrated liquidity’ more aggressively by posting liquidity in narrower ranges, 2) Positions remain in an active state for a greater proportion of time, and 3) Positions do not need to be adjusted as frequently, reflected in their longer life span.

4.3 Liquidity provision by retail and sophisticated investors

This section compares the behaviour of retail and sophisticated investors when they provide liquidity. Recall that, if liquidity provision was fully decentralised we would not expect significant differences between retail and sophisticated LPs. However, as our definition of sophisticated investors is based on some of their characteristics it is indeed possible that such differences emerge. A regression analysis allows us to quantify the relative effect of being a retail agent. In a first step, we use a simple model, which takes the following form:

$$Y_i = \alpha + \beta_1 Retail_i + \varepsilon_i \tag{1}$$

Where the constant can be interpreted as the value of the variable for sophisticated investors and β_1 as the differential impact of being retail. Y_i is a metric characterizing the liquidity provisioning at the LP or position level i . At the LP level, $\#Positions$ is the number of liquidity positions a wallet address has created and $\#Pools$ the number of distinct pools a wallet has provided liquidity in. At the position level, $Size\ at\ Mint$ is the USD size of the position at time of the mint, $Duration$ is the time in days between mint and burn (how long the liquidity remains posted), $Proportion\ Active$ is the proportion of time during which a liquidity position is in an active price range, $Tickrange\ Spread$ is the width of the position’s liquidity range in percentage and $\#Interactions$ is the number of interactions undertaken with a specific liquidity position.

Table 3 summarises the results of these regressions. They clearly expose large differences between sophisticated and retail LPs, providing the first piece of evidence for the emergence of distinct subgroups of market participants. At the LP level, retail LPs create considerably fewer positions than retail ones. Specifically,

the results suggest that while a sophisticated LP creates 41 positions, a retail LP creates just over two. Retail LPs are also active across a narrower range of pools. Whereas the average retail LP provides liquidity in 1.4 distinct pools, sophisticated LPs do so in over five pools. We further find significant differences at the position level. Retail LPs have considerably smaller position sizes: 29,000 USD compared to 3.7 million USD respectively (a two orders of magnitude difference). Importantly, the difference between sophisticated and retail liquidity providers is also evident in areas that are not directly related to our definition of a sophisticated LP. For example, sophisticated LPs provide liquidity for a shorter amount of time. Compared to the average sophisticated position that has a duration of about two weeks (16 days), liquidity positions for retail LPs exist on average four months (120 days) longer. As Figure 5 already indicated, our analysis also confirms that the liquidity positions of sophisticated LPs are posted in considerably tighter ranges, at tickrange spreads that are less than half the size of retail widths (23% vs. 63% for retail). Importantly, despite the narrower ranges, sophisticated LPs are able to keep their positions in an active state for a higher proportion of time, about 9 percentage points more frequently than retail positions. Sophisticated LPs also interact with their positions more often, highlighting their more active management style. To summarise, sophisticated market participants provide liquidity across more positions, for shorter time periods, in narrower ranges, are active across a larger number of pools and engage in a more active liquidity management style. Appendix B shows that these results do not depend on our preferred definition of a sophisticated LP meeting at least two of our criteria: The stricter the definition of an LP, meaning the more criteria that are met, the greater the differences compared to a retail investor - with near monotonic increases in the observed differences as the number of sophisticated traits increases.

Table 3: The differences between retail and sophisticated LPs. This table reports results of an OLS regression analysing differences between liquidity positions of retail and sophisticated market participants. At the LP level, the dependent variables are the number of distinct liquidity positions an address has minted and the number of distinct pools an address has provided liquidity in. At the position level, the dependent variables are the Size of the position at mint in US\$, the duration how long the position has been posted in days, the tickrange spread measuring the width of the position in percentage, the proportion of time the position has been active in percentage and the number of interactions with the liquidity position, Retail is a binary variable identifying if the liquidity position owner is considered retail (fulfilling at most one criteria). Standard errors reported in parentheses are robust and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	LP level		Position level				
	#Positions	#Pools	Size at Mint (\$)	Duration (days)	Tickrange Spread (%)	Proportion Active (%)	#Interactions
Intercept	41.04*** (2.16)	5.55*** (0.079)	3,696,713*** (18,174)	15.52*** (0.13)	22.77*** (0.073)	90.27*** (0.047)	4.58*** (0.043)
Retail	-38.85*** (2.16)	-4.10*** (0.080)	-3,667,867*** (18,175)	120.2*** (0.58)	40.13*** (0.15)	-9.42*** (0.087)	-1.18*** (0.044)
N	88,299	88,299	430,799	430,410	430,799	430,799	430,410
Adj. R2	0.05	0.24	0.06	0.12	0.16	0.03	0.00

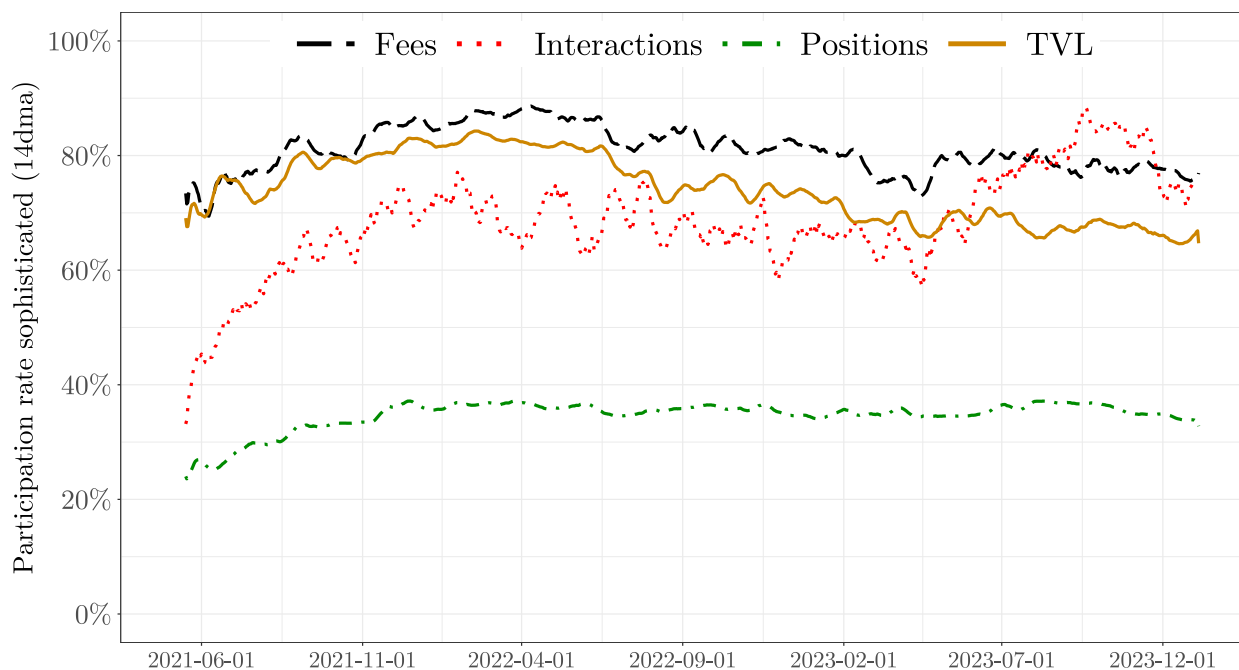
4.3.1 The prevalence of sophisticated LPs

The previous section highlighted the different characteristics in the behaviour of retail and sophisticated liquidity providers. This section analyses the extent to which retail LPs are able to compete with sophisticated ones or whether the forces that have driven to the emergence of different kinds of liquidity providers in traditional finance are resulting in a similar outcome in DeFi.

As a first step, we measure the *participation rate*, or dominance of sophisticated LPs at the protocol level by calculating the share of fees or TVL accruing to sophisticated LPs.

Figure 7 shows the participation rate of sophisticated investors. We find that roughly 80% of TVL and accrued fees can be attributed to them, despite them holding only around 20-30% of positions and representing only 7% of LPs. Furthermore, while at the inception of Uniswap V3 sophisticated LPs were responsible for 40-50% of interactions with liquidity positions, this number has steadily grown to 70-80% at the end of 2023. We conclude that the participation rate of sophisticated market participants is both substantial and increasing.

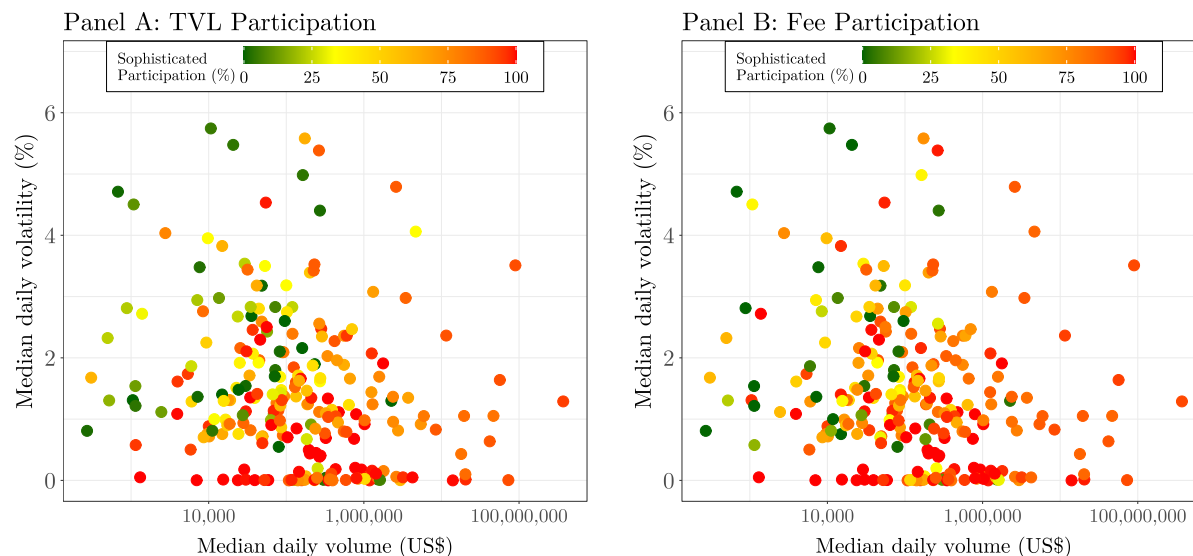
Figure 7: Participation of sophisticated LPs at the protocol level. This figure shows the participation rate of sophisticated market participants across multiple dimensions. Daily pool values are aggregated using a 14 days moving average. Fees is the USD of accrued fees, Interactions is the number of liquidity related transactions, positions is the proportion of fee-collecting liquidity positions and TVL is the total-value locked in USD.



Next, we investigate the preferences of sophisticated LPs with regards to their choice of liquidity pools. For each liquidity pool, we measure the median daily trading volume, volatility and the participation rate of sophisticated LPs for TVL and fees during our sample period.

Figure 8 visualizes the relationship between a pool’s trading volume, volatility and sophisticated dominance with regards to TVL (Panel A) and accrued fees (Panel B). A number of interesting patterns emerge: first sophisticated participation is higher in pools with higher volumes. This dominance is striking: where daily trading volume exceeds 10m USD, sophisticated LPs provide essentially all the liquidity and earn most of the fees. Retail liquidity providers are much more prevalent in pools where daily trading volume is smaller than 100,000 USD. In other words, where there is significant money at stake, sophisticated investors provide the lion’s share of liquidity and reap most of the profits. Second, sophisticated LPs tend to focus on the less volatile pools: red dots (indicating sophisticated LP dominance) are relatively more prominent in the bottom half of both panels of Figure 8 - representing less volatile asset pairs.

Figure 8: Sophisticated participation at the liquidity pool level. This figure illustrates the dominance of sophisticated market participants in terms of total-value locked (Panel A) and accrued fees (Panel B). Each dot represents one of the 250 liquidity pools and relates the median daily volatility expressed as a percentage with the median USD daily trading volume of a pool. The x-axis is log-scaled. The color represents the sophisticated participation rate, with green indicating low dominance of sophisticated liquidity providers in a liquidity pool and red indicating high sophisticated dominance.



To validate this visual evidence more rigorously, we conduct a logistic regression as specified in Equation (2), which we run separately for TVL and fee dominance. The response variable $Y_{i,t}$ corresponds to whether a liquidity pool i on day t is dominated by sophisticated market participants. It takes the value of 1 if more than 50% of the respective measure (TVL or accrued fees) is attributed to sophisticated LPs, and is 0 otherwise.

$$Y_{i,t} = \alpha + \beta_1 \log(\$TVL_{i,t}) + \beta_2 \log(\$Volume_{i,t}) + \beta_3 \log(Volatility_{i,t}) + \varepsilon_{i,t} \quad (2)$$

The results, shown in Table 4, confirm that the influence of sophisticated LPs increases with the size of a liquidity pool (both in terms of TVL and volume) and decreases with volatility, for both measures of TVL and fee dominance. Specifically, we find that each e-fold (i.e. approximately 2.7 times) increase in TVL increases the odds of TVL being dominated by sophisticated agents by 60% ($\exp(0.62)=1.6$). Similarly, we find an odds-increase of 28% ($\exp(0.25)=1.28$) for an e-fold increase in volume. In contrast, the likelihood of sophisticated dominance diminishes with increased volatility. Here, an e-fold increase in volatility reflects a 16% lower likelihood ($\exp(-0.18)=0.84$). The results of fee dominance are consistent. Notably, an e-fold increase in trading volume raises the odds of sophisticated LPs dominating the pool by 93% ($\exp(0.66)=1.93$), which aligns with sophisticated LPs' preference for high-revenue pools. Overall, sophisticated market participants tend to dominate liquidity pools that are economically more significant (measured through TVL and volume) and exhibit lower volatility.

Table 4: Logistic regression results on the dominance of sophisticated agents in liquidity pools. This table reports results of a logistic regression analysing the likelihood of liquidity pools being dominated by sophisticated LPs. The dependent variables *TVL dominance* and *Fee dominance* are binary variables taking the value of 1 if more than 50% of TVL or generated fees are attributed to sophisticated LPs and 0 otherwise. TVL and Volume are the respective pool, day TVL and trading volume in US\$. Volatility is the absolute value of the daily high-low range scaled by its midpoint. Standard errors are reported in parentheses are clustered at the day level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	TVL dominance				Fee dominance			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(\$TVL)	0.62*** (0.007)			0.47*** (0.007)	0.74*** (0.008)			0.31*** (0.010)
log(\$Volume)		0.25*** (0.002)		0.14*** (0.003)		0.66*** (0.004)		0.56*** (0.005)
log(Volatility)			-0.18*** (0.004)	-0.12*** (0.004)			-0.022*** (0.005)	-0.009** (0.004)
N	165,872	164,826	159,378	159,202	165,944	164,806	159,259	159,083
FE: Date	X	X	X	X	X	X	X	X
Adj. R2	0.13	0.08	0.02	0.15	0.17	0.38	0.02	0.32

4.3.2 Analysing the profitability of liquidity provision

LPs engage in liquidity provision with the expectation of earning a profit on their provided capital. Analysing the magnitude and source of these profits is thus important to understand the motivations of LPs. We decompose an LP’s total return into three components. Similar to Heimbach et al. (2022), we derive the profit earned through fees and an adverse selection cost component. Additionally, we source the transaction costs (gas fees) that LPs pay to manage their liquidity positions from the Ethereum blockchain.

Profits earned through fees depend on three parameters: the overall trading volume, the amount and distribution of liquidity in a pool and the transaction costs (gas fees) paid. Each LP will earn revenues depending on its strategic liquidity provision and the amount of trading that takes place in the liquidity pools where its positions are active. We derive the fee revenue per position per day and use both an absolute measure, the accrued fee revenue F , and a relative yield metric, which is the absolute fee revenue over the value of the provided capital, calculated as $FeeYield = \frac{F}{V_{hold}}$.

The second component, an adverse selection cost, is often referred to as *Impermanent Loss* or *Divergence Loss*. Impermanent loss arises because the LP’s asset mix changes unfavorably as the trading process evolves. As the AMM cannot update its quotes without trading, by design the LP is ‘selling low’ and ‘buying high’ because the AMM provides ‘stale’ quotes that arbitrageurs trade against. The LP’s impermanent loss is calculated as the difference between the actual value of the LP’s assets in the pool and the hypothetical value of the same assets, had the LP simply held assets outside the liquidity pool.

To calculate impermanent loss we calculate two metrics. The value of an LP’s assets after the provision of liquidity (V_{liq}), and the hypothetical value of the position had the LP simply held the two assets in a portfolio (V_{hold}). The former can be defined as $V_{liq} = P_{x,y} * x_1 + y_1$ where $P_{x,y}$ is the price of the pair in terms of token y ; and x_1 and y_1 are the amounts of token x and y held by the LP after liquidity provision. The latter is given by $V_{hold} = P_{x,y} * x_0 + y_0$ where x_0 and y_0 are the quantities the LP held before providing liquidity. The Impermanent Loss is then derived as $IL = \frac{V_{liq} - V_{hold}}{V_{hold}}$. Notably, impermanent loss is a negative return in the interval $[-1, 0]$ where lower values indicate a higher loss.

The third component is the transaction costs incurred by LPs that are paid to manage the liquidity position. We calculate the median amount of *gasUsed* across the transaction type (mint, collect or burn) and multiply it by the *gasPrice* paid, which we source from the transaction receipt. We then convert the ETH-denominated metric to USD using hourly ETH prices from Binance.²⁴ More details on the computation of transaction costs are provided in Appendix G.

²⁴We use the median gas used because sophisticated participants use liquidity provision as part of more complex strategies. Attributing the total amount of gas used to their liquidity provision activity would overestimate their costs.

In the following analysis we use the term ‘Total Return’ for the sum of fee yield and impermanent loss. The ‘Net Return’ is the ‘Total Return’ reduced by gas fees as in Equation (3).

$$\begin{aligned}
 R_{net} &= \overbrace{FeeYield + ImpermanentLoss}^{\text{Total Return}} - GasFees \\
 &= \frac{F}{V_{hold}} + \frac{V_{liq} - V_{hold}}{V_{hold}} - \frac{G}{V_{hold}} = \frac{F + V_{liq} - V_{hold} - G}{V_{hold}}
 \end{aligned} \tag{3}$$

We assess Fee Yield and Impermanent Loss on a daily basis, i.e. an LP can decide whether or not to provide liquidity at the end of each day. Gas fees are assigned to the day they are incurred.²⁵ Summary statistics on the profitability of liquidity positions are reported in Appendix E. They highlight that retail positions do make positive daily excess returns on average, but also that this is driven by a smaller number of highly profitable positions. However, the negative median indicates that the positions of retail LPs lose money for the majority of days. In contrast, positions of sophisticated investors do considerably better: median daily excess returns are essentially zero and more profitable on average compared to retail.

Our models to examine the profitability differences between sophisticated and retail LPs take the form:

$$Y_{i,t} = \alpha + \beta_1 Retail_{i,t} + \beta_2 \log(Volatility_{i,t}) + \beta_3 \log(\$TVL_{i,t}) + \beta_4 \log(\$Volume_{i,t}) + \varepsilon_{i,t} \tag{4}$$

Where $Y_{i,t}$ is a profitability metric at the liquidity pool, day or liquidity position, day level. At the pool level, we use the total amount of fees in USD terms earned by the group of sophisticated or retail LPs on a given day t in the liquidity pool i . For the position-level regressions, we use the absolute fees, fee yield, impermanent loss and return metrics as per Equation (3) on a given day t for position i . The binary variable *Retail* takes a value of 1 if the observation belongs to the group of retail LPs and 0 for sophisticated LPs. We control for *Volatility* that is the pool’s high-low range scaled by its midpoint on a given day, as well as a pool’s *TVL* and *Volume* in USD.

The results shown in Table 5 indicate that the profitability of sophisticated and retail market participants differs significantly. Despite retail representing 93% of all liquidity providers, they earn about 6,000 USD less in each pool, every day, on average.

²⁵To increase tractability, we exclude observations of liquidity positions with a value below USD 1 and winsorize our metrics at the 99th percentile.

Table 5: The profitability of sophisticated and retail LPs. This table reports results of a fixed effect OLS regression analysing profitability measures across liquidity pools and liquidity positions. The dependent variable Fees (\$) is the total amount of fees the group of sophisticated or retail has accrued on a pool, day level or the position has accrued on a day. The dependent variables Fee Yield (bps), Impermanent Loss (bps), Total Returns (bps) and Net Return (bps) are as outlined before. Retail is a binary variable identifying if the metric belongs to the group of retails or the liquidity position owner is considered retail, where retail fulfills at most one classification criterium. Volatility is the absolute value of the daily high-low range scaled by its midpoint. TVL and Volume are the respective pool, day TVL and trading volume in USD. Fee Yield, Impermanent Loss and Total Return are winsorized at the 99th percentile. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	Pool, Day level		Position, Day level			
	Fees (\$)	Fees (\$)	Fee Yield (bps)	Impermanent Loss (bps)	Total Return (bps)	Net Return (bps)
Retail	-6,014*** (1,695)	-262.7*** (30.26)	-3.54*** (0.34)	0.18*** (0.046)	-3.36*** (0.31)	-2.84*** (0.27)
log(Volatility)	181.6 (134.8)	-4.96 (4.64)	0.100 (0.43)	-0.39*** (0.11)	-0.29 (0.36)	-0.25 (0.33)
log(\$TVL)	-204.7 (189.3)	-28.65** (11.43)	-3.82*** (0.41)	0.24*** (0.076)	-3.58*** (0.36)	-3.33*** (0.32)
log(\$Volume)	1,615*** (189.7)	40.06*** (5.16)	3.78*** (0.34)	-0.34*** (0.036)	3.44*** (0.31)	3.15*** (0.28)
N	316,690	28,317,758	28,317,758	28,317,758	28,317,758	28,317,758
FE: Pool	X	X	X	X	X	X
FE: Date	X	X	X	X	X	X
Adj. R2	0.24	0.01	0.27	0.16	0.24	0.22

Retail market participants are also less profitable at the position level. A retail liquidity position earns 263 USD less on average, per day, than a position belonging to a sophisticated agent. Importantly, this is not just driven by the fact that retail liquidity providers have smaller positions to begin with. Our results show that retail positions generate a daily fee yield of almost 3.5bps, which translates into a 14 percentage points lower annual relative fee revenue compared to sophisticated LPs. Interestingly, retail LPs experience lower adverse selection compared to retail ones: they lose approximately 0.2bps less through impermanent loss. Part of the reason why the impermanent loss is lower for retail positions can be reconciled with wider average position widths, that make positions less prone to impermanent loss. However, the lower impermanent loss does not make retail LPs more profitable than sophisticated ones, as the lower fee revenue outweighs the decreased adverse selection.

Overall, we find the daily total return to be 3.4bps lower on average for retail LPs. Sophisticated LPs seem to be rationally choosing to incur higher impermanent losses because they are more than compensated by

the trading fees they receive to provide liquidity. Including transaction costs, the underperformance of retail, measured through the net return, shrinks to about 2.9bps, indicating that sophisticated LPs spend a higher amount on gas fees, which is unsurprising given their more active liquidity management style. Due to the greater dilution of gas fees relative to position size, the profits of sophisticated LPs are less sensitive to higher gas prices. However, our robustness analysis in Appendix H demonstrates that the underperformance of retail LPs remains consistent across both high and low gas fee environments, indicating that gas fee levels do not significantly contribute to the observed profitability gap between retail and sophisticated market participants.

4.3.3 Excess returns and dynamic liquidity provision behaviour

Section 4.3.2 has shown that liquidity positions of retail LPs are less profitable than those of sophisticated market participants on both an absolute and relative level. In this section, we analyse the drivers of those results. We also look at returns on a risk-adjusted basis by benchmarking an LP’s daily return against the prevailing risk-free rate. We further focus on the LPs dynamic liquidity provision behaviour and analyse what role an LP’s tickrange spread and active liquidity management plays.

We derive an *Excess Return* metric for position i on day t that is the differential between the return derived from liquidity provisioning (after transaction costs) and the respective daily risk-free yield: $ExcessReturn_{i,t} = NetReturn_{i,t} - RiskFreeRate_{i,t}$.²⁶ We provide summary statistics on excess returns in Appendix E. Notably, we find that the median daily excess return is negative, indicating that on the majority of days LPs lose money on a risk-adjusted basis. However, given the positive skew, mean daily excess returns are positive, or about 3.5bps per day. We also observe a significant disparity in profitability between the two groups of liquidity providers: sophisticated LPs achieve substantially higher mean daily excess returns of 8.4bps, compared to just 2.7bps for retail LPs.

To investigate the drivers behind more or less successful liquidity positions, we quantify the absolute differences in excess returns based on investor classification (retail/sophisticated) as well as liquidity provision behaviour. Our models take the following form:

$$Y_{i,t} = \alpha + \beta_1 Retail_{i,t} + \beta_2 Interacted_{i,t} + \beta_3 Tickrange_i + \sum_{j=1}^n \gamma_j Controls_{j,it} + \varepsilon_{i,t} \quad (5)$$

Where $Y_{i,t}$ is the *Excess Return* in bps. *Retail* is a binary variable taking the value of 1 if the liquidity

²⁶Specifically, we use daily market rates of four-week treasury bills from the Federal Reserve Economic Data website under the ticker DTB4K sourced from <https://fred.stlouisfed.org/series/DTB4WK>.

position belongs to an address classified as retail. *Interacted* is a binary variable taking the value of 1 if the liquidity position has been interacted with over the past three days. We use this metric as a proxy for the active liquidity management of a position.²⁷ *Tickrange* is the width of the position's tickrange as previously defined. We further control for volatility, volume and TVL at the pool level.

The results shown in Table 6 suggest that liquidity positions of retail LPs generally underperform those of retail by about 3bps on a risk-adjusted basis (Column 1) after controlling for pool-specific characteristics. We also identify active liquidity management as proxied through the *Interacted* variable as a driver for higher excess returns. Specifically, we find that liquidity positions that have recently been interacted with outperform others on average by almost 5bps, all else equal (Column 2). Liquidity positions that leverage the feature of concentrated liquidity more aggressively, i.e. having a narrower (or lower) *Tickrange* coefficient equally see improved returns (Column 3). This result is unsurprising given that more concentrated positions essentially leverage their provided capital. However, more leveraged positions also run a higher risk of being adversely selected and going out of range. Consequently, it is even more important to actively supervise or manage these positions. We find supportive evidence for additional positive effects of the active management of narrow positions through the interaction term used in Column 4.

Importantly, the implications of our overall results equally apply to liquidity positions of both sophisticated and retail (Columns 5-8). While the previous discussion has emphasized the underperformance of retail positions, we find that retail market participants can sustain higher excess returns by mimicking the liquidity provision behaviour of sophisticated LPs, i.e. providing liquidity in narrower ranges and utilizing a more active management style (Columns 7 and 8).

²⁷Three days is also the median lifetime of a position.

Table 6: The profitability of sophisticated and retail LPs (extended analysis). This table reports results of an OLS regressions analysing drivers of risk-adjusted returns. The dependent variable Excess return is the differential of the daily net return and risk-free rate in bps. Retail is a binary variable identifying if the position's owner is classified as retail, where retail fulfills at most one classification criteria. Interacted is a binary variable taking the value of 1 if the position has been interacted with in the past three days. Tickrange is the tickrange spread of a position. Volatility is the absolute value of the daily high-low range scaled by its midpoint. TVL and Volume are the daily TVL and trading volume in the respective pool in USD. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	Excess return (bps)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Overall	Overall	Overall	Overall	Sophisticated	Sophisticated	Retail	Retail
Retail	-2.88*** (0.27)	-2.20*** (0.23)	-2.65*** (0.27)	-1.88*** (0.23)				
Interacted		4.77*** (0.86)		7.63*** (0.96)	5.67*** (0.60)	7.24*** (0.91)	1.80* (1.12)	4.56*** (1.05)
Tickrange			-1.29*** (0.25)	-0.91*** (0.23)	-4.20*** (0.33)	-3.68*** (0.43)	-0.68*** (0.16)	-0.57*** (0.16)
Interacted x Tickrange				-5.75*** (1.75)		-3.84*** (1.28)		-4.18** (1.90)
log(Volatility)	-0.43 (0.34)	-0.37 (0.33)	-0.41 (0.33)	-0.37 (0.33)	-0.31 (0.46)	-0.34 (0.46)	-0.33 (0.33)	-0.34 (0.33)
log(\$TVL)	-3.50*** (0.35)	-3.32*** (0.34)	-3.47*** (0.34)	-3.32*** (0.34)	-5.26*** (0.33)	-5.28*** (0.33)	-2.92*** (0.38)	-2.92*** (0.38)
log(\$Volume)	3.24*** (0.23)	3.09*** (0.21)	3.20*** (0.23)	3.09*** (0.21)	4.59*** (0.22)	4.62*** (0.21)	2.80*** (0.21)	2.82*** (0.21)
N	28,317,758	28,317,758	28,317,758	28,317,758	4,014,096	4,014,096	24,303,662	24,303,662
FE: Pool	X	X	X	X	X	X	X	X
Adj. R2	0.21	0.21	0.21	0.22	0.27	0.27	0.18	0.18

4.3.4 The effects of market volatility on LP profitability

In traditional financial markets, volatility is negatively related to the provision of liquidity. Pagano (1989) highlights that high volatility pushes risk-averse investors out of the market thereby reducing the likelihood of trades taking place; a series of inventory models that go back to Demsetz (1968) shows that volatility increases the liquidation costs of market makers, thereby reducing their willingness to intermediate trades. More recent empirical studies show that market makers are still able to profit during periods of high volatility (Anand & Venkataraman, 2016; Brogaard et al., 2018).

We investigate this issue in our setting, to assess whether LPs in AMMs behave similarly, and if sophisticated LPs differ from retail in their response to market volatility. Our model is shown in Equation (6) below:

$$Y_{i,t} = \alpha + \beta_1 Retail_{i,t} + \beta_2 HighVolatility_{i,t} + \beta_3 HighVolatility_{i,t} \times Retail_{i,t} + \sum_{j=1}^n \gamma_j Controls_{j,it} + \varepsilon_{i,t} \quad (6)$$

Where the variables $Y_{i,t}$, *Retail* and controls are as defined in Equation (4). We further define a binary dummy variable *HighVolatility* that takes the value of 1 if the volatility measured at the pool-day level is above or equal to the 95th percentile of its distribution and 0 otherwise. The results presented in Table 7 show that the profitability gap between sophisticated and retail widens on days with high volatility.

At the pool level, the earnings of retail positions reduce by an additional 14,200 USD on average on days with high volatility. Similarly, a single retail liquidity position earns 835 USD less in absolute fees, more than tripling the differential on days without excess volatility (223 USD). Importantly, the difference also widens on the calculated relative return metrics. Specifically, the fee yield difference increases by a factor of 2.5 compared to days without high volatility and reaches 8.5bps. We also find evidence that the increased volatility has a significantly more negative effect on impermanent loss for retail LPs. All else equal, we find total daily returns of retail liquidity positions underperform an additional 6.4bps on days of high volatility (4.7bps after accounting for transaction costs).

We conclude that sophisticated liquidity positions on average generate a 2.5x higher profit (7.6bps increase) on days with high volatility compared to the average daily net return of a retail position (approx. 3bps), further increasing their outperformance during times of heightened price fluctuations.

Table 7: The profitability of LPs during times of high volatility. This table reports results of a fixed effect OLS regression analysing profitability measures across liquidity pools and liquidity positions with a focus on volatility. The dependent variable Fees (\$) is the total amount of fees the group of sophisticated or retail has accrued on a pool, day level or the position has accrued on a day. The dependent variables Fee Yield (bps), Impermanent Loss (bps), Total Returns (bps) and Net Return (bps) are as outlined before. Retail is a binary variable identifying if the metric belongs to the group of retails or the liquidity position owner is considered retail (fulfilling at most one criterium). High Volatility is a binary variable taking the value of 1 if the pool, day observation if the observation's high-low range scaled by its midpoint is in the 95th percentile or above. TVL and Volume are the respective pool, day TVL and trading volume in USD. Fee Yield, Impermanent Loss, Total Return and Net Return are winsorized at the 99th percentile. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	Pool, Day level		Position, Day level			
	Fees (\$)	Fees (\$)	Fee Yield (bps)	Impermanent Loss (bps)	Total Return (bps)	Net Return (bps)
Retail x HighVolatility	-14,212*** (4,047)	-835.1*** (106.9)	-5.20*** (1.21)	-1.17** (0.47)	-6.38*** (0.99)	-4.74*** (0.90)
Retail	-5,080*** (1,483)	-223.5*** (27.07)	-3.30*** (0.30)	0.23*** (0.042)	-3.06*** (0.28)	-2.62*** (0.25)
HighVolatility	16,712*** (4,128)	810.2*** (100.8)	15.06*** (1.27)	-1.90*** (0.27)	13.16*** (1.20)	10.97*** (1.01)
log(\$TVL)	373.3* (205.3)	-13.24 (10.71)	-3.02*** (0.45)	0.34*** (0.066)	-2.67*** (0.39)	-2.51*** (0.36)
log(\$Volume)	987.1*** (117.0)	29.73*** (4.01)	3.20*** (0.29)	-0.39*** (0.037)	2.80*** (0.27)	2.58*** (0.24)
N	328,616	28,467,238	28,467,238	28,467,238	28,467,238	28,467,238
FE: Pool	X	X	X	X	X	X
FE: Date	X	X	X	X	X	X
Adj. R2	0.24	0.02	0.27	0.17	0.25	0.22

Overall, as in traditional finance, sophisticated LPs demonstrate the ability to remain in the market and exploit the profitable opportunities that increased volatility presents to them, while retail LPs struggle to do so.

To further investigate the source of this difference, we regress the number of interactions (as a measure of activity in the provision of liquidity) and the *tickrange spread* (as a measure of the posted spread) on the retail and volatility dummies and their interaction terms. Our focus is on the top 25 liquidity pools, as those are where most of the trading takes place, especially during high volatility periods.

The results are summarised in Table 8. They show that retail LPs are substantially less active on high volatility days: on average they have 17 less interactions with their positions. At the same time they also provide liquidity in narrower price ranges. On the other hand, sophisticated LPs *widen the spread* (compared to retail) in a similar manner to what their counterparts in traditional finance do to limit the negative impact of adverse selection, resulting in an increased adverse selection for retail LPs.

While the results demonstrate that sophisticated liquidity providers earn higher fees during periods of extreme market volatility, Section 4.3.1 revealed that they are not particularly active in asset pairs with consistently higher volatility. This suggests that sophisticated LPs are able to capitalize on short-term price fluctuations in otherwise less volatile and more liquid pools, rather than maintaining a constant presence in inherently volatile asset pairs, which typically carry significant adverse selection costs that outweigh the available fee revenue.

Table 8: The differences of liquidity management during times of high volatility. This table reports results of a fixed effect OLS regression analysing characteristics of liquidity management between retail and sophisticated market participants. The dependent variable #Interactions is the number of interactions with liquidity positions, i.e. mints or burns of the group of sophisticated or retail. The dependent variable Interactions (%) is the proportion of interactions with liquidity positions of the group of sophisticated or retail. The dependent variable Tickrange Spread is the daily mean tickrange spread of positions of the group of sophisticated or retail. Retail is a binary variable identifying if the metric belongs to the group of retail or the liquidity position owner is considered retail (fulfilling at most one criteria). High Volatility is a binary variable taking the value of 1 if the pool, day observation if the observation's high-low range scaled by its midpoint is in the 95th percentile or above. TVL and Volume are the respective pool, day TVL and trading volume in US\$. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	Top 25 pools			Full sample		
	#Interactions	Interactions (%)	Tickrange Spread (%)	#Interactions	Interactions (%)	Tickrange Spread (%)
Retail x HighVolatility	-17.37*** (5.51)	-9.20** (4.04)	-5.10*** (1.67)	-2.37** (1.02)	-5.04*** (1.67)	-1.65* (0.85)
Retail	-16.06* (9.11)	-22.63*** (6.05)	16.24*** (3.21)	-3.18*** (1.19)	-17.83*** (2.36)	12.83*** (1.82)
HighVolatility	22.37*** (7.34)	4.36** (2.00)	2.66* (1.41)	7.43*** (1.33)	2.86*** (0.82)	3.03*** (0.75)
log(\$TVL)	3.65 (2.74)	-0.68** (0.29)	0.41 (0.99)	0.93** (0.36)	-0.84*** (0.12)	-0.009 (0.72)
log(\$Volume)	2.12 (1.33)	-0.22* (0.11)	-1.00** (0.43)	1.10*** (0.14)	-0.10*** (0.037)	-0.91*** (0.21)
N	40,209	36,465	40,209	328,616	170,827	328,616
FE: Pool	X	X	X	X	X	X
FE: Date	X	X	X	X	X	X
Adj. R2	0.41	0.12	0.84	0.36	0.05	0.76

5 Conclusion

The development of decentralised finance had the stated intent to dis-intermediate the functions that, in traditional financial services, are fulfilled by intermediaries such as broker-dealers, custodian banks and market makers. In this study, we focus on a specific subset of DeFi applications: decentralised exchanges.

We parse logs of the Ethereum blockchain and reconstruct the state of the top 250 liquidity pools (representing approximately 96% of total volume traded) of the largest decentralized exchange, Uniswap V3 at the transaction level. This computationally-intensive exercise allows us to identify the changes in the asset composition of each liquidity position and makes it possible to calculate the individual LP's fees earned from each swap. Using this data, we separate LPs into two sets of sophisticated and retail market participants and examine the differences in their behaviour.

We show that liquidity provision, apart from the technological implementation, is not decentralized. Instead, we find that retail LPs are outcompeted by a small group of sophisticated agents that dominate liquidity provision. These players hold about 80% of total value locked and focus their attention on liquidity pools that have the most trading volume and are less volatile. Retail LPs capture a much lower share of trading fees than would be expected if liquidity provision was fully decentralised. They also manage their liquidity much less actively, and earn lower relative returns on their invested capital. While the average return for retail LPs is positive, this is driven by the high profitability of a select few. In more than half of the days in our sample, retail LPs lose money on a risk-adjusted basis.

The implications of our study extend beyond the immediate dynamics of liquidity provision on Uniswap V3. Our research looks at a single (large) DEX, but there is no reason to believe that Uniswap V3 is 'special'. Future work should shed light on the activity of sophisticated and retail market participants across multiple applications (e.g. lending and borrowing). The dominance of sophisticated liquidity providers challenges the fundamental ethos of DEXs, which is to democratize financial systems by removing intermediaries and to provide equal opportunities for all participants. Our findings highlight a trend where the ability to provide liquidity is becoming consolidated in the hands of a few sophisticated participants, with retail participants unable to effectively compete with them.

Not all of the reasons that underpin the emergence of different classes of intermediaries in traditional finance are present in DeFi. Regulatory, operational and technological barriers are clearly less prevalent. However, our results suggest that many of the economic forces that give rise to centralisation in traditional finance, where a relatively small set of intermediaries provides services in many markets, are likely inherent characteristics of the financial system, even in DeFi (Aramonte et al., 2021). Simply allowing all participants

access to a protocol does not appear to eliminate such forces and has not - to date - resulted in a truly disintermediated market.

References

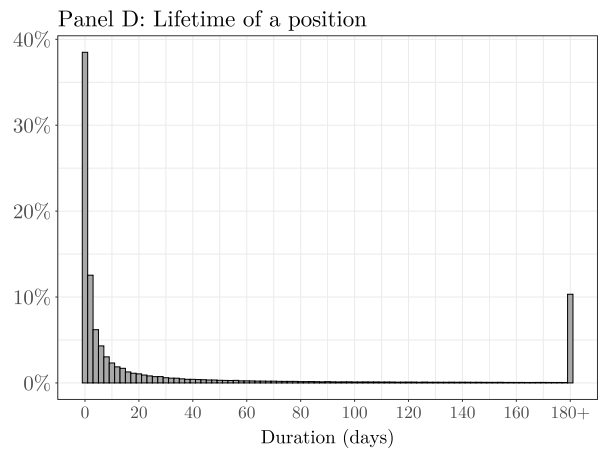
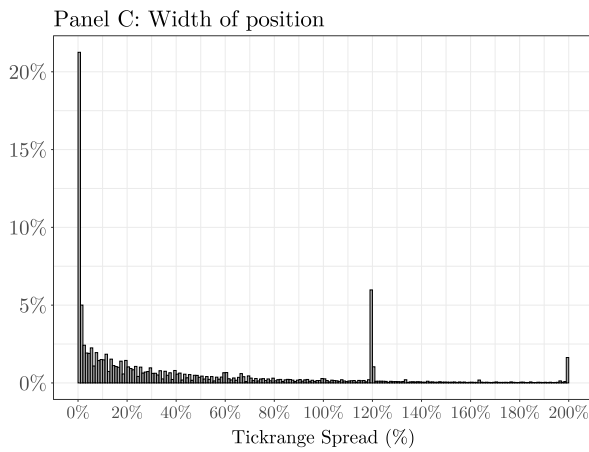
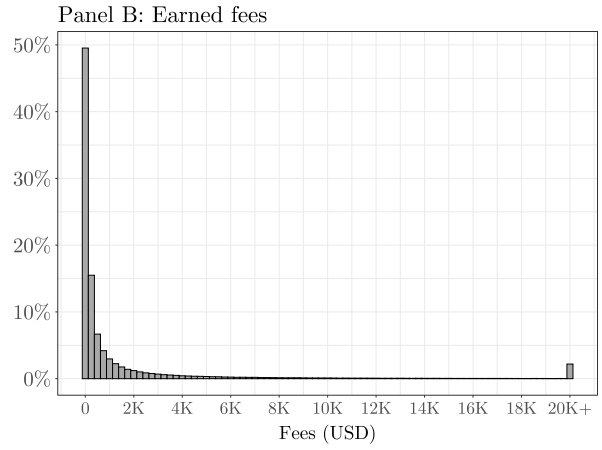
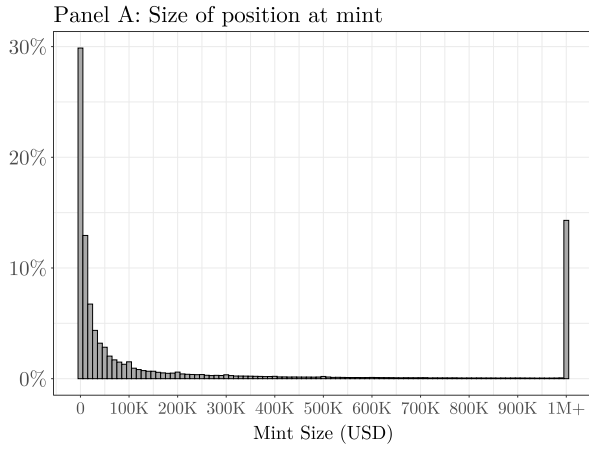
- Adams, H., Zinsmeister, N., Salem, M., Keefer, R., & Robinson, D. (2021). Uniswap v3 core. *Uniswap Labs*.
- Anand, A., & Venkataraman, K. (2016). Market conditions, fragility, and the economics of market making. *Journal of Financial Economics*, 121(2), 327–349.
- Angeris, G., & Chitra, T. (2020). Improved price oracles: Constant function market makers. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 80–91.
- Aquilina, M., Budish, E., & O’Neill, P. (2021). Quantifying the High-Frequency Trading “Arms Race”. *The Quarterly Journal of Economics*, 137(1), 493–564. <https://doi.org/10.1093/qje/qjab032>
- Aquilina, M., Frost, J., & Schrimpf, A. (2024). Decentralized Finance (DeFi): A Functional Approach. *Journal of Financial Regulation*, fjad013. <https://doi.org/10.1093/jfr/fjad013>
- Aramonte, S., Huang, W., & Schrimpf, A. (2021). Defi risks and the decentralisation illusion. *Bank of International Settlements Quarterly Review*.
- Aspris, A., Foley, S., Svec, J., & Wang, L. (2021). Decentralized exchanges: The “wild west” of cryptocurrency trading. *International Review of Financial Analysis*, 77, 101845.
- Barbon, A., & Ranaldo, A. (2023). On the quality of cryptocurrency markets: Centralized versus decentralized exchanges.
- Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., & Sokolov, K. (2018). High frequency trading and extreme price movements. *Journal of Financial Economics*, 128(2), 253–265.
- Budish, E., Cramton, P., & Shim, J. (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics*, 130(4), 1547–1621. <https://doi.org/10.1093/qje/qjv027>
- Budish, E., Lee, R., & Shim, J. J. (2024). A theory of stock exchange competition and innovation: Will the market fix the market? *Journal of Political Economy*, 132(4), 1209–1246. <https://EconPapers.repec.org/RePEc:ucp:jpolec:doi:10.1086/727284>
- Caparros, B., Chaudhary, A., & Klein, O. (2023). Blockchain scaling and liquidity concentration on decentralized exchanges. *arXiv preprint arXiv:2306.17742*.
- Capponi, A., & Jia, R. (2021). The adoption of blockchain-based decentralized exchanges.
- Capponi, A., Jia, R., & Zhu, B. (2023). *The Paradox Of Just-in-Time Liquidity in Decentralized Exchanges: More Providers Can Sometimes Mean Less Liquidity*.

- Capponi, A., Ruizhe, J., & Shiao, Y. (2023). Price Discovery on Decentralized Exchanges. *Working Paper w30949*.
- Cong, L., Tang, K., Wang, Y., & Zhao, X. (2023). Inclusion and Democratization Through Web3 and DeFi? Initial Evidence from the Ethereum Ecosystem. *NBER Working Paper w30949*.
- Cornelli, G., Gambacorta, L., Garratt, R., & Reghezza, A. (2024). *Why DeFi lending? Evidence from Aave V2*, Bank for International Settlements.
- Demsetz, H. (1968). The Cost of Transacting. *The Quarterly Journal of Economics*, 82(1), 33–53. <https://doi.org/10.2307/1882244>
- Duffie, D., Gârleanu, N., & Pedersen, L. H. (2005). Over-the-counter markets. *Econometrica*, 73(6), 1815–1847.
- Foley, S., Krekel, W., & Kwan, A. (2024). Fixed spreads and flexible inventories: Examining amm venue fragmentation. *Unpublished Working Paper*.
- Foley, S., Krekel, W., Mollica, V., & Svec, J. (2023). Not so fast: Identifying and remediating slow and imprecise cryptocurrency exchange data. *Finance Research Letters*, 51, 103401.
- Foley, S., O’Neill, P., & Putnins, T. (2023). *A better market design? applying “automated market makers” to traditional financial markets*.
- Foley, S., & Putniņš, T. J. (2016). Should we be afraid of the dark? dark trading and market quality. *Journal of Financial Economics*, 122(3), 456–481.
- Han, J., Huang, S., & Zhong, Z. (2022). Trust in defi: An empirical study of the decentralized exchange. *Available at SSRN 3896461*.
- Heimbach, L., Schertenleib, E., & Wattenhofer, R. (2022). Risks and returns of uniswap v3 liquidity providers. *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 89–101.
- Lehar, A., Parlour, C., & Zoican, M. (2023). Liquidity fragmentation on decentralized exchanges. *arXiv preprint arXiv:2307.13772*.
- Lehar, A., & Parlour, C. A. (2021). Decentralized exchanges. *Investments eJournal*. <https://api.semanticscholar.org/CorpusID:237189812>
- Loesch, S., Hindman, N., Richardson, M. B., & Welch, N. (2021). Impermanent loss in uniswap v3. *arXiv preprint arXiv:2111.09192*.
- Pagano, M. (1989). Endogenous Market Thinness and Stock Price Volatility. *The Review of Economic Studies*, 56(2), 269–287.
- Park, A. (2023). The conceptual flaws of decentralized automated market making. *Management Science*, 69(11), 6731–6751.

- Schär, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review*. <https://doi.org/https://doi.org/10.20955/r.103.153-74>
- Wan, X., & Adams, A. (2022). Just-in-time liquidity on the uniswap protocol. *Available at SSRN 4382303*.
- Xu, J., Paruch, K., Cousaert, S., & Feng, Y. (2023). Sok: Decentralized exchanges (dex) with automated market maker (amm) protocols. *ACM Computing Surveys*, 55(11), 1–50.

Appendix

A Histograms



B Robustness: Summary statistics sophisticated LPs

Table B.1 provides additional summary statistics across the various sophisticated classification mechanisms. The ‘retail’ group entails all wallet addresses that are not classified as sophisticated through any of the six mechanisms. The results show that the liquidity provisioning differs greatly between ‘retail’ and any ‘sophisticated’ group.

Table B.1: Summary statistics sophisticated wallet addresses: Robustness. This table reports summary statistics of wallet addresses classified as retail or sophisticated according to different metrics. A wallet address is considered retail if it is not classified as sophisticated through any of the metrics. Positions is the number of distinct liquidity positions, a wallet address has created. Pools is the number of distinct liquidity pools, a wallet address has provided liquidity in. MintSize (\$) is the mean mint size of a position in thousand US\$. Duration (days) is the mean time a wallet’s positions are open. Fees (\$) is the sum of earned fees in US\$. Return (bps) is the average of mean daily total returns (fee yield - impermanent loss) in bps. Interactions is the mean number of interactions with a liquidity position. TickRange (%) is the mean tickrange spread of a wallet’s positions in percentage. Stable/Stable, Stable/Token and Token/Token indicate the proportion of the respective pool classification across an LP’s positions.

	Retail	MintUSD 95th	MintUSD 1m	Positions 95th	Pools 95th	Interactions 95th	Tagged
	(n=78238)	(n=4415)	(n=2134)	(n=4855)	(n=4517)	(n=4506)	(n=848)
Variable	Mean	Mean	Mean	Mean	Mean	Mean	Mean
Positions	1.95	33.45	49.06	51.21	42.30	53.24	48.94
Pools	1.35	3.77	4.39	6.31	8.01	6.33	4.24
MintSize (k\$)	12.85	2,667	5,131	1,018	1,142	770.8	7,937
Duration (days)	209.1	49.55	45.41	26.33	61.89	36.26	75.03
Fees (\$)	563.1	8,870	10,840	2,422	2,153	3,114	4,124
Return (bps)	13.69	13.57	10.92	26.66	21.93	26.09	8.66
Interactions	2.98	5.86	6.62	4.35	4.14	9.58	25.42
TickRange (%)	73.25	39.41	34.10	34.64	51.63	38.08	27.64
TimeActive (%)	78.45	86.92	88.16	87.13	86.94	86.89	87.51
Stable/Stable (%)	5.35	14.51	14.08	3.22	4.68	4.39	8.71
Stable/Token (%)	48.72	46.42	50.59	46.77	37.97	43.95	54.50
Token/Token (%)	45.93	39.07	35.33	50.01	57.35	51.66	36.78

C Robustness: Regressions Sophisticated vs. Retail

This section further expands on the findings of Table 3 that analysed differences in liquidity provision of retail and sophisticated market participants. Instead of identifying sophisticated LPs only through the main classification mechanism (‘at least two criteria’), Table C.1 provides alternative classification thresholds based on how many different criteria are met, from 1 to all 6. We find monotonic increases in differences, the more stringent the sophisticated sample is.


$$Y_{i,t} = \alpha + \beta_1 \text{Sophisticated}_{i,t} + \varepsilon_{i,t} \quad (7)$$

Table C.1: The differences between retail and sophisticated LPs: Robustness This table reports results of an OLS regression analysing differences between liquidity positions of retail and sophisticated market participants. At the position level, the dependent variables are the US\$ size of the position at mint, the number of interactions with the liquidity position, the duration how long the position has been posted in days and the proportion of time the position has been active in percentage. At the LP level, the dependent variables are the number of distinct liquidity positions an address has minted and the number of distinct pools an address has provided liquidity in. The different models contrast classification mechanisms of sophisticated LPs. For example, ‘2 or more’ indicates that a wallet is classified as sophisticated if two of the six criteria are fulfilled. Sophisticated is a binary variable identifying if the liquidity position owner is considered sophisticated (fulfilling at least two criteria). Standard errors reported in parentheses are robust and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	1 or more	2 or more	3 or more	4 or more	5 or more	6 or more
Mint Size (US\$)	3,329,594***	3,678,810***	4,496,273***	7,165,927***	9,862,177***	13,192,926***
Duration (days)	-125.3***	-131.0***	-133.2***	-137.1***	-139.5***	-140.4***
Time Active (%)	9.43***	10.02***	10.74***	12.33***	13.64***	12.57***
Tickrange Spread (%)	-38.77***	-41.76***	-43.33***	-50.12***	-55.29***	-63.04***
#Interactions	1.25***	1.30***	1.29***	1.43***	1.36***	2.22***
#Positions	25.70***	39.09***	59.83***	103.4***	165.1***	306.6***
#Pools	3.38***	4.20***	6.51***	7.25***	9.72***	16.44***

D Example Etherscan label

This figure provides an example of a labelled ‘sophisticated’ crypto wallet address on Etherscan. Here, the wallet address `0x073Dca8ACbC11ffB0b5Ae7ef171e4c0b065FfA47` has been associated with ‘Alameda Research’, the principal trading firm related to the now defunct FTX ecosystem. Besides the label, Etherscan, where known, provides the source of the label as well as an overview of the account balance.



The screenshot shows the Etherscan interface for a specific wallet address. At the top, the address `0x073Dca8ACbC11ffB0b5Ae7ef171e4c0b065FfA47` is displayed with a globe icon and utility icons. Below this, a grey box contains the text: "This address is associated with [Alameda Research](#) and is provided by [Larry Cermak's](#) list of Alameda Research's Wallets." Underneath, three tags are shown: "Alameda Research 1" (with a link icon), "# Alameda Research", and "# Fund". A red box highlights these tags, and an arrow points to them with the label "Label information". The page is divided into two main sections: "Overview" on the left and "More Info" on the right. The "Overview" section shows the "ETH BALANCE" as `0.025845195017698644 ETH`. The "More Info" section is titled "PRIVATE NAME TAGS" and includes a "+ Add" button.

E Summary statistics: Profitability of liquidity positions

Table E.1: Summary statistics: Profitability of liquidity positions (timeseries). This table reports summary statistics of profitability characteristics of liquidity positions per day. Panel A shows the full sample. Panel B differentiates between retail and sophisticated liquidity positions, where positions are classified as retail if the owner fulfills at most one of the six criteria. Fees (\$) is the total absolute accrued fee in US \$. Fee Yield (bps) is the mean daily fee yield calculated as the ratio of accrued fees on invested capital in basis points. IL (bps) is the mean daily impermanent loss calculated as the ratio of the value of assets in a liquidity pool on their value if held outside the pool in basis points. Total Return (bps) is the daily total return (fee yield - impermanent loss) in basis points. Net Return (bps) is the Total Return reduced by transaction costs in the form of gas fees. Excess Return is the differential of the Net Return and risk-free rate. The metrics are winsorized at the 99th percentile level.

Panel A: Overall sample								
Variable	N	Mean	Std. Dev.	Q5	Q25	Median	Q75	Q95
Fees (\$)	28,759,123	3,131	56,148	0.06	0.96	14.76	254.1	5,405
Fee Yield (bps)	28,759,123	5.82	15.28	0.00	0.00	0.88	4.39	27.25
IL (bps)	28,759,123	-1.13	4.38	-5.92	-0.10	0.00	0.00	0.00
Total Return (bps)	28,759,123	4.69	13.95	0.00	0.00	0.40	3.29	22.87
Net Return (bps)	28,759,123	4.30	13.65	-0.18	0.00	0.34	3.18	22.06
Excess Return (bps)	28,759,123	3.45	13.75	-1.45	-1.10	-0.05	2.35	21.52
Panel B: Sophisticated / Retail								
Variable	N	Mean	Std. Dev.	Q5	Q25	Median	Q75	Q95
Group: Sophisticated								
Fees (\$)	4,087,150	17,104	146,446	0.27	66.67	533.9	3,288	46,520
Fee Yield (bps)	4,087,150	11.31	23.50	0.00	0.00	1.28	9.80	67.95
IL (bps)	4,087,150	-1.55	5.49	-9.62	-0.07	0.00	0.00	0.00
Total Return (bps)	4,087,150	9.76	21.99	0.00	0.00	0.65	7.78	62.27
Net Return (bps)	4,087,150	8.72	21.17	-1.41	0.00	0.43	7.14	59.63
Excess Return (bps)	4,087,150	7.98	21.25	-2.18	-1.10	-0.01	6.52	59.09
Group: Retail								
Fees (\$)	24,671,973	816.7	9,182	0.05	0.69	7.76	128.3	2,522
Fee Yield (bps)	24,671,973	4.91	13.22	0.00	0.00	0.84	3.92	21.39
IL (bps)	24,671,973	-1.06	4.16	-5.46	-0.10	0.00	0.00	0.00
Total Return (bps)	24,671,973	3.85	11.91	0.00	0.00	0.37	2.96	17.48
Net Return (bps)	24,671,973	3.57	11.80	-0.08	0.00	0.33	2.90	17.05
Excess Return (bps)	24,671,973	2.70	11.90	-1.45	-1.10	-0.08	2.00	16.50

F Robustness Regression: Profitability

This section further expands on the findings of Table 5 that analysed differences in profitability of retail and sophisticated market participants. Instead of identifying sophisticated LPs only through the main classification mechanism ('at least two criteria'), Table F.1 provides alternative classification thresholds based on how many different thresholds are met, from 1 to all 6.

$$Y_{i,t} = \alpha + \beta_1 \text{Sophisticated}_{i,t} + \beta_2 \log(\text{Volatility}_{i,t}) + \beta_3 \log(\$TVL_{i,t}) + \beta_4 \log(\$Volume_{i,t}) + \varepsilon_{i,t} \quad (8)$$

Table F.1: The profitability of sophisticated and retail LPs: Robustness. This table reports results of a fixed effect OLS regression analysing profitability measures across liquidity pools and liquidity positions. Coefficients of the controls are not shown to improve readability and the displayed data is the coefficient of the sophisticated dummy variable in the shown equation. The dependent variables Position Fees (\$), Fee Yield (bps), Impermanent Loss (bps), Total Return (bps) and Net Return (bps) are calculated at the position, day level as outlined before. The columns represent the classification scheme of the sophisticated, i.e. how many criteria have to be met for an LP to be assigned to the sophisticated group. Sophisticated is a binary variable identifying if the liquidity position owner is considered sophisticated. Volatility is the absolute value of the daily high-low range scaled by its midpoint. TVL and Volume are the respective pool, day TVL and trading volume in US\$. Fee Yield, Impermanent Loss, Total Return and Net Return are winsorized at the 99th percentile. Excess Returns equal Net Returns when using date fixed-effects. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	1 or more	2 or more	4 or more	5 or more	6 or more
Position Fees (US\$)	185.2***	268.5***	675.6***	1,021***	2,502***
Fee Yield (bps)	2.59***	3.63***	3.15***	3.02***	5.94***
Impermanent Loss (bps)	-0.18***	-0.20***	0.029	0.11	0.79***
Total Return (bps)	2.41***	3.42***	3.18***	3.13***	6.74***
Net Return (bps)	2.06***	2.90***	2.80***	2.81***	3.72***

G Transaction Costs Computation

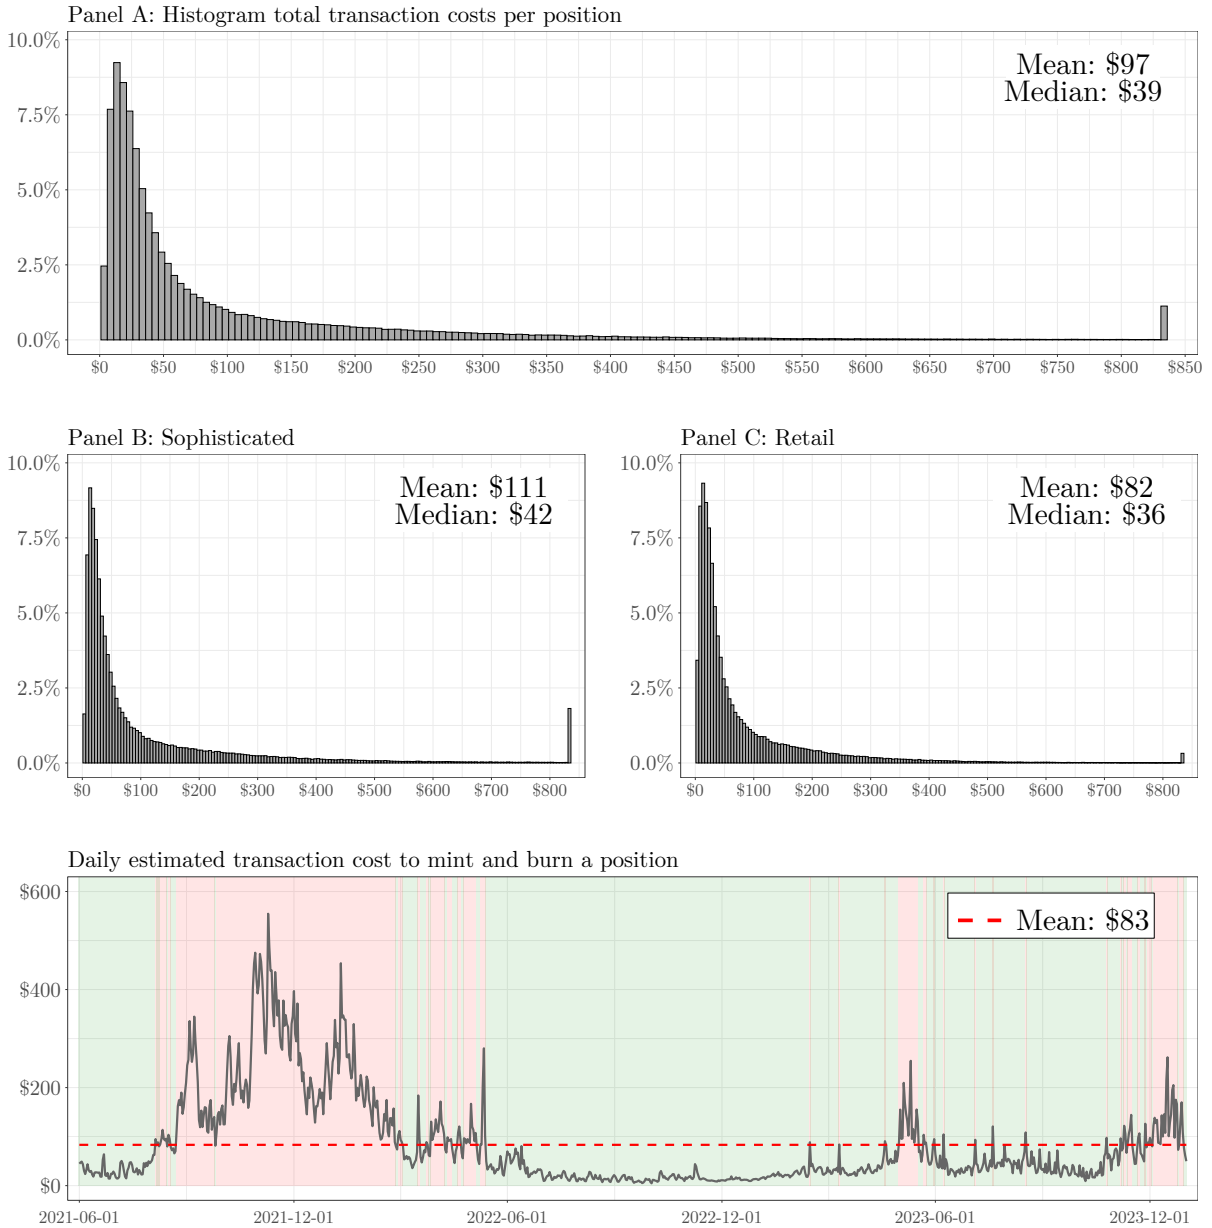
Figure G.1, Panel A shows the histogram of total transaction costs per liquidity position. LPs spend on average USD 97 (median USD39) on their liquidity position management. Notably, the distribution has a fat right-tail, in which liquidity providers paid considerably high gas prices: either as part of a trading strategy (e.g. JIT liquidity provision) or seemingly in error.

Panel B and C plot the transaction cost distribution for the positions of ‘Sophisticated’ and ‘Retail’. It becomes evident that sophisticated LPs spend a higher amount on gas fees. We find two potential explanations. First, retail Lps are more sensitive to higher gas prices. For retail positions, transaction costs relative to their smaller position size are higher than for sophisticated LPs. For the mean sized retail position (about USD 30,000) a USD 100 transaction cost equals a drag of 30bps, whereas this number is less than 1bps for the average sophisticated position, and as such is negligible. It is then unsurprising that sophisticated LPs favor certainty and immediacy to be included in the next block or even at a certain block position, especially when deploying certain trading strategies such as JIT liquidity.²⁸

We further estimate the USD it would cost to mint and burn a liquidity position on a given day during our sample period using the daily median *gasPrice* paid and the previously calculated amounts of *gas*. The time series is depicted in Figure G.1, Panel B. The volatility is entirely due to the variable *gasPrice*, which in USD terms is dependent on the USD price of ETH and the *gasPrice* in ETH terms. We find the daily estimated transaction costs to be higher during the earlier periods of our sample, in which both the USD price of ETH and the average *gasPrice* was higher.

²⁸While the higher transaction costs of sophisticated market participants could also stem from them minting and burning liquidity positions during the earlier part of our sample in which transaction costs were higher (see Figure G.1, Panel D), we do not find such evidence. We refer to Lehar et al. (2023) for an extensive discussion on how gas costs affect liquidity provision in AMMs.

Figure G.1: This figure illustrates the network or ‘gas’ fees of liquidity providers. Panel A shows the histogram of paid gas fees for all liquidity position. For a single position, gas fees are derived by totaling fees of all mint and burn transactions of a liquidity position. Panel B and C only include positions of classified sophisticated or retail market participants. Panel D showcases the timeseries of gas fees that would have been paid to mint and burn a liquidity position on a given day, which is calculated by multiplying the mean gas used for a mint ($gas_{mint} = 424,300$) and burn ($gas_{burn} = 238,300$) times the daily mean gas price of all mint and burns. The red (green) shaded areas correspond to days where the transaction costs are above (below) the mean. Values are converted to USD through hourly or day-end Ether prices sourced from CCData. Transaction costs are winsorized at the 99th percentile.



H Profitability during high and low gas costs

Table H.1: The profitability of sophisticated and retail LPs during times of high and low gas costs. This table reports results of a fixed effect OLS regression analysing profitability measures of liquidity positions during times of high and low gas costs. The dependent variables Fee Yield (bps), Total Returns (bps) and Net Return (bps) are as outlined before. Institutional is a binary variable identifying if the metric belongs to the group of sophisticated or the liquidity position owner is considered sophisticated, where sophisticated fulfills at least two classification criteria. Volatility is the absolute value of the daily high-low range scaled by its midpoint. TVL and Volume are the respective pool, day TVL and trading volume in USD. Fee Yield, Impermanent Loss and Total Return are winsorized at the 99th percentile. Standard errors reported in parentheses are clustered at the liquidity pool level and ***, **, * denote the statistical significance at the 1, 5, and 10% level, respectively.

	Overall			High gas regime			Low gas regime		
	Fee Yield (bps)	Total Return (bps)	Net Return (bps)	Fee Yield (bps)	Total Return (bps)	Net Return (bps)	Fee Yield (bps)	Total Return (bps)	Net Return (bps)
Retail	-3.54*** (0.34)	-3.36*** (0.31)	-2.84*** (0.27)	-3.50*** (0.36)	-3.34*** (0.34)	-2.58*** (0.30)	-3.42*** (0.35)	-3.23*** (0.32)	-2.83*** (0.28)
log(Volatility)	0.10 (0.43)	-0.29 (0.36)	-0.25 (0.33)	0.92 (0.65)	0.22 (0.52)	0.18 (0.47)	0.01 (0.33)	-0.32 (0.28)	-0.27 (0.26)
log(\$TVL)	-3.82*** (0.41)	-3.58*** (0.36)	-3.33*** (0.32)	-4.90*** (0.61)	-4.65*** (0.54)	-4.30*** (0.49)	-3.55*** (0.45)	-3.34*** (0.40)	-3.08*** (0.36)
log(\$Volume)	3.78*** (0.34)	3.44*** (0.31)	3.15*** (0.28)	4.79*** (0.49)	4.38*** (0.45)	4.01*** (0.40)	3.14*** (0.35)	2.84*** (0.32)	2.60*** (0.27)
N	28,317,758	28,317,758	28,317,758	7,428,795	7,428,795	7,428,795	20,641,588	20,641,588	20,641,588
FE: Pool	X	X	X	X	X	X	X	X	X
FE: Date	X	X	X	X	X	X	X	X	X
Adj. R2	0.27	0.24	0.22	0.33	0.31	0.28	0.24	0.21	0.19

Previous volumes in this series

1226 November 2024	How does fiscal policy affect the transmission of monetary policy into cross-border bank lending? Cross-country evidence	Swapan-Kumar Pradhan, Előd Takáts and Judit Temesvary
1225 October 2024	Trade fragmentation, inflationary pressures and monetary policy	Ludovica Ambrosino, Jenny Chan and Silvana Tenreyro
1224 October 2024	A theory of economic coercion and fragmentation	Chris Clayton, Matteo Maggiori and Jesse Schreger
1223 October 2024	Carbon prices and reforestation in tropical forests	José A Scheinkman
1222 October 2024	Artificial intelligence and big holdings data: opportunities for central banks	Xavier Gabaix, Ralph S J Koijen, Robert Richmond and Motohiro Yogo
1221 October 2024	Bank geographic diversification and funding stability	Sebastian Doerr
1220 October 2024	The road to net zero: a fund flow investigation	Louisa Chen and Koji Takahashi
1219 October 2024	Stablecoins, money market funds and monetary policy	Iñaki Aldasoro, Giulio Cornelli, Massimo Ferrari Minesso, Leonardo Gambacorta, Maurizio Michael Habib
1218 October 2024	Bank specialisation and corporate innovation	Hans Degryse, Olivier De Jonghe, Leonardo Gambacorta and Cédric Huylebroek
1217 October 2024	Global inflation, inflation expectations and central banks in emerging markets	Ana Aguilar, Rafael Guerra and Berenice Martinez
1216 October 2024	Trade credit and exchange rate risk pass-through	Bryan Hardy, Felipe Saffie and Ina Simonovska
1215 October 2024	CB-LMs: language models for central banking	Leonardo Gambacorta, Byeungchun Kwon, Taejin Park, Pietro Patelli and Sonya Zhu
1214 September 2024	The impact of financial crises on industrial growth: lessons from the last 40 years	Carlos Madeira

All volumes are available on our website www.bis.org.