

Liquidity fragmentation on decentralized exchanges

Alfred Lehar
Christine A. Parlour
Marius Zoican*

October 3, 2023

Abstract

We study economies of scale in liquidity provision on decentralized exchanges, focusing on the impact of fixed transaction costs such as gas prices on liquidity providers (LPs). Small LPs are disproportionately affected by the fixed cost, resulting in liquidity supply fragmentation between low- and high-fee pools. Analyzing Uniswap data, we find that high-fee pools attract 58% of liquidity supply but execute only 21% of trading volume. Large (institutional) LPs dominate low-fee pools, frequently adjusting positions in response to substantial trading volume. In contrast, small (retail) LPs converge to high-fee pools, accepting lower execution probabilities to mitigate smaller liquidity management costs.

Keywords: FinTech, decentralized exchanges (DEX), liquidity, fragmentation, retail trading

JEL Codes: G11, G12, G14

*Alfred Lehar (alfred.lehar@haskayne.ucalgary.ca) is affiliated with Haskayne School of Business at University of Calgary. Christine A. Parlour (parlour@berkeley.edu) is with Haas School of Business at UC Berkeley. Marius Zoican (marius.zoican@rotman.utoronto.ca) is affiliated with University of Toronto Mississauga and Rotman School of Management. We have greatly benefited from discussion on this research with Michael Brolley, Itay Goldstein, Sang Rae Kim (discussant), Olga Klein, Katya Malinova (discussant), Uday Rajan, Thomas Rivera, Gideon Saar, Andriy Shkilko (discussant), and Shihao Yu. We are grateful to conference participants at the Edinburgh Economics of Technology, Financial Intermediation Research Society 2023, the Northern Finance Association 2023, the UNC Junior Faculty Finance Conference, as well as to seminar participants at the Microstructure Exchange. Marius Zoican gratefully acknowledges funding support from the Rotman School of Managements' FinHub Lab and the Canadian Social Sciences and Humanities Research Council (SSHRC) through an Insight Development research grant (430-2020-00014).

Liquidity fragmentation on decentralized exchanges

Abstract

We study economies of scale in liquidity provision on decentralized exchanges, focusing on the impact of fixed transaction costs such as gas prices on liquidity providers (LPs). Small LPs are disproportionately affected by the fixed cost, resulting in liquidity supply fragmentation between low- and high-fee pools. Analyzing Uniswap data, we find that high-fee pools attract 58% of liquidity supply but execute only 21% of trading volume. Large (institutional) LPs dominate low-fee pools, frequently adjusting positions in response to substantial trading volume. In contrast, small (retail) LPs converge to high-fee pools, accepting lower execution probabilities to mitigate smaller liquidity management costs.

Keywords: FinTech, decentralized exchanges (DEX), liquidity, fragmentation, retail trading

JEL Codes: G11, G12, G14

1 Introduction

The way in which economic agents react to transaction costs affects the extent to which such costs are detrimental to economic efficiency. Or, as Demsetz (1968) puts it “the question that is relevant for efficiency is whether or not the cost is appropriately economized.” In this paper, we investigate a unique environment that allows us to investigate, theoretically and empirically, how explicit transactions costs affect liquidity supply.

We investigate liquidity provision on a decentralized exchange (Uniswap v3). Automated exchanges such as Uniswap present a unique laboratory to investigate how transaction costs affect the market for liquidity. While there are various new institutional details that animate these exchanges, for our purposes two are economically important. First, in automated exchanges, liquidity demand and supply can easily be distinguished: users either post liquidity or initiate a trade. So, we can isolate the effect of transactions costs on each side of the market for liquidity. Second, prices are not set by market participants, but are automatically calculated as a function of liquidity demand and supply. Thus, costs are not passed from liquidity suppliers to demanders through prices.

Briefly, Uniswap v3 allows liquidity suppliers and demanders to select into trading places (called pools) that differ on transaction fees. These proportional fees are paid by liquidity demanders and are the only source of remuneration to liquidity suppliers. As noted, prices are calculated mechanically based on posted liquidity, so liquidity suppliers do not earn price impact. In addition, all participants pay a fixed transaction cost (called a “gas fee”) to access the markets. Our theory and empirical work investigates the effect of different proportional fees and fixed fees on liquidity supply.

To do so, we present a simple model of liquidity supply in which there are two markets – one has high fees, and one has low fees. Liquidity demanders, who can either be small or large, choose a venue based on the fees each venue charges. The decisions of the liquidity demanders determine the payoff to the liquidity suppliers. Liquidity suppliers have heterogeneous endowments, interpretable as different capital constraints — low-endowment liquidity providers are akin to retail traders, whereas high-endowments stand in for large institutional investors or quantitative funds. Further, liquidity providers incur a fixed cost (i.e., gas price) each time they update their position.

Traders demanding liquidity route their orders first to the low-fee pool to minimize transaction costs. The high-fee pool executes the residual order flow from large traders who exhaust all liquidity at the cheaper venue. As a result, low-fee markets are actively traded and require frequent liquidity updates whereas high-fee pools have a longer liquidity update cycle since they absorb fewer trades.

We establish conditions under which there is fragmentation or consolidation. Specifically, even in this simple framework there is a robust parameter range in which liquidity does not naturally

concentrate on one of the exchanges. Not only does liquidity fragment, but it differs in both use and type across the two markets. That is, one exchange features liquidity provided by large suppliers while on the other smaller (retail) suppliers predominate. Finally, changes in the common fixed market access fee differentially affects the liquidity supply on the two pools. Specifically it reduces market quality (in the sense of lower posted liquidity) on the low fee pool.

Our paper also shows that fixed costs translate to significant economies of scale, which lead to market fragmentation and the emergence of liquidity provider clienteles. Specifically, our findings reveal that a small number of highly active large liquidity providers, potentially institutional investors and hedge funds, primarily trade against numerous small incoming trades on pools with low fees. In contrast, high-fee pools involve less frequent trading between a substantial number of capital-constrained passive liquidity providers (e.g., retail market makers) on one side and a few sizeable incoming orders on the other. The underlying economic mechanism is that smaller liquidity providers strategically trade off a lower execution probability against higher liquidity fee and less frequent fixed costs, such as gas fees.

Both pools can attract a positive market share if the gas price is large enough due to economies of scale. Liquidity providers trade off a higher revenue per unit of time in the low-fee pool (driven by the larger trading volume) against the additional gas cost required for active liquidity management. As a result, market maker clienteles emerge in equilibrium. Liquidity providers with large endowments gravitate towards low-fee markets, as they are best positioned to frequently update their position. In contrast, smaller market makers choose to passively provide liquidity on high-fee markets where they only trade against large orders being routed there. They optimally trade off a lower execution probability against higher fees per unit of volume and a lower liquidity management cost per unit of time.

Uniswap v3 launched in May 2021 and quickly established itself as the leading decentralized exchange. Daily trading volume on Uniswap v3 typically exceeds US \$1 billion, whereas the total liquidity supply varies between US \$2-4 billion in our sample, across hundreds of cryptocurrencies. For major pairs such as Ether against USD stablecoins, Uniswap boasts twice or three times better liquidity than centralized exchanges such as Binance.¹

The market design explicitly allows for fragmentation. Each asset pair to be traded on up to four liquidity pools that only differ in the compensation for liquidity providers: in particular, liquidity fees can be equal to 1, 5, 30, or 100 basis points and the corresponding tick sizes are 1, 10, 60, or 200 basis points. The liquidity pools are otherwise identical and, importantly, they share the common infrastructure of the Ethereum blockchain – i.e., they are similarly impacted by variation in gas prices due to, say, network congestion. At launch, Uniswap Labs conjectured that trading and liquidity should consolidate in equilibrium on a single “canonical” pool for which the liquidity

¹See [The Dominance of Uniswap v3 Liquidity](#); May 5, 2022.

fee is just enough to compensate the marginal market maker for adverse selection and inventory costs. That is, activity in low-volatility pairs such as stablecoin-to-stablecoin trades should naturally gravitate to low fee liquidity pools, whereas speculative trading in more volatile pairs will consolidate on high fee markets.²

Using the model for guidance, we analyze more than 28 million interactions with Uniswap v3 liquidity pools – that is, all liquidity updates and trades from the inception of v3 in May 2021 until July 2023. We first document liquidity fragmentation in 32 out of 242 asset pairs in our sample, which account for 95% of liquidity committed to Uniswap v3 smart contracts and 93% of trading volume. For each of the fragmented pairs, trading consolidates on two pools with adjacent fee levels: either 1 and 5 basis points (e.g., USDC-USDT), 5 and 30 basis points (ETH-USDC), or 30 and 100 basis points (USDC-CRV).

We then document that high-fee pools are on average larger – with aggregate end-of-day liquidity of \$46.50 million relative to \$33.78 million, the average size of low-fee pools. Nevertheless, three quarters of daily trading volume executes on low-fee pools. In line with the model predictions, low-fee pools are more active as they capture many small trades. There are five times as many trades on low- than on high-fee pools (610 versus 110). However, the average trade on the high fee pool is twice as large: \$14,490 relative to \$6,340. Unsurprisingly, liquidity cycles – measured as the time between the submission and update of posted liquidity – are 20% shorter on the highly active low-fee pool.

We find robust evidence of liquidity supply clienteles across pools. The average liquidity deposit is 107.5% larger on the low-fee pool, after controlling for daily volume and return volatility. At the same time, high-fee pools’ market share is 21 percentage points higher. The results point to an asymmetric match between liquidity supply and demand: large liquidity suppliers are matched with small liquidity demanders on low-fee pools, whereas small liquidity suppliers trade with a few large orders on the high-fee pool.

We then turn to the common fixed cost of accessing the market, or gas fees. The degree of market fragmentation depends on the magnitude of gas costs on the Ethereum blockchain. In the model, a higher gas price leads to a shift in liquidity supply from the low- to the high-fee pool as active position management becomes relatively more costly for the marginal liquidity supplier. We find that a one standard deviation increase in gas prices corresponds to a 4.63 percentage points decrease in the low-fee pool market share, and a 29% drop in liquidity inflows on days when gas costs are elevated. As liquidity suppliers leave the low-fee pool, the time between position updates becomes 4.19% shorter since incoming order flow depletes liquidity at a faster pace. The effect reinforces the expected gas cost difference between the two pools, and more liquidity suppliers switch to the high-fee venue.

²See *Flexible fees* paragraph at <https://uniswap.org/blog/uniswap-v3>; accessed September 14, 2022.

Our paper is related to various literatures. The classical study of [Pagano \(1989\)](#) shows that if an asset is traded on two identical exchanges with equal transaction costs, in equilibrium market participants gravitate to a single exchange due to network effects. In practice, exchanges are rarely identical: fragmentation can emerge between fast and slow exchanges ([Pagnotta and Philippon, 2018](#); [Brolley and Cimon, 2020](#)) or between lit and dark markets ([Zhu, 2014](#)). In our model, fragmentation on decentralized exchanges is driven by variation in liquidity fees as well as different economies of scale due to heterogeneity in liquidity provider capital. We find that liquidity fragmentation driven by high gas fees implies larger transaction costs on incoming orders. We note that there is no time priority on decentralized exchanges, which clear in a pro rata fashion. On markets with time priority, [Foucault and Menkveld \(2008\)](#) and [O’Hara and Ye \(2011\)](#) find that market segmentation in equity markets improves liquidity (by allowing queue jumping) and price discovery.

Fixed costs for order submission are uncommon in traditional markets. However, in 2012, the Canadian regulator IROC implemented an “integrated fee model” that charged traders for all messages sent to Canadian marketplaces. [Korajczyk and Murphy \(2018\)](#) document that this measure disproportionately affected high-frequency traders, resulting in wider bid-ask spreads but lower implementation shortfall for large traders, possibly due to a reduction in back-running activity. Our study contributes additional insights by highlighting that the introduction of a fixed cost, even when applied across exchanges, can lead to market fragmentation.

We also relate to a rich literature on market fragmentation and differential fees. Closest to our paper, [Battalio, Corwin, and Jennings \(2016\)](#) and [Cimon \(2021\)](#) study the trade-off between order execution risk and compensation for liquidity provision in the context of make-take fee exchanges. However, [Battalio, Corwin, and Jennings \(2016\)](#) specifically addresses the issue of the broker-customer agency problem, whereas our study focuses on liquidity providers who trade on their own behalf. In traditional securities markets, make-take fees are contingent on trade execution and proportional to the size of the order. On the other hand, gas costs on decentralized exchanges are independent of order execution, highlighting the significance of economies of scale (lower proportional costs for larger liquidity provision orders) and dynamic liquidity cycles (managing the frequency of fixed cost payments). Strategic brokers in [Cimon \(2021\)](#) provide liquidity alongside exogenous market-makers in a static setting. We complement this approach by modelling network externalities inherent in the coordination problem of heterogeneous liquidity suppliers. In our dynamic setup, this allows us to pin down the equilibrium duration of liquidity cycles and the relative importance of gas fixed costs.

Our paper relates to a nascent and fast-growing literature on the economics of decentralized exchanges. Many studies (e.g., [Aoyagi, 2020](#); [Aoyagi and Ito, 2021](#); [Park, 2022](#)) focus on the economics of constant-function automated market makers, which do not allow liquidity providers to set price limits. In this restrictive environment, [Capponi and Jia \(2021\)](#) argue that market makers have little incentives to update their position upon the arrival of news to avoid adverse selection,

since pro-rata clearing gives an advantage to arbitrageurs. [Lehar and Parlour \(2021\)](#) solve for the equilibrium pool size in a setting where liquidity providers fully internalize information costs without rushing to withdraw positions at risk of being sniped. We argue that on exchanges that allow for limit or range orders, the cost of actively managing positions becomes a first-order concern, as liquidity suppliers need to re-set the price limits once posted liquidity no longer earns fees. Our empirical result on economies of scale echoes the argument in [Barbon and Rinaldo \(2021\)](#), who compare transaction costs on centralized and decentralized exchanges and find that high gas prices imply that the latter only become competitive for transactions over US\$100,000. [Hasbrouck, Rivera, and Saleh \(2022\)](#) argue that liquidity suppliers require remuneration. We complement the argument by stating that high fees might be *necessary* for some liquidity providers to cover the fixed costs of managing their position. In line with our theoretical predictions, [Caparros, Chaudhary, and Klein \(2023\)](#) find that liquidity providers reposition their quotes more often on Uniswap V3 pools built on Polygon, which features substantially lower gas fees. Finally, [Heimbach, Schertenleib, and Wattenhofer \(2022\)](#) document that after accounting for price impact, concentrated liquidity on Uniswap v3 pools results in increased returns for sophisticated participants but losses for retail traders.

Despite higher gas costs, decentralized exchanges may hold advantages over centralized venues. [Han, Huang, and Zhong \(2022\)](#) demonstrate Uniswap frequently leads price discovery compared to centralized exchanges such as Binance, despite the latter having higher trading volume. [Capponi, Jia, and Yu \(2023\)](#) find that the fee paid traders to establish execution priority unveils their private information, and therefore contributes to price discovery. [Aspris, Foley, Svec, and Wang \(2021\)](#) argue that decentralized exchanges offer better security than their centralized counterparts since assets are never transferred to the custody of a third party such as an exchange wallet. In turn, [Brolley and Zoican \(2022\)](#) make the point that decentralized exchanges may be able to reduce overall computational costs associated with latency arbitrage races, as they eliminate long-term co-location subscriptions.

2 Model

Consider the following continuous time model of trade in a single token, \mathbf{T} . The expected value of the token, $v > 0$, is common knowledge. Two risk neutral trader types consummate trade in this market: a continuum of liquidity providers (\mathbf{LPs}) and a continuum of liquidity takers (\mathbf{LTs}). Trade occurs because market participants have heterogeneous private values for the asset. In particular, liquidity providers have no private value for the token, while liquidity takers value the token at $v(1 + \mathcal{I}\Delta)$. Here $\Delta > 0$ are the gains from trade and \mathcal{I} is an indicator that takes on the value of 1 if the taker buys and -1 if the taker sells. In what follows for expositional simplicity, as in [Foucault, Kadan, and Kandel \(2013\)](#), we focus on a one sided market in which liquidity takers act as buyers.

Liquidity providers differ in their endowments of the token. Each provider i can supply at most q_i of the token, where q_i follows a truncated Pareto distribution on $[1, Q]$. So,

$$\varphi(q) = \left(\frac{Q}{Q-1}\right) \frac{1}{q^2} \quad \text{for } q \in [1, Q]. \quad (1)$$

The right skew of the Pareto distribution captures the idea that there are many low-endowment liquidity providers such as retail traders, but few high-capital **LPs** such as sophisticated quantitative funds. Heterogeneity in **LP** size is captured by Q , where a larger Q naturally corresponds to a larger dispersion of endowments. Given the endowment distribution, collectively **LPs** supply at most

$$S = \int_1^Q q\varphi(q) dq = \frac{Q}{Q-1} \log Q \quad (2)$$

tokens.

There are two types of liquidity takers: small and large. Small **LTs** arrive at the market at constant rate θdt , and each demand one unit of the token. The large **LT** arrival time follows a Poisson process with rate $\lambda > 0$. Conditional on arrival, a large **LT** demands Θ units of the token, where $\Theta > S$. That is, the large **LT** liquidity demand exceeds the maximum liquidity supply. Let D_t denote aggregate liquidity demand. Then,

$$dD_t = \theta dt + \Theta dJ_t(\lambda), \quad (3)$$

where $J_t(\lambda)$ is the Poisson arrival process.

Liquidity demanders and suppliers can interact in two liquidity pools in which token trade occurs against a numéraire asset (cash). We assume that the terms of trade are fixed, so that all trades occur at the expected value of the token, v . We do this by assuming prices in both pools satisfy a linear bonding curve, so there is no price impact of trade.³ So, for a pool with T tokens and N of the numéraire good,

$$vT + N = \text{constant}. \quad (4)$$

It is well known that price impact costs can lead to order splitting and fragmentation. Fixing the terms of trade allows us to focus on how fees affect liquidity fragmentation.

Fees are levied on liquidity takers as a fraction of the value of the trade and distributed pro rata to liquidity providers. The pools have different fees. One pool charges a low fee, and one pool charges a high fee which we denote ℓ and h respectively. Specifically, to purchase τ units of the token on the low fee pool, the total cost to a taker is $\tau(v + \ell)$. The **LPs** in the pool receive $\tau\ell$ in fees.

³In practice, most decentralized exchanges use convex bonding curves, for example constant product pricing.

In addition, consistent with gas costs on Ethereum, any interaction with a liquidity pool (for example, trading or managing liquidity) incurs a fixed execution cost $\Gamma > 0$. It is important to note that gas fees on decentralized exchanges differ from trading fees on traditional exchanges. First, they are not set by individual exchanges to compete with each other, but are a common transaction cost *across* trading venues. Second, they are levied on a per-order rather than per-share basis: this implies economies of scale for larger orders. Third, there is significant time variation in gas fees which will allow us to identify the impact of transaction costs on liquidity pool market shares.⁴

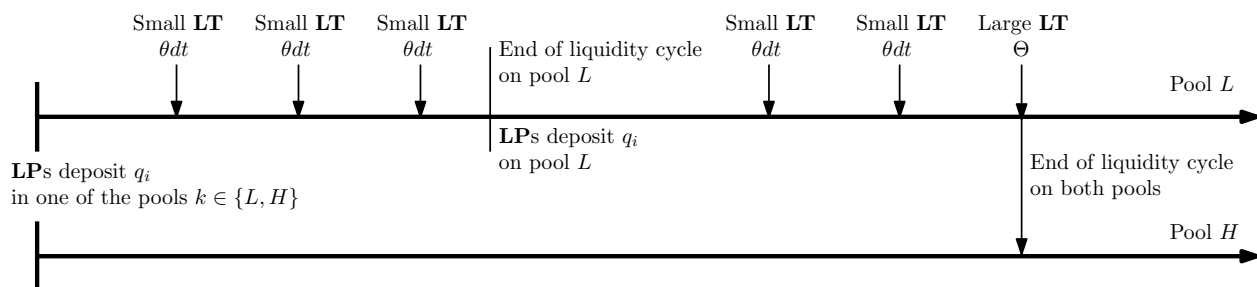
To ensure that both small and large liquidity traders participate in the market, we assume that the gains from trade are larger than the aggregate transaction cost, including the pool fee and the gas price. To rule out trivial cases, we further assume that $Q\ell - \Gamma > 0$. The condition ensures that there are at least some liquidity providers who can earn positive expected profit on the low fee pool.

Assumption 1: The gains from trade are sufficiently large so that all liquidity traders participate in the market, and the fixed costs are sufficiently low so that both pools attract liquidity.

- i. Gains from trade are sufficiently large $\Delta > h + \frac{\Gamma}{v}$.
- ii. The low fee pool attracts some liquidity $Q\ell - \Gamma > 0$

We follow [Foucault, Kadan, and Kandel \(2013\)](#) and partition the continuous timeline into *liquidity cycles*. A liquidity cycle starts with an empty pool (zero liquidity offered) which triggers **LP** token deposits and ends when incoming trades deplete the liquidity supply. The first liquidity cycle starts at $t = 0$, and the sequence of events within a cycle is as follows: each liquidity provider deposits their token into the high fee pool, the low fee pool and pay the gas cost Γ , or withdraw from the market. The **LPs** do not interact with the pool until their liquidity is consumed at some random time $\tilde{\tau}$ and the next cycle begins. [Figure 1](#) illustrates the model timing.

Figure 1: **Model timing**



⁴Empirical properties of gas fees are exhibited in [Lehar and Parlour \(2021\)](#), while [Capponi and Jia \(2021\)](#), consider traders who affect the gas price.

2.1 Equilibrium

First, consider the liquidity traders' decisions. Faced with pool sizes of \mathcal{L}_ℓ and \mathcal{L}_h in the low and high pool respectively, they choose the pool which minimizes their trading costs. Conditional on trading, the small liquidity takers choose the ℓ fee pool. Thus, small traders arrive at the ℓ fee pool at rate θ , and each trade one unit. By assumption the large liquidity taker wants to trade more than the posted liquidity and exhausts both pools. Thus, liquidity is consumed on the low fee pool at rate $\frac{\theta}{\mathcal{L}_\ell} dt$, while liquidity is only consumed on the high fee pool if a large trader arrives. Once each pool is empty, liquidity providers refill it and restart the cycle. Let $d_k, k = \ell, h$ denote the duration of a liquidity cycle on the low and high pool respectively. The the expected duration of a cycle on the low fee pool is

$$d_\ell = e^{-\lambda \frac{\mathcal{L}_\ell}{\theta}} \frac{\mathcal{L}_\ell}{\theta} + \int_0^{\frac{\mathcal{L}_\ell}{\theta}} t \lambda e^{-\lambda t} dt = \frac{1}{\lambda} - \frac{1}{\lambda} e^{-\frac{\mathcal{L}_\ell}{\theta} \lambda}, \quad (5)$$

while the expected duration on a high fee pool is $d_h = \frac{1}{\lambda}$. Thus,

Lemma 1. *The expected duration of a liquidity cycle is shorter in the low fee pool than the high fee pool. Or, $d_\ell < d_h$.*

Lemma 1 is intuitive. A liquidity cycle on the high fee pool only ends with the arrival of a large trader. Conversely, a liquidity cycle on the low fee pool ends either because a large liquidity taker arrived or the cumulative small liquidity trader orders exhaust the pool.

Notice that the low fee pool provides liquidity for a large share of the order flow, as both small and large **LTs** trade there. However, the liquidity providers on that pool earn a low fee per traded unit. Additionally, as the expected cycle duration is shorter, they need to manage their liquidity more often, which leads to larger gas costs per unit of time. Liquidity providers face a choice between the low and high fee pool or not participating in the market.

$$\max \left[\frac{q_i \ell - \Gamma}{d_\ell}, \frac{q_i h - \Gamma}{d_h}, 0 \right]. \quad (6)$$

First, consider the choice between pools. Rearranging Equation 6, liquidity provider i chooses the low-fee pool if and only if

$$(d_h \ell - d_\ell h) q_i > \Gamma (d_h - d_\ell). \quad (7)$$

Equation 7 highlights the tradeoff between the expected fee revenue per unit time and the fixed cost of accessing the market. If $\frac{h}{d_h} > \frac{\ell}{d_\ell}$, then the expected liquidity fee per unit of time on the high fee pool is larger than that of the low fee pool, and the left hand side of Equation 7 is negative. From Lemma 1, the expected duration is higher on the high fee pool and the right hand side is always positive. In this case, all liquidity providers choose the high fee pool. The more natural case

to consider is if $\frac{h}{d_h} < \frac{\ell}{d_\ell}$ so that liquidity providers face a trade-off between a higher liquidity fee per unit of time on the low fee pool against lower gas costs on the high fee pool. Clearly, the larger a liquidity providers' endowment, the more important is the fee revenue. We have

Lemma 2. *For any liquidity pools, $\{\mathcal{L}_\ell, \mathcal{L}_h\}$, for which $\frac{h}{d_h} < \frac{\ell}{d_\ell}$, if a liquidity provider of size q prefers the low fee pool, then any liquidity provider with a larger endowment, $\tilde{q} > q$, also prefers the low fee pool.*

Now consider the choice of participating in the market. An agent only provides liquidity if she is able to break even on the high-fee pool – that is, if her endowment q_i is large enough. The participation constraint follows from equation (6):

$$q_i h - \Gamma \geq 0. \quad (8)$$

Define $\underline{q} = \frac{\Gamma}{h}$. Recall, that the lower bound of the Pareto distribution is 1. Thus, if $\frac{\Gamma}{h} > 1$, liquidity provision is competitive: all **LPs** with $q_i > \underline{q}$ enter the market, and the marginal entrant earns zero expected profit. Conversely, if $\frac{\Gamma}{h} \leq 1$, capital scarcity in the economy implies that liquidity provision is not competitive; all **LPs** enter the market and earn strictly positive profits.

Following Lemma 2, let q_t be the threshold endowment such that all **LPs** with $q_i > q_t$ post liquidity on the low-fee pool and all **LPs** with $q_i \leq q_t$ choose the high-fee pool. This allows us to characterize pool sizes \mathcal{L}_ℓ and \mathcal{L}_h as a function of the endowment for the marginal **LP**:

$$\begin{aligned} \mathcal{L}_\ell &= \int_{q_t}^Q q_i \varphi(q_i) di = \frac{Q}{Q-1} (\log Q - \log q_t) \quad \text{and} \\ \mathcal{L}_h &= \int_{\underline{q}}^{q_t} q_i \varphi(q_i) di = \frac{Q}{Q-1} (\log q_t - \log \underline{q}) \end{aligned} \quad (9)$$

From equations (7) through (9) it follows that the expected profit difference between the low- and high-fee pools can be written as an increasing function of the marginal **LP**'s endowment, that is

$$\begin{aligned} \pi_\ell - \pi_h &= \frac{1}{\lambda d_h d_\ell} \left[\exp\left(-\frac{\lambda}{\theta} \mathcal{L}_\ell\right) \underbrace{(q_i h - \Gamma)}_{>0} - q_i (h - \ell) \right] \\ &= \frac{1}{\lambda d_h d_\ell} \left[q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}} \underbrace{(q_i h - \Gamma)}_{>0} - q_i (h - \ell) \right]. \end{aligned} \quad (10)$$

Proposition 1 characterizes the equilibrium liquidity provision.

Proposition 1. *i. If $\frac{h-\ell}{h} Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} < 1$ and $\frac{\Gamma}{h} < 1 - \frac{h-\ell}{h} Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}}$, then all **LPs** deposit liquidity on the low fee pool.*

ii. If $\Gamma > Q\ell$, then all **LPs** deposit liquidity on the high fee pool.

iii. Otherwise, there exists a unique fragmented equilibrium characterized by marginal trader q_t^* which solves


$$q_t = \Gamma \frac{q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}}}{h \left[q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}} \right] - (h - \ell)} \in [q, Q] \quad (11)$$

such that all **LPs** with $q_i \leq q_t^*$ deposit liquidity in the high fee pool and all **LPs** with $q_i > q_t^*$ choose the low fee pool.

Figure 2 illustrates the equilibrium regions in Proposition 1. If the gas price is low and **LPs** are homogeneous (low Γ and Q), then all liquidity providers choose the low fee pool because managing liquidity is relatively cheap. If more high-endowment **LPs** enter the market (i.e., there is an increase in Q) the more liquidity is posted the low fee pool. Keeping the **LT** arrival rate fixed, a larger pool depletes more slowly and thus the liquidity cycle on the pool becomes longer. As a consequence, the liquidity fee per unit of time on pool L drops and smaller **LPs** switch to the high-fee pool. As a result, liquidity becomes fragmented across the two pools.

Figure 2: Fragmented and single-pool equilibria

This figure plots the existence conditions for a fragmented market equilibrium, as described in Proposition 1, for various values of gas cost (Γ) on the y-axis and liquidity endowment Q on the x-axis. Parameter values: $h = 1$, $\ell = 0.75$, $\lambda = 0.5$ $\theta = 0.66$.



An equilibrium in which all liquidity consolidates on the high fee pool is only sustainable for

very high gas costs $\Gamma > Q\ell$. In this case, none of the **LPs** breaks even on pool L . For intermediate values of gas price, both pools co-exist with positive market share.

Proposition 2 establishes comparative statics for the two pools' liquidity market shares. From equation (9), we can compute the liquidity market share of the low-fee pool at the beginning of each cycle as

$$w_\ell = \frac{\mathcal{L}_\ell}{\mathcal{L}_\ell + \mathcal{L}_h} = \frac{\log Q - \log q_t^*}{\log Q - \log \underline{q}}. \quad (12)$$

Proposition 2. *In equilibrium, the market share of the low fee pool w_ℓ*


- i. decreases in the gas cost (Γ), the arrival rate of large trades (λ), and the fee on pool H (h).*
- ii. increases in the fee on pool L (ℓ) and the arrival rate of small trades (θ)*

The results in Proposition 2 are intuitive. The market share of the low fee pool increases if the fee gap $h - \ell$ is narrower, since this reduces **LP** incentives to switch to the high-fee pool. If the small **LT** arrival rate is large, then liquidity cycles in the low-fee pool are shorter, increasing the revenue per unit of time and consequently the market share of pool L . Conversely, if large trades arrive more often (high λ), then the high fee pool attracts a higher share of incoming order flow and becomes more appealing for liquidity providers.

Figure 3 shows that the market share of the low fee pool (weakly) decreases in the gas cost Γ . A larger gas price increases the costs of active liquidity management, everything else equal, and incentivizes smaller **LPs** to switch from the low fee pool to the high fee pool, since the latter has a lower turnover. For $\Gamma \leq h$, any increase in gas costs leads to a *redistribution* of liquidity from one pool to another; the aggregate liquidity across both pools is constant since all **LPs** participate in the market.

Figure 3: Liquidity shares and gas costs

This figure illustrates the equilibrium liquidity market shares (left panel) and absolute liquidity deposits on the two pools (right panel), as a function of the gas fee Γ . Parameter values: $h = 1$, $\ell = 0.75$, $\lambda = 0.5$, $\theta = 0.66$, $Q = 3$, and $\Gamma = 1$.



If gas prices increase beyond a threshold ($\Gamma > h$), then the aggregate liquidity falls since **LPs** with $q_i < \frac{\Gamma}{h}$ are shut out of the market. Both the low and high fee pool experience a decrease in liquidity deposits. However, the liquidity drop is sharper for the low fee pool which further depresses its market share.

2.2 Model implications and empirical predictions

Prediction 1: The liquidity market share of the low-fee pool decreases in the gas fee Γ .

Prediction 1 follows directly from Proposition 2 and Figure 3. A higher gas price increases the fixed cost of active liquidity management, particularly so for smaller liquidity providers. In response, **LPs** with lower endowments migrate to the high-fee pool where they trade less often.

Prediction 2: **LPs** on the low-fee pool make larger liquidity deposits than **LPs** on the high-fee pool.


Prediction 2 follows from the equilibrium discussion in Section 2.1. Liquidity providers with large token endowments ($q_i > q_t$) deposit them in the low-fee pool since they are better positioned to actively manage liquidity due to economies of scale. **LPs** with lower endowments ($q_i \leq q_t$) either stay out of the market or choose pool H which allows them to offer liquidity in a more passive manner.

Figure 4 illustrates this prediction through a Monte Carlo simulation. We plot the equilibrium

liquidity supply decisions of 100,000 **LPs** with endowments drawn from density (1) and $Q = 3$. The top panel highlights three groups of liquidity providers: low-endowment **LPs** (in green) that are being rationed out of the market due to high gas cost, medium-endowment **LPs** (blue) that deposit liquidity on pool H , and high-endowment **LPs** (orange) that choose the low-fee pool L .

Figure 4: Liquidity supply on fragmented markets

This figure illustrates equilibrium liquidity supply on a fragmented market. We simulate 100,000 liquidity endowments from the truncated Pareto distribution in (1). The top panel illustrates the simulated distributions: **LPs** to the left (right) of the marginal trader q_i^* provide liquidity on pool H (pool L , respectively). The bottom left panel shows the gap in average liquidity deposits for **LPs** participating on the two pools. The bottom right panel shows the gap in the number of **LPs** participating on the two pools. Parameter values: $h = 1$, $\ell = 0.75$, $\lambda = 0.5$, $\theta = 0.66$, $Q = 3$, and $\Gamma = 1$.




Prediction 3: The average trade size is higher on pool H than on pool L . At the same time, trading volume is higher on pool L than on pool H .

Next, Prediction 3 deals with differences between incoming trades on the two liquidity pools. For a wide range of model parameters, incoming order flow on the low fee pool consists of a large number of small trades, and occasional large trades. In contrast, there are few trades on the high fee pool, but they are all relatively large. The model can therefore reconcile two apparently conflicting patterns: one liquidity pool captures most of the trading volume, while the largest trades are executed on the competitor. Figure 5 illustrates the prediction through a Monte Carlo simulation of the model.

Figure 5: Trade volume and trade size on fragmented markets

The left panel illustrates the average trading volume per unit of time in both pools, using 100,000 Monte Carlo simulations with a discrete time increment of $\Delta t = 0.1$. The right panel illustrates the average trade size on a given liquidity pool, conditional on having a non-zero volume in the discrete time interval Δt . If a large liquidity trader arrives, it depletes both pools – i.e., she trades amounts \mathcal{L}_ℓ and \mathcal{L}_h on the low- and high-fee pools, respectively. Parameter values: $h = 1$, $\ell = 0.75$, $\lambda = 0.5$, $\theta = 0.66$, $Q = 3$, and $\Gamma = 1$.



Prediction 4: The average liquidity deposit on both the low- and high fee pool increases with gas costs.

An increase in the gas cost Γ has two effects: first, the **LPs** with the lowest endowments on pool L switch to pool H . As a result, the average deposit on pool L increases. Second, the **LPs** with low endowments on pool H may leave the market. Both channels translate to a higher average deposit on pool H , which experiences an inflow (outflow) of relatively high (low) endowment LP following an increase in gas costs.

The bottom left panel highlights the clientele effect: that is, the average deposit is higher on pool L . Due to the skew of the Pareto distribution, however, there are more **LP** accounts active on pool H than on pool L for a wide range of parameter values.

Prediction 5: **LPs** update liquidity more frequently on the low-fee than on the high-fee pool.

Prediction 5 is a consequence of Lemma 1. Liquidity cycles on pool L are shorter than on pool H , since **LPs** on the low-fee pool trade against both small and large orders, rather than only against large orders on the high-fee pool. Consequently, we expect **LPs** on pool L to actively manage their liquidity positions.

Prediction 6: A larger gas price leads to more frequent liquidity updates on the low-fee pool.

An increase in gas price leads to some **LPs** switching from the low- to the high-fee pool. As a result, liquidity supply on the low-fee pool drops, leading to a shorter cycle as incoming order flow depletes the pool at a faster pace.

3 Data and descriptive statistics

3.1 Sample construction

We obtain data from the Uniswap V3 Subgraph, covering all trades, liquidity deposits (referred to as “mints”), and liquidity withdrawals (referred to as “burns”) on 4,069 Uniswap v3 pools. The data spans from the protocol’s launch on May 4, 2021, up until July 15, 2023. Each entry in our data includes a transaction hash that uniquely identifies each trade and liquidity update on the Ethereum blockchain. Additionally, it provides details such as trade price, direction, and quantity, along with quantities and price ranges for each liquidity update. Moreover, the data also includes wallet addresses associated with initiating each transaction, akin to anonymous trader IDs. The Subgraph data we obtained also provides USD-denominated values for each trade and liquidity mint. We further collect daily pool snapshots from the Uniswap V3 Subgraph, including the end-of-day pool size in Ether and US Dollars, and summary price information (e.g., open, high, low, and closing prices for each pool).

To enhance our dataset, we combine the Subgraph data with public Ethereum data available on [Google Big Query](#) to obtain the position of each transaction in its block, as well as the gas price limit set by the trader and the amount of gas used.

There are no restrictions to list a token pair on Uniswap. Some pools might therefore be used for experiments, or they might include untrustworthy tokens. Following [Lehar and Parlour \(2021\)](#), we remove pools that are either very small or that are not attracting an economically meaningful trading volume. We retain liquidity pools that fulfill the following four criteria: (i) have at least one interaction in more than 100 days in the sample, (ii) have more than 500 liquidity interactions throughout the sample, (iii) have an average daily liquidity balance in excess of US\$100,000, and (iv) capture more than 1% of trading volume for a particular asset pair. We exclude burn events with zero liquidity withdrawal in both base and quote assets, as traders use them solely to collect fees without altering their liquidity position.

These basic screens give us a baseline sample of 274 liquidity pools covering 242 asset pairs, with combined daily dollar volume of \$1.12 billion and total value locked (i.e., aggregate liquidity supply) of \$2.53 billion as of July 15, 2023. We capture 24,202,803 interactions with liquidity pool smart contracts (accounting for 86.04% of the entire universe of trades and liquidity updates). Trading and liquidity provision on Uniswap is heavily concentrated: the five largest pairs (USDC-WETH,

WETH-USDT, USDC-USDT, WBTC-WETH, and DAI-USDC) account on average for 86% of trading volume and 63% of supplied liquidity.⁵

3.2 Liquidity fragmentation patterns

For 32 out of the 242 asset pairs in our baseline sample, liquidity supply is fragmented across two pools with different fees – either with 1 and 5 bps fees (5 pairs), 5 and 30 bps fees (6 pairs), or 30 and 100 bps fees (21 pairs).⁶ Despite being fewer in number, fragmented pairs are economically important: they account on average for 95% of the capital committed to Uniswap v3 and for 93% of its dollar trading volume. All major token pairs such as WETH-USDC, WETH-USDT, or WBTC-WETH trade on fragmented pools.

For each fragmented liquidity pair, we label the *low* and the *high* fee liquidity pool to facilitate analysis across assets. For example, the low and high liquidity fees for USDC-WETH are 5 and 30 bps, respectively, but only 1 and 5 bps for a lower volatility pair such as USDC-USDT. We refer to non-fragmented pools as *single* (i.e., the unique pool for an asset pair).

We aggregate all interactions with Uniswap smart contracts into a panel across days and liquidity pools. To compute the end-of-day pool size, we account for all changes in token balances, across all price ranges. There are three possible interactions: A deposit or “mint” adds tokens to the pool, a withdrawal or “burn” removes tokens, whereas a trade or “swap” adds one token and removes the other. We track these changes across to obtain daily variation in the quantity of tokens on each pool. We obtain dollar values for the end-of-day liquidity pool sizes, intraday trade volumes, and liquidity events from the Uniswap V3 Subgraph. To determine a token’s price in dollars, the Subgraph searches for the most liquid path on Uniswap pools to establish the token’s price in Ether and subsequently converts the Ether price to US dollars.

Table 1 reports summary statistics across pools with different fee levels. High-fee pools attract on average 58% of total liquidity supply, significantly more than their low-fee counterparts (\$46.50 million and \$33.78 million, respectively), but only capture 20.74% percent of the trading volume (computed as $\frac{8,071.24}{(8,071.24+30,848.79)}$ from the first column of Table 1). Consistent with our theoretical predictions, low-fee pools attract five times as many trades as high-fee competitors (610 versus 110 average trade count per day). At the same time, the average trade on a high-fee pool is twice as large (\$14,490) than on a low-fee pool (\$6,340).

The distribution of mint sizes is heavily skewed to the right, with 6.6% of deposits exceeding

⁵WETH and WBTC stand for “wrapped” Bitcoin and Ether. Plain vanilla Bitcoin and Ether are not compliant with the ERC-20 standard for tokens, and therefore cannot be directly used on decentralized exchanges’ smart contracts. USDC (USD Coin), USDT (Tether), and DAI are stablecoins meant to closely track the US dollar.

⁶In some cases, more than two pools are created for a pair – e.g., for USDC-WETH there are four pools with 1, 5, 30, and 100 bps liquidity fees. In all cases however, two pools heavily dominate the others: As described in Section 3.1 we filter out small pools with less than 1% volume share or less than \$100,000 liquidity deposits.

\$1 million. There are large differences across pools – the median **LP** deposit on the low-fee pool is \$15,680, twice as much as the median deposit on the high-fee pool (\$7,430). At the same time, the number of liquidity providers on high-fee pools is 51% higher than on low-fee pools (10.08 unique addresses per day on high-fee pools versus only 6.68 unique address on high-fee pools).

Table 1: Descriptive statistics

This table reports descriptive statistics for variables used in the empirical analysis. *Pool size* is defined as the total value locked in the pool’s smart contract at the end of each day. We compute the balance on day t as follows: we take the balance at day $t - 1$ and add (subtract) liquidity deposits (withdrawals) on day t , as well as accounting for token balance changes due to trades. The liquidity balance on the first day of the pool is taken to be zero. End of day balances are finally converted to US dollars. *Daily volume* is computed as the sum of US dollar volume for all trades in a given pool and day. *Liquidity share* (*Volume share*) is computed as the ratio between a pool size (trading volume) for a given fee level and the aggregate size of all pools (trading volumes) for the same pair in a given day. *Trade size* and *Mint size* are the median trade and liquidity deposit size on a given pool and day, denominated in US dollars. *Trade count* represents the number of trades in a given pool and day. *LP wallets* counts the unique number of wallet addresses interacting with a given pool in a day. The *liquidity yield* is computed as the ratio between the daily trading volume and end-of-day TVL, multiplied by the fee tier. The *price range* for every mint is computed as the difference between the top and bottom of the range, normalized by the range midpoint – a measure that naturally lies between zero and two. The *impermanent loss* is computed as in [Heimbach, Schertenleib, and Wattenhofer \(2022\)](#) for a position in the range of 95% to 105% of the current pool price, with a forward-looking horizon of one hour. Finally, *mint-to-burn* and *burn-to-mint* times are defined as the time between a mint (burn) and a subsequent burn (mint) by the same address in the same pool, measured in hours. *Mint-to-burn* and *burn-to-mint* are recorded on the day of the final interaction with the pool.

Statistic	Pool fee	Mean	Median	St. Dev.	Pctl(25)	Pctl(75)	N
Pool size (\$M)	Low	33.78	2.05	96.91	0.30	14.12	20,151
	High	46.50	3.85	95.73	1.43	27.51	20,151
	Single	3.89	0.84	13.56	0.26	2.62	130,767
Liquidity share (%)	Low	39.52	35.52	32.53	7.37	72.16	20,151
	High	60.48	64.48	32.53	27.84	92.63	20,151
Daily volume (\$000)	Low	30,848.79	619.77	118,908.80	6.18	5,697.30	20,151
	High	8,071.24	114.96	36,777.38	7.83	1,882.12	20,151
	Single	915.73	36.07	6,059.78	1.93	277.00	130,767
Volume share	Low	66.51	88.38	38.50	29.43	98.48	18,001
	High	42.20	23.83	41.18	3.19	95.03	18,058
Trade size (\$000)	Low	6.34	2.20	13.36	0.61	6.03	18,001
	High	14.49	2.76	33.19	0.82	10.48	18,060
	Single	4.12	1.32	11.03	0.45	3.79	113,362
Mint size (\$000)	Low	820.84	15.68	13,114.83	3.78	58.98	10,640
	High	1,001.10	7.43	13,807.10	1.55	30.52	10,370
	Single	96.97	6.93	622.12	1.42	30.39	45,300
Trade count	Low	610.61	95	1,518.52	12	414	20,151
	High	110.59	26	490.29	8	89	20,151
	Single	63.94	19	194.03	4	55	130,767
LP wallets	Low	6.68	1	16.01	0	6	20,151
	High	10.08	1	37.79	0	5	20,151
	Single	1.57	1.17	1.19	1.00	1.85	55,580
Liquidity yield (bps)	Low	11.72	2.58	56.31	0.16	9.08	20,122
	High	9.69	1.65	51.44	0.15	6.40	20,130
	Single	17.90	1.94	90.18	0.18	8.58	130,433
Price range	Low	0.39	0.30	0.37	0.13	0.56	11,866
	High	0.61	0.54	0.44	0.32	0.84	12,195
	Single	0.68	0.58	0.52	0.27	1.02	55,580
Impermanent loss (bps)	Low	8.46	1.84	27.88	0.06	7.23	20,118
	High	7.37	1.33	27.21	0.05	5.93	20,132
	Single	17.20	2.44	71.34	0.17	11.37	130,340
Mint-to-burn (hrs)	Low	450.40	59.82	1,341.67	19.70	243.83	10,186
	High	952.14	165.61	2,076.42	39.66	711.67	9,979
	Single	760.26	126.64	1,778.62	27.01	563.50	39,735
Burn-to-mint (hrs)	Low	105.29	0.20	521.62	0.08	5.63	8,279
	High	224.26	0.32	941.31	0.10	27.78	7,289
	Single	177.74	0.23	803.40	0.07	20.64	27,477

One concern with measuring average mint size is just-in-time liquidity provision (JIT). JIT liquidity providers submit very large and short-lived deposits to the pool to dilute competitors on an incoming large trade; they immediately withdraw the balance in the same block after executing the trade. In our sample, JIT liquidity provision is not economically significant, accounting for less than 1% of aggregate trading volume. However, it has the potential to skew mint sizes to the right, particularly in low-fee pools, without providing liquidity to the market at large. We address this issue by (i) filtering out JIT mints using the algorithm in Appendix D and (ii) taking the median liquidity mint size at day-pool level rather than the mean.

Further, we follow [Augustin, Chen-Zhang, and Shin \(2022\)](#) to compute the daily liquidity fee yield as the product between pool’s fee tier and the ratio between trading volume and the lagged total value locked (TVL). That is,

$$\text{Liquidity yield} = \text{liquidity fee}_i \times \frac{\text{Volume}_{i,t}}{\text{TVL}_{i,t-1}}, \quad (13)$$


for pool i and day t . The average daily yield is slightly higher on low-fee pools, at 11.72 basis points, compared to 9.69 basis points on high-fee pools.

A salient observation in Table 1 is that non-fragmented pairs (“single” pools) are significantly smaller – on average less than 10% of the pool size and trading volume of fragmented pairs. Average trade and mint sizes are correspondingly lower as well. The evidence suggests that pairs for which there is significant trading interest, and therefore potentially a broader cross-section of potential liquidity providers, are more likely to become fragmented.


Figure 6 plots the distributions of our empirical measures across low- and high-fee liquidity pools. It suggest a sharp segmentation of liquidity supply and trading across pools. High-fee pools attract smaller liquidity providers by mint size, and end up with a larger *aggregate* size than their low-fee counterparts. Trading volume is similarly segmented: most small value trades are executed on the cheaper low-fee pools, making up the majority of daily volume for a given pair. High-value trades, of which there are fewer, are more likely to (also) execute on high-fee pools.

Figure 6: Liquidity supply on decentralized exchanges


This figure plots the empirical distributions of variables in the pool-day panel, across low and high fee pools (for fragmented pairs) as well as single pools in pairs that are not fragmented. In each box plot, the median is marked as a vertical line; the box extends to the quartiles of the data set, whereas the whiskers extend to an additional 1.5 times the inter-quartile range.



(a) Pool size and trading volume



(b) Average liquidity mint and trade size



(c) Number of liquidity providers and trades

Our theoretical framework in Section 2 implies that liquidity suppliers manage their positions more actively in the low- than the high-fee pool. Figure 7 provides suggestive evidence for liquidity cycles of different lengths in the cross-section of pools. Liquidity on decentralized exchanges is significantly more passive than on traditional equity markets. That is, liquidity providers do not often manage their positions at high frequencies. The median time from a mint (deposit) to a subsequent burn (withdrawal) from the same wallet on the same pool ranges from 59.82 hours, or 2.49 days, on low-fee pools to 165.61 hours, or 6.9 days on high-fee pools.

When do LPs re-balance their positions? In 53% of cases, liquidity providers only withdraw tokens from the pool when their position exits the price range that allows them to collect fees. Concretely, LPs specified price range for liquidity provision does not straddle the most recent reference price of the pool. The scenario mirrors a limit order market where a liquidity provider's


outstanding limit orders are deep in the book, such that she doesn't stand to earn the spread on the marginal incoming trade. In this case, a rational market maker might want to cancel their outstanding order and place a new one at the top of the book. This is exactly the pattern we observe on Uniswap: the subsequent mint following a burn straddles the new price 77% of the time – **LPs** reposition their liquidity around the current prices to keep earning fees on incoming trades. Moreover, re-balancing is swift – the median time between a burn and a subsequent mint is just 12 minutes (0.20 hours).

The empirical pattern in Figure 7 echoes liquidity cycles as described in Section 2. Liquidity deposit tokens in Uniswap pools and then patiently wait for days until incoming order flow exhausts their position (i.e., posted liquidity no longer earns fees). Once this happens, **LPs** quickly re-balance their position in a matter of minutes – by removing stale liquidity and adding a new position around the current price. The cycle is longer on high-fee pools for which trading volume is lower and liquidity takes longer to deplete.

Importantly, **LPs** do not seem to “race” to update liquidity upon information arrival as in [Budish, Cramton, and Shim \(2015\)](#). First, they very rarely manage their position intraday. Second, **LPs** on Uniswap typically do not remove in-range liquidity that stands to trade first against incoming order flow and therefore bears the highest adverse selection risk. Our results are consistent with [Capponi and Jia \(2021\)](#) who theoretically argue that **LPs** have low incentives to manage liquidity on news arrival, as well as with [Capponi, Jia, and Yu \(2022\)](#) who find no evidence of traders racing to trade on information on Uniswap v2.

Figure 7: Liquidity cycles on high- and low-fee pools

The top panel plots the distribution of liquidity cycle times from mint to subsequent burn (left) and from burn to subsequent mint (right) for the same LP wallet address in the same pool. In each box plot, the median is marked as a vertical line; the box extends to the quartiles of the data set, whereas the whiskers extend to an additional 1.5 times the inter-quartile range. The bottom panel plots the probability that the LP position is out of range and therefore does not earn fees. A position is considered to be “out of range” when the minimum and maximum prices at which the LP is willing to provide liquidity do not straddle the current price on the pool. We plot the probability separately for low- and high- fees, as well as conditional on whether the event is a burn (liquidity withdrawal) or mint (liquidity deposit).




Measuring gas prices. Each interaction with smart contracts on the Ethereum blockchain requires computational resources, measured in units of “gas.” Upon submitting a mint or burn transaction to the decentralized exchange, each liquidity provider specifies their willingness to pay per unit of gas, that is they bid a “gas price.” Traders are likely to bid higher prices for more complex transactions or if they require a faster execution. To generate a conservative daily benchmark for the gas price, we compute the average of the lowest 1000 user gas bids for mint and burn interactions on day t , across all liquidity pools in the benchmark sample. These transactions are more likely to be plain vanilla deposits or withdrawals of liquidity, capturing the cost of a simple interaction with the decentralized exchange smart contract.

Figure 8 showcases the significant fluctuation in gas costs for Uniswap liquidity transactions over time. Gas costs denominated in USD are influenced by two primary factors: network congestion, which leads to variations in gas prices measured in Ether, and the fluctuation of Ether’s value relative to the US dollar. On a monthly average, gas costs peaked at above US\$100 in November

2021 and have since plummeted to around US\$6 from the second half of 2022, albeit with occasional spikes.

Figure 8: Gas costs for Uniswap v3 mint/burn transactions

The figure illustrates the daily average gas cost on mint/burn transactions in Uniswap v3 pools. The gas cost is computed as the average of the lowest 1000 user gas bids for mint and burn interactions on each day, across all liquidity pools in the benchmark sample.



4 Empirical results

4.1 Liquidity supply on high- and low-fee pools

To formally test the model predictions and quantify the differences in liquidity supply across fragmented pools, we build a panel data set for the 32 fragmented pairs in our sample where the unit of observation is pool-day. We estimate linear regressions of liquidity and volume measures on liquidity fees and gas costs:

$$y_{ijt} = \alpha + \beta_0 d_{low-fee, ij} + \beta_1 \text{GasPrice}_{jt} + \beta_2 \text{GasPrice}_{jt} \times d_{low-fee, ij} + \sum \beta_k \text{Controls}_{ijt} + \theta_j + \delta_w + \varepsilon_{ijt}, \quad (14)$$

where y is a variable of interest, i indexes liquidity pools, j runs over asset pairs, and t and w indicates days and weeks, respectively. The dummy $d_{low-fee, ij}$ takes the value one for the pool with the lowest fee in pair j and zero else.

Further, our set of controls includes pair and week fixed effects, the log aggregate trading volume and log liquidity supply (i.e., total value locked) for day t across all pools i . Volume and liquidity are measured in US dollars. We also control for daily return volatility, computed as the range between

the daily high and low prices for a given pair j (following Alizadeh, Brandt, and Diebold, 2002):

$$\text{Volatility}_{jt} = \frac{1}{2\sqrt{\log 2}} \log \left(\frac{\text{High}_{jt}}{\text{Low}_{jt}} \right). \quad (15)$$

To measure volatility for fragmented pairs that actively trade in multiple pools, we select the pool with the highest trading volume for a given day.

Consistent with Figure 6, most of the capital deployed to provide liquidity for a given pair is locked in high-fee pools. At the same time, low-fee pools attract much larger trading volume. Models (1) and (5) show that the average low-fee pool attracts 39.5% of liquidity supply for the average pair (that is, equal to $(100-20.92)/2$) while it executes 62% (i.e., $(100+24.62)/2$) of the total trading volume. At a first glance, it would seem that a majority of capital on decentralized exchanges is inefficiently deployed in pools with low execution probability. We will show that, in line with our model, the difference is driven by heterogeneous liquidity cycles across pools, leading to the formation of **LP** clienteles.

The regression results in Table 2 support Prediction 1, stating that market share differences between pools are linked to variation in fixed transaction costs on the blockchain. A one-standard deviation increase in gas prices leads to a 4.63 percentage point increase in the high-fee liquidity share. The results suggests that blockchain transaction costs have an economically meaningful and statistically significant impact on liquidity fragmentation. In line with the theoretical model in Section 2, a jump in gas prices leads to a reshuffling of liquidity supply from low- to high-fee pools.

Evidence suggests that a higher gas price leads to a 6.52% lower volume share for the low-fee pool. This outcome is natural, as the incoming order flow is optimally routed to the high-fee pool, following the liquidity providers.

Table 2: Liquidity pool market shares and gas prices

This table reports the coefficients of the following regression:

$$\text{MarketShare}_{ijt} = \alpha + \beta_0 d_{\text{low-fee}, ij} + \beta_1 \text{GasPrice}_{jt} + \beta_2 \text{GasPrice}_{jt} \times d_{\text{low-fee}, ij} + \sum \beta_k \text{Controls}_{ijt} + \theta_j + \varepsilon_{ijt}$$

where the dependent variable is the liquidity or trading volume market share for pool i in asset pair j on day t . $d_{\text{low-fee}, ij}$ is a dummy that takes the value one for the pool with the lowest fee in pair j and zero else. GasPrice_{jt} is the average of the lowest 100 bids on liquidity provision events across all pairs on day t , standardized to have a zero mean and unit variance. Volume is the natural logarithm of the sum of all swap amounts on day t , expressed in thousands of US dollars. $\text{Total value locked}$ is the natural logarithm of the total value locked on Uniswap v3 pools on day t , expressed in millions of dollars. Volatility is computed as the daily range between high and low prices on the most active pool for a given pair. All regressions include pair and week fixed-effects. Robust standard errors in parenthesis are clustered by week and ***, **, and * denote the statistical significance at the 1, 5, and 10% level, respectively. The sample period is from May 4, 2021 to July 15, 2023.

	Liquidity market share (%)				Volume market share (%)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$d_{\text{low-fee}}$	-20.92*** (-27.42)	-20.92*** (-27.41)	-20.92*** (-27.42)	-20.94*** (-23.95)	24.62*** (20.55)	24.63*** (20.56)	24.62*** (20.55)	24.71*** (18.54)
Gas price $\times d_{\text{low-fee}}$	-4.63*** (-7.32)	-4.62*** (-7.32)	-4.63*** (-7.32)		-6.52*** (-5.92)	-6.52*** (-5.92)	-6.52*** (-5.92)	
Gas price	2.31*** (7.32)	2.31*** (7.32)	2.31*** (7.32)		3.63*** (7.33)	3.61*** (7.30)	3.61*** (7.26)	
Volume	0.00 (0.65)	0.00 (1.33)	0.00 (0.65)	0.00 (0.66)	-0.19** (-2.54)	-0.20** (-2.61)	-0.19** (-2.50)	-0.12 (-1.56)
Total value locked	-0.00 (-0.58)	-0.00 (-0.06)		-0.00 (-0.64)	0.58 (1.44)	0.58 (1.44)		0.44 (1.10)
Volatility	-0.29 (-0.90)		-0.29 (-0.90)	-0.28 (-0.82)	-1.15*** (-2.74)		-1.15*** (-2.74)	-1.13** (-2.56)
Constant	60.45*** (158.00)	60.46*** (158.46)	60.45*** (158.00)	60.46*** (137.54)	41.96*** (69.99)	41.99*** (70.22)	41.96*** (70.02)	41.96*** (62.81)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,288	40,288	40,288	40,288	36,059	36,059	36,059	36,059
R-squared	0.10	0.10	0.10	0.09	0.13	0.13	0.13	0.12

Robust t-statistics in parentheses. Standard errors are clustered at week level. *** p<0.01, ** p<0.05, * p<0.1

What drives the market share gap across fragmented pools? In Table 3 we document stark differences between the characteristics of individual orders supplying or demanding liquidity on pools with low and high fees. On the liquidity supply side, model (1) in Table 3 shows that the average liquidity mint is 107.5% larger on low-fee pools, which supports Prediction 2 of the model.⁷ At the same time, there are 3.40 fewer unique wallets (Model5) providing liquidity on the low-fee pool – that is, a 34% relative difference between high- and low-fee pools.

⁷Since all dependent variables are measured in natural logs, the marginal impact of a dummy coefficient β is computed $(e^\beta - 1) \times 100$ percent.

Table 3: Fragmentation and order flow characteristics

This table reports the coefficients of the following regression:

$$y_{ijt} = \alpha + \beta_0 d_{low-fee, ij} + \beta_1 GasPrice_{jt} d_{low-fee, ij} + \beta_2 GasPrice_{jt} \times d_{high-fee, ij} + \sum \beta_k Controls_{ijt} + \theta_j + \varepsilon_{ijt}$$

where the dependent variable y_{ijt} can be (i) the log median mint size, (ii) the log median trade size, (iii) the log trading volume, (iv) the log trade count $\log(1 + \#trades)$, (v) count of unique **LP** wallets interacting with a pool in a given day, (vi) the liquidity yield in bps for pool i in asset j on day t , computed as in equation (13), and (vii) the average liquidity mint price range for pool i in asset j on day t . Price range is computed as the difference between the top and bottom of the range, normalized by the range midpoint – a measure that naturally lies between zero and two. $d_{low-fee, ij}$ is a dummy that takes the value one for the pool with the lowest fee in pair j and zero else. $d_{high-fee, ij}$ is defined as $1 - d_{low-fee, ij}$. $GasPrice_{jt}$ is the average of the lowest 100 bids on liquidity provision events across all pairs on day t , standardized to have a zero mean and unit variance. $Volume$ is the natural logarithm of the sum of all swap amounts on day t , expressed in thousands of US dollars. $Total\ value\ locked$ is the natural logarithm of the total value locked on Uniswap v3 pools on day t , expressed in millions of dollars. $Volatility$ is computed as the daily range between high and low prices on the most active pool for a given pair. All regressions include pair and week fixed-effects. Robust standard errors in parenthesis are clustered by week and ***, **, and * denote the statistical significance at the 1, 5, and 10% level, respectively. The sample period is from May 4, 2021 to July 15, 2023.

	Mint size (1)	Trade size (2)	Volume (3)	# Trades (4)	# Wallets (5)	Liquidity yield (6)	Price range (7)
$d_{low-fee}$	0.73*** (12.27)	-0.30*** (-10.05)	0.89*** (14.23)	1.02*** (32.95)	-3.40*** (-5.00)	2.03*** (3.60)	-0.18*** (-41.84)
Gas price $\times d_{low-fee}$	0.37*** (4.96)	0.08*** (3.75)	-0.03 (-0.95)	-0.22*** (-7.29)	-3.00*** (-3.43)	3.57** (2.30)	-0.00 (-0.47)
Gas price $\times d_{high-fee}$	0.58*** (7.52)	0.17*** (8.81)	0.24*** (5.95)	0.07** (2.46)	-2.89*** (-3.15)	5.57*** (2.83)	-0.03*** (-4.65)
Volume	0.37*** (8.68)	0.16*** (21.38)	0.43*** (15.27)	0.20*** (13.85)	1.22*** (6.56)	1.01 (0.81)	-0.01** (-2.56)
Total value locked	-0.16 (-1.30)	0.11*** (3.54)	0.23** (1.99)	-0.01 (-0.18)	-1.86 (-0.99)	-13.42 (-1.09)	-0.02 (-0.99)
Volatility	-0.04 (-1.11)	-0.01 (-1.34)	-0.07 (-1.38)	0.01 (0.88)	-0.09 (-1.03)	1.18** (2.21)	0.02*** (3.98)
Constant	1.88*** (58.27)	1.64*** (111.47)	5.27*** (168.58)	3.26*** (209.84)	10.12*** (28.65)	10.01*** (26.04)	0.59*** (184.91)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	21,000	36,059	36,059	40,288	40,288	40,252	24,058
R-squared	0.26	0.53	0.55	0.52	0.37	0.09	0.42

Robust t-statistics in parentheses. Standard errors are clustered at week level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

On the liquidity demand side, trades on the low-fee pool are 25.91% smaller (Model 2), consistent with Prediction 3. However, the low-fee pool executes almost three times the number of trades (i.e., trade count is 177% higher from Model 4) and has 143% higher volume than the high-fee pool (Model 3). On average, liquidity providers on low-fee pools earn 2.03 basis points higher revenue than their counterparts on high-fee pools (Model 6), indicating significant positive returns resulting from economies of scale.

Our findings (Model 7) indicate that liquidity providers on low-fee pools select price ranges that are 30% ($=0.18/0.59$) narrower when minting liquidity compared to those on high-fee pools. This pattern aligns with the capability of large LPs to adjust their liquidity positions frequently, enabling more efficient capital concentration. Similarly, [Caparros, Chaudhary, and Klein \(2023\)](#) report a higher concentration of liquidity in pools on alternative blockchains like Polygon, known for lower transaction costs than Ethereum.

The results point to an asymmetric match between liquidity supply and demand across pools. On low-fee pools, a few **LPs** provide large chunks of liquidity for the vast majority of incoming small trades. Conversely, on high-fee pools there is a sizeable mass of small liquidity providers that mostly trade against a few large incoming trades.

How does variation in fixed transaction costs impact the gap between individual order size across pools? We find that increasing the gas price by one standard deviation leads to higher liquidity deposits on both the low- and the high-fee pools (14.2% and 30.1% higher, respectively).⁸ The result supports Prediction 4 of the model. Our theoretical framework implies that a larger gas price leads to some (marginal) **LPs** switching from the low- to the high-fee pool. The switching **LPs** have low capital endowments relative to their low-fee pool peers, but higher than **LPs** on the high-fee pool. Therefore, the gas-driven reshuffle of liquidity leads to a higher average endowment on both high- and low-fee pools. Consistent with the model, a higher gas price leads to fewer active liquidity providers, particularly on low-fee pools. Specifically, a one-standard increase in gas costs leads to a significant decrease in the number of **LP** wallets interacting daily with low- and high-fee pools, respectively (Model 5).

While a higher gas price is correlated with a shift in liquidity supply, it has a muted impact on liquidity demand on low-fee pools. A higher gas cost is associated with 6% larger trades (Model 2), likely as traders aim to achieve better economies of scale. At the same time, the number of trades on the low-fee pool drops by 5.1% (Model 4) – since small traders might be driven out of the market. The net of gas prices effect on aggregate volume on the low-fee pool is small and not statistically significant (Model 3). The result matches our model assumption that the aggregate order flow on low-fee pool is not sensitive to gas prices.⁹

On the high-fee pool, a higher gas price is also associated with a higher trade size, but a much smaller relative increase in the number of trades. The implication is that large-size traders who typically route orders to the high-fee pool are unlikely to be deterred by an increase in fixed costs, while they do adjust quantities to achieve better economies of scale. In the context of our model, the large **LT** arrival rate λ does not depend on the gas price Γ .

⁸The relative effects are computed as $0.37/(1.88+0.73) = 13.8\%$ for low pools and $0.58/1.88 = 30.85\%$ for high-fee pools, respectively.

⁹Formally, one could extend the model to assume that small **LTs** arrive at the market at rate $\tilde{\theta}(\Gamma) dt$ and demand $f(\Gamma)$ units each, where $\tilde{\theta}(\Gamma)$ decreases in Γ and $f(\Gamma)$ increases in Γ such that $\tilde{\theta}(\Gamma) f(\Gamma) = \theta$.

In Table 4, we shift the analysis from individual orders to aggregate daily liquidity flows to Uniswap pools. We find that higher gas prices lead to a decrease in liquidity inflows, but only on the low fee pools. A one standard deviation increase in gas prices leads to a 29.5% drop in new liquidity deposits by volume (Model 1) and an 6.02% drop in probability of having at least one mint (Model 4) on the low-fee pool. However, the slow-down in liquidity inflows is less evident in high fee pools. While an increase in gas prices reduce the probability of liquidity inflows by 1.43%, it actually leads to a 11.6% increase in the daily dollar inflow to the pool. Together with the result in Table 3 that the size of individual mints increases with gas prices, our evidence is consistent with the model implication that higher fixed transaction costs change the composition of liquidity supply on the high-fee pool, with small **LP** being substituted by larger **LPs** switching over from the low-fee pool.

Table 4: Liquidity flows and gas costs on fragmented pools

This table reports the coefficients of the following regression:

$$y_{ijt} = \alpha + \beta_0 d_{low-fee, ij} + \beta_1 GasPrice_{jt} d_{low-fee, ij} + \beta_2 GasPrice_{jt} \times d_{high-fee, ij} + \sum \beta_k Controls_{ijt} + \theta_j + \varepsilon_{ijt}$$

where the dependent variable y_{ijt} can be (i) the aggregate dollar value of mints (in logs), or (vi) a dummy variable taking value one hundred if there is at least one mint on liquidity pool i in asset j on day t . $d_{low-fee, ij}$ is a dummy that takes the value one for the pool with the lowest fee in pair j and zero else. $d_{high-fee, ij}$ is defined as $1 - d_{low-fee, ij}$. $GasPrice_{jt}$ is the average of the lowest 100 bids on liquidity provision events across all pairs on day t , standardized to have a zero mean and unit variance. $Volume$ is the natural logarithm of the sum of all swap amounts on day t , expressed in thousands of US dollars. $Total\ value\ locked$ is the natural logarithm of the total value locked on Uniswap v3 pools on day t , expressed in millions of dollars. $Volatility$ is computed as the daily range between high and low prices on the most active pool for a given pair. All regressions include pair and week fixed-effects. Robust standard errors in parenthesis are clustered by week and ***, **, and * denote the statistical significance at the 1, 5, and 10% level, respectively. The sample period is from May 4, 2021 to July 15, 2023.

	Daily mints (log US\$)			Prob (at least one mint)		
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{low-fee}$	0.43*** (6.07)	0.43*** (6.07)	0.43*** (6.07)	1.38* (1.71)	1.37* (1.71)	1.38* (1.71)
Gas price $\times d_{low-fee}$	-0.35*** (-8.50)	-0.35*** (-8.50)	-0.46*** (-7.14)	-6.02*** (-9.13)	-6.01*** (-9.13)	-4.58*** (-6.76)
Gas price $\times d_{high-fee}$	0.11** (2.15)	0.11** (2.15)		-1.43** (-2.57)	-1.43** (-2.57)	
Volume	0.26*** (14.78)	0.26*** (14.77)	0.26*** (14.78)	0.96*** (3.93)	0.96*** (3.93)	0.96*** (3.93)
Total value locked	-0.07 (-0.78)	-0.07 (-0.78)	-0.07 (-0.78)	1.47 (1.01)	1.47 (1.00)	1.47 (1.01)
Volatility	-0.01 (-0.68)		-0.01 (-0.68)	0.26 (0.59)		0.26 (0.59)
Gas price			0.11** (2.15)			-1.43** (-2.57)
Constant	2.61*** (73.46)	2.61*** (73.49)	2.61*** (73.46)	51.44*** (126.67)	51.43*** (127.28)	51.44*** (126.67)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,288	40,288	40,288	40,288	40,288	40,288
R-squared	0.47	0.47	0.47	0.28	0.28	0.28

Robust t-statistics in parentheses. Standard errors are clustered at week level.

*** p<0.01, ** p<0.05, * p<0.1

4.2 Liquidity cycles on high- and low-fee pools

Next, we test Predictions 5 and 6 on the duration of liquidity cycles on fragmented pools. Since the descriptive statistics in Table 1 suggest that LPs manage their positions over multiple days, we cannot accurately measure liquidity cycles in a pool-day panel. Instead, we use intraday data on liquidity events (either mints or burns) to measure the duration between two consecutive opposite-sign interactions by the same Ethereum wallet with a liquidity pool: either a mint followed by a burn, or vice-versa.

To complement our previous analysis, we additionally control for whether each liquidity position is out of range (i.e., the price range set by the LP does not straddle the current price and therefore the LP does not earn fees). We further introduce wallet fixed effects to soak up variation in reaction times across traders.

Table 5 presents the results. Liquidity updates on decentralized exchanges are very infrequent, as times elapsed between consecutive interactions are measured in days or even weeks. In line with Prediction 5, we find evidence for shorter liquidity cycles on low-fee pools. The average time between consecutive mint and burn orders is 20.06% shorter on the low-fee pool (from Model 1, the relative difference is 99.74 hours/497.18 hours).

Liquidity cycles are in part driven by fixed Blockchain transaction costs. A one standard deviation increase in gas prices speeds up the liquidity cycle on low-fee pools by a further 15.80 hours, a result which supports the intuition behind Prediction 6. When the gas price spikes, the liquidity supply on the low-fee pool decreases at a higher rate than the liquidity demand. As a result, liquidity deposits deplete faster (i.e., they move outside the fee-earning range), triggering the need for more frequent updates. At the same time, a higher gas fee speeds up the liquidity cycle on the high-fee pool as well. The result is consistent with small liquidity providers on the high-fee pool being crowded out by the high fixed costs, leading to a lower liquidity supply.

Table 5: Liquidity cycles on fragmented pools

This table reports the coefficients of the following regression:

$$y_{ijtk} = \alpha + \beta_0 d_{low-fee, ij} + \beta_1 GasPrice_{jt} d_{low-fee, ij} + \beta_2 GasPrice_{jt} \times d_{high-fee, ij} + \sum \beta_k Controls_{ijt} + \theta_j + \varepsilon_{ijt}$$

where the dependent variable y_{ijtk} can be (i) the mint-to-burn time, (ii) the burn-to-mint time, measured in hours, for a transaction initiated by wallet k on day t and pool i trading asset j . The mint-to-burn and burn-to-mint times are computed for consecutive interactions of the same wallet address with the liquidity pool. $d_{low-fee, ij}$ is a dummy that takes the value one for the pool with the lowest fee in pair j and zero else. $d_{high-fee, ij}$ is defined as $1 - d_{low-fee, ij}$. $GasPrice_{jt}$ is the average of the lowest 100 bids on liquidity provision events across all pairs on day t , standardized to have a zero mean and unit variance. $Volume$ is the natural logarithm of the sum of all swap amounts on day t , expressed in thousands of US dollars. $Total\ value\ locked$ is the natural logarithm of the total value locked on Uniswap v3 pools on day t , expressed in millions of dollars. $Volatility$ is computed as the daily range between high and low prices on the most active pool for a given pair. $Position\ out-of-range$ is a dummy taking value one if the position being burned or minted is out of range, that is if the price range selected by the LP does not straddle the current pool price. All variables are measured as of the time of the second leg of the cycle (i.e., the burn of a mint-burn cycle). All regressions include pair, week, and trader wallet fixed-effects. Robust standard errors in parenthesis are clustered by day and ***, **, and * denote the statistical significance at the 1, 5, and 10% level, respectively. The sample period is from May 4, 2021 to July 15, 2023.

	Mint-burn time			Burn-mint time		
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{low-fee}$	-99.74*** (-8.86)	-100.17*** (-8.94)	-104.09*** (-9.22)	-157.95*** (-10.59)	-159.71*** (-10.81)	-159.69*** (-10.80)
Gas price $\times d_{low-fee}$	-16.65** (-2.13)	-15.41* (-1.98)	-15.80** (-2.02)	-11.29 (-1.65)	2.95 (0.40)	2.94 (0.39)
Gas price $\times d_{high-fee}$	-14.44** (-2.04)	-13.42* (-1.89)	-13.98* (-1.96)	-10.52* (-1.69)	1.96 (0.32)	1.95 (0.32)
Volume		-5.87 (-1.15)	-7.45 (-1.41)		-24.84*** (-4.10)	-24.82*** (-4.09)
Total value locked		-53.17* (-1.70)	-51.72* (-1.66)		-12.71 (-0.52)	-12.71 (-0.52)
Volatility		-2.11*** (-2.75)	-2.26*** (-2.86)		-2.99*** (-3.36)	-2.98*** (-3.36)
Position out-of-range			37.09*** (6.43)			-1.53 (-0.22)
Constant	497.18*** (91.65)	497.00*** (90.60)	479.22*** (82.13)	248.00*** (29.91)	250.13*** (30.27)	250.47*** (30.32)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Trader wallet FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	405,586	405,584	405,584	265,848	265,848	265,848
R-squared	0.82	0.82	0.82	0.37	0.37	0.37

Robust t-statistics in parentheses. Standard errors are clustered at week level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The reaction time of a trader may depend on capital constraints, network congestion, and other confounding factors that can correlate with gas prices. As a robustness check, we repeat the analysis above with burn-to-mint times as the dependent variables. The burn-to-mint time measures the speed at which LPs deposit liquidity at updated prices after removing (out-of-range) positions, and should not depend on the rate at which liquidity is consumed. Consistent with the theory, we find no significant relationship between gas price and the burn-to-mint duration (Models 4 through 6).

4.3 Adverse selection costs across low- and high-fee pools

Our findings indicate that larger aggressive orders tend to execute on high-fee liquidity pools. Given that, at least in equity markets, larger incoming orders tend to be better informed (Hasbrouck, 1991), a natural question is whether liquidity providers on high-fee pools face higher adverse selection.

A widely-used measure for gauging informational costs for liquidity providers on decentralized exchanges employing automated market makers (AMMs) is the “impermanent loss,” the equivalent of adverse selection measures in traditional limit order markets (see, for example, Aoyagi, 2020; Barbon and Ranaldo, 2021). The impermanent loss (IL) is defined as the negative return from providing liquidity as opposed to holding the assets outside the exchange and marking them to market as the price evolves. Intuitively, if the fundamental value of the asset changes, an arbitrageur trades at the stale price in the direction of the price change, thus minimizing the value of the liquidity pool (Milinois, Moallemi, Roughgarden, and Zhang, 2023).

To illustrate, let’s consider a straightforward example. At some initial time t_0 , a liquidity provider deposits 1 ETH (i.e., a token quantity of $T_0 = 1$) and 3000 USDC (i.e., a numeraire $N_0 = 3000$) into a liquidity pool, establishing a token price of $p_0 = 3000$. If at the next time interval ($t = 1$), the intrinsic value surges to $v_1 = 3500$, an arbitrageur has the incentive to remove ETH from the pool and deposit USDC until the price reflects the intrinsic value, or $\frac{N_1}{T_1} = 3500$.

Given that the AMM requires the product of token and numeraire quantities to be constant ($T_0 N_0 = T_1 N_1$), this leads to $T_1 = 0.926$ and $N_1 = 3240.37$. In terms of the numeraire equivalent, the liquidity provider’s position is valued at $V_{\text{pos}} = \$3500 \times 0.926 + \$3240.37 = \$6480.74$. Had the liquidity provider not deposited their tokens on the exchange and instead held it in its own account, the value would be $V_{\text{hold}} = \$3500 \times 1 + \$3000 = \$6500$ (that is, equal to $T_0 \times p_1 + N_0$).

Hence, disregarding gas cost and liquidity fees, the impermanent loss can be calculated as:

$$\text{ImpermanentLoss} = \frac{V_{\text{hold}} - V_{\text{pos}}}{V_{\text{hold}}} = 29 \text{ bps.} \quad (16)$$

However, liquidity providers on Uniswap V3 pools can set a price range for their orders. While this feature caps the potential adverse selection cost, it simultaneously introduces a new layer of computational complexity. In Appendix E, we present the exact formulas for calculating

impermanant loss on Uniswap V3, based on the methodology described by [Heimbach, Schertenleib, and Wattenhofer \(2022\)](#).

We obtain hourly liquidity snapshots from the Uniswap V3 Subgraph, which allow us to calculate the impermanant loss for symmetric liquidity positions within a price range of $[\frac{1}{\alpha}p, \alpha p]$ centered around the current pool price p . We consider different ranges of α , specifically $\alpha \in \{1.01, 1.05, 1.1, 1.25\}$. The liquidity management horizon, which represents the time delay between measuring p_0 and p_1 , is set to one hour.

Table 6 presents the empirical results. In Models 1, 3, 5, and 7, where we exclude control variables, the impermanant loss appears to be higher in low-fee pools. However, the link is mechanical: since high-fee pools exhibit lower trading activity, the price updates less frequently leading to a lower measured impermanant loss. Once we account for trading activity in our analysis (Models 2, 4, 6, and 8), we find that impermanant loss is between 1.09 and 1.85 basis points lower in low-fee pools, although the effect is not statistically significant in all specifications. This finding aligns with the intuition that larger incoming orders on high-fee pools tend to impose higher adverse selection costs on liquidity providers.

Table 6: Impermanent loss across high- and low-fee pools

This table reports the coefficients of the following regression:

$$\text{Impermanent Loss}_{ijt} = \alpha + \beta_0 d_{\text{low-fee}, ij} + \sum \beta_k \text{Controls}_{ijt} + \theta_j + \varepsilon_{ijt}$$

where the dependent variable is the daily average impermanent loss for a liquidity position with price range $[\frac{p}{\alpha}, \alpha p]$ around the current pool price p , for $\alpha \in \{1.01, 1.05, 1.1, 1.25\}$. The average impermanent loss is computed across all Ethereum blocks mined within the day. To compute the impermanent loss, we use a liquidity provider horizon of one hour: that is, we compare the current pool price with the pool price one hour later. $d_{\text{low-fee}, ij}$ is a dummy that takes the value one for the pool with the lowest fee in pair j and zero else. GasPrice_{jt} is the average of the lowest 100 bids on liquidity provision events across all pairs on day t , standardized to have a zero mean and unit variance. TradeCount_{jt} is number of trades on pool j and day t , standardized to have a zero mean and unit variance. Volume is the natural logarithm of the sum of all swap amounts on day t , expressed in thousands of US dollars. $\text{Total value locked}$ is the natural logarithm of the total value locked on Uniswap v3 pools on day t , expressed in millions of dollars. Volatility is computed as the daily range between high and low prices on the most active pool for a given pair. All regressions include pair and week fixed-effects. Robust standard errors in parenthesis are clustered by week and ***, **, and * denote the statistical significance at the 1, 5, and 10% level, respectively. The sample period is from May 4, 2021 to July 15, 2023.

Impermanent loss for a liquidity position with range $[\frac{p}{\alpha}, \alpha p]$ around price p								
	$\alpha = 1.01$		$\alpha = 1.05$		$\alpha = 1.10$		$\alpha = 1.25$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$d_{\text{low-fee}}$	2.59*** (11.26)	-1.38 (-1.57)	1.08*** (5.72)	-1.85** (-2.28)	0.71*** (4.28)	-1.56** (-2.18)	0.37** (2.58)	-1.09* (-1.98)
Gas price		4.75*** (3.97)		3.68*** (3.96)		2.72*** (3.86)		1.55*** (3.42)
Trade count		4.82*** (4.59)		3.56*** (3.71)		2.78*** (3.30)		1.79*** (2.83)
Volume		3.03*** (7.00)		1.19*** (3.87)		0.60** (2.45)		0.22 (1.25)
Total value locked		0.43 (0.16)		1.78 (0.79)		2.02 (1.05)		1.83 (1.38)
Volatility		6.98*** (2.69)		6.65** (2.59)		6.39** (2.51)		6.06** (2.40)
Constant	15.52*** (134.72)	15.87*** (103.02)	7.37*** (77.84)	7.65*** (60.07)	4.63*** (55.47)	4.87*** (43.73)	2.45*** (34.58)	2.66*** (29.33)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,250	40,248	40,250	40,248	40,250	40,248	40,250	40,248
R-squared	0.17	0.23	0.09	0.15	0.06	0.12	0.03	0.08

Robust t-statistics in parentheses. Standard errors are clustered at week level.

*** p<0.01, ** p<0.05, * p<0.1

5 Conclusion

This paper argues that fixed costs associated with liquidity management drive a wedge between large (institutional) and small (retail) market makers. In the context of blockchain-based decentralized exchanges, the most evident fixed cost is represented by gas fees, where market makers compensate

miners and validators for transaction processing in proof-of-work, respectively in proof-of-stake blockchains. Innovative solutions such as Proof of Stake (PoS) consensus algorithms and Layer 2 scaling aim to address the concern of network costs. However, even if gas fees were eliminated entirely, individual retail traders still encounter disproportionate fixed costs in managing their liquidity, such as the expenditure of time and effort.

Our paper highlights a trade-off between capital efficiency and the fixed costs of active management. During the initial phase of decentralized exchanges, such as Uniswap V2, liquidity providers were not able to set price limits, resulting in an even more passive liquidity supply and fewer incentives for active position management. However, the mechanism implied that incoming trades incurred significant price impact. To enhance the return on liquidity provision and reduce price impact on incoming trades, modern decentralized exchanges (DEXs) have evolved to enable market makers to fine-tune their liquidity positions, albeit at the expense of more active management.

We show, both theoretically and empirically, that fixed costs of liquidity management promote market fragmentation across decentralized pools and generate clienteles of liquidity providers. Large market makers, likely institutions and funds, have stronger economies of scale and can afford to frequently manage their positions on very active low fee markets. On the other hand, smaller retail liquidity providers become confined to high fee markets with scant activity, trading off a lower execution probability against lower gas costs to update their positions. Since large liquidity providers can churn their position at a faster pace, two thirds of the trading volume interacts with less than half the capital locked on Uniswap V3.

Our findings indicate that substantial fixed costs can hinder the participation of small market makers in the forefront of liquidity provision, where active order management is crucial. Instead, smaller liquidity providers tend to operate on the market maker “fringe,” opting for a lower execution probability in exchange for better prices. The results are particularly relevant the context of a resurgence in retail trading activity and the ongoing evolution of technology that fosters market structures aimed at enhancing broader access to financial markets.

References

- Adams, Hayden, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson, 2021, Uniswap v3 core, .
- Alizadeh, Sassan, Michael W. Brandt, and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *The Journal of Finance* 57, 1047–1091.
- Aoyagi, Jun, 2020, Liquidity Provision by Automated Market Makers, *Working paper*.

- , and Yuki Ito, 2021, Coexisting Exchange Platforms: Limit Order Books and Automated Market Makers, *Working paper*.
- Aspris, Angelo, Sean Foley, Jiri Svec, and Leqi Wang, 2021, Decentralized exchanges: The “wild west” of cryptocurrency trading, *International Review of Financial Analysis* 77, 101845.
- Augustin, Patrick, Roy Chen-Zhang, and Donghwa Shin, 2022, Reaching for yield in decentralized financial markets, LawFin Working Paper No. 39. Available at SSRN: <https://ssrn.com/abstract=4063228>.
- Barbon, Andrea, and Angelo Rinaldo, 2021, On the quality of cryptocurrency markets: Centralized versus decentralized exchanges, .
- Battalio, Robert, Shane A. Corwin, and Robert Jennings, 2016, Can brokers have it all? on the relation between make-take fees and limit order execution quality, *The Journal of Finance* 71, 2193–2237.
- Brolley, Michael, and David Cimon, 2020, Order flow segmentation, liquidity and price discovery: The role of latency delays, *Journal of Financial and Quantitative Analysis* 55, 2555–2587.
- Brolley, Michael, and Marius Zoican, 2022, On-demand fast trading on decentralized exchanges, *Finance Research Letters* p. 103350.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics* 130, 1547–1621.
- Caparros, Basile, Amit Chaudhary, and Olga Klein, 2023, Blockchain scaling and liquidity concentration on decentralized exchanges, Available at SSRN: <https://ssrn.com/abstract=4475460> or <http://dx.doi.org/10.2139/ssrn.4475460>.
- Capponi, Agostino, and Ruizhe Jia, 2021, The Adoption of Blockchain-based Decentralized Exchanges, *Working paper*.
- , and Shihao Yu, 2022, The Information Content of Blockchain Fees, *Working paper*.
- , 2023, Price discovery on decentralized exchanges, .
- Cimon, David A., 2021, Broker routing decisions in limit order markets, *Journal of Financial Markets* 54, 100602.
- Demsetz, Harold, 1968, The Cost of Transacting*, *The Quarterly Journal of Economics* 82, 33–53.
- Foucault, Thierry, Ohad Kadan, and Eugene Kandel, 2013, Liquidity cycles and make/take fees in electronic markets, *Journal of Finance* 68, 299–341.


- Foucault, Thierry, and Albert J. Menkveld, 2008, Competition for order flow and smart order routing systems, *The Journal of Finance* 63, 119–158.
- Han, Jianlei, Shiyang Huang, and Zhuo Zhong, 2022, Trust in defi: An empirical study of the decentralized exchange, .
- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *The Journal of Finance* 46, 179–207.
- , Thomas Rivera, and Fahad Saleh, 2022, The Need for Fees at a DEX: How Increases in Fees Can Increase DEX Trading Volume, *Working paper*.
- Heimbach, Lioba, Eric Schertenleib, and Roger Wattenhofer, 2022, Risks and Returns of Uniswap V3 Liquidity Providers, Papers 2205.08904 arXiv.org.
- Korajczyk, Robert A, and Dermot Murphy, 2018, High-Frequency Market Making to Large Institutional Trades, *The Review of Financial Studies* 32, 1034–1067.
- Lehar, Alfred, and Christine Parlour, 2021, Decentralized Exchanges, *Working paper*.
- Milinois, Jason, Ciamac Moallemi, Tim Roughgarden, and Anthony Lee Zhang, 2023, Automated market making and loss-versus-rebalancing, Manuscript.
- O’Hara, Maureen, and Mao Ye, 2011, Is market fragmentation harming market quality?, *Journal of Financial Economics* 100, 459–474.
- Pagano, Marco, 1989, Trading volume and asset liquidity, *The Quarterly Journal of Economics* 104, 255–274.
- Pagnotta, Emiliano, and Thomas Philippon, 2018, Competing on speed, *Econometrica* 86, 1067–1115.
- Park, Andreas, 2022, Conceptual flaws of decentralized automated market making, *Working paper, University of Toronto*.
- Wan, Xin, and Austin Adams, 2022, Just-in-time liquidity on the uniswap protocol, Working paper.
- Zhu, Haoxiang, 2014, Do dark pools harm price discovery?, *Review of Financial Studies* 27, 747–789.

A Liquidity provision mechanism on Uniswap v3

In this appendix, we walk through a numerical example to illustrate the mechanism of liquidity provision and trading on Uniswap V3 liquidity pools. To facilitate understanding, we highlight the similarities and differences between the Uniswap mechanism and the familiar economics of a traditional limit order book.

Let $p_c = 1500.62$ be the current price of the ETH/USDT pair. Traders can provide liquidity on Uniswap V3 pools at prices on a log-linear tick space. In particular, consecutive prices are always θ basis point apart: $p_i = 1.0001^{\theta i}$, where θ is the tick spacing. For the purpose of the example, we take $\theta = 60$. Consequently, the current price of 1500.62 corresponds to a tick index of $c = 73140$. Figure A.1 illustrates three ticks on grid below and above the current price of ETH/USDT 1500.62.

Figure A.1: ETH/USDT price grid around p_c



Two-sided liquidity provision. Trader **A** starts out with a capital of USDT 20,000 and wants to provide liquidity over the price range $[1491.64, 1527.87]$, a range which spans four ticks. Liquidity provision over a range that includes the current price corresponds to posting quotes on both the bid and ask side of a traditional limit order book, where the current price of the pool corresponds to the mid-point of the book.

1. *Bid quotes:* trader **A** deposits USDT over the price range $[1491.64, 1500.62)$. This action is equivalent to submitting a buy limit order with a bid price of 1491.64. An incoming Ether seller can swap their ETH for the USDT deposited by **A**, generating price impact until the limit price of 1491.64 is reached.
2. *Ask quotes:* at the same time, trader **A** deposits ETH over three ticks: $[1500.62, 1509.65)$, $[1509.65, 1518.73)$, and $[1518.73, 1527.87)$. The action corresponds to submitting *three* sell limit orders with ask prices 1509.65, 1518.73, and 1527.87, respectively. Incoming Ether buyers can swap USDT for trader **A**'s ETH.

In the Uniswap V3 protocol, deposit amounts over each tick $[p_i, p_{i+1})$ must satisfy

$$\text{ETH deposit over } [p_i, p_{i+1}): x_i = L \left(\frac{1}{\sqrt{p_i}} - \frac{1}{\sqrt{p_{i+1}}} \right) \quad (\text{A.1})$$

$$\text{USDT deposit over } [p_i, p_{i+1}): y_i = L (\sqrt{p_{i+1}} - \sqrt{p_i}), \quad (\text{A.2})$$

where L (“liquidity units”) is a scaling factor proportional to the capital committed to the liquidity position. The scaling factor L is pinned down by setting the total committed capital equal to the sum of the positions (in USDT), that is $p_c \sum_i x_i + \sum_i y_i$. In our example,

$$1500.62 \times L_A \times \left(\frac{1}{\sqrt{1500.62}} - \frac{1}{\sqrt{1527.87}} \right) + L_A \times \left(\sqrt{1500.62} - \sqrt{1491.64} \right) = 20000, \quad (\text{A.3})$$

leading to $L_A = 43188.6$. We the value of L_A into (A.1) and conclude that trader **A** deposits 5,013.38 USDT over [1491.64, 1500.62) and ETH 9.99 over [1500.62, 1527.87) (approximately ETH 3.33 over each tick size covered).

One-sided liquidity provision. Trader **B** has USDT 20,000 and wants to post liquidity over the range [1509.65, 1527.87), which does not include the current price. This action corresponds to posting ask quotes to sell ETH deep in the book, at price levels 1518.73 and 1527.87. Liquidity is not “active” – that is, the quotes are not filled – until the existing depth at 1509.65 is consumed by incoming trades.

We use equation (A.1) to solve for the amount of liquidity units provided by **B**:

$$1500.62 \times L_B \times \left(\frac{1}{\sqrt{1509.65}} - \frac{1}{\sqrt{1527.87}} \right) = 20000, \quad (\text{A.4})$$

which leads to $L_B = 86589.4$. Trader **B** deposits 6.67 ETH on each of the two ticks covered by the chosen range.

Figure A.2: ETH/USDT pool state after liquidity provision choices




Figure A.2 illustrates market depth after A and B deposit liquidity in the pool. The current price of the pool is equivalent to a midpoint in traditional limit order markets. The “ask side” of the pool is deeper, consistent with both liquidity providers choosing ranges skewed towards prices above the current midpoint. Liquidity is uniformly provided over ticks – that is, each trader deposits an equal share of their capital at each price tick covered by their price range.

Trading, fees, and price impact. Suppose now that a trader C wants to buy 10 ETH from the pool. For each tick interval $[p_i, p_{i+1})$, price impact is computed using a constant product function over virtual reserves:

$$\underbrace{\left(x + \frac{L}{\sqrt{p_{i+1}}}\right)}_{\text{Virtual ETH reserves}} \underbrace{(y + L\sqrt{p_i})}_{\text{Virtual USDT reserves}} = L^2, \quad (\text{A.5})$$

where x and y are the actual ETH and USDT deposits in that tick range, respectively. Virtual reserves are just a mathematical artifact: they extend the physical (real) deposits as if liquidity would be uniformly distributed over all possible prices on the real line. Working with constant product functions over real reserves is not feasible: in our example, the product of real reserves is zero throughout the order book (since only one asset is deposited in each tick range).

Let $\tau = 1\%$ denote the pool fee that serves as an additional compensation for liquidity providers. That is, if the buyer pays to pay Δy USDT to purchase a quantity Δx ETH, he needs to effectively pay $\Delta y(1 + \tau)$. As per the Uniswap V3 white paper, liquidity fees are not automatically deposited back into the pool.

1. **Tick 1:** [1500.62, 1509.65). Trader **C** first purchases 3.33 ETH at the first available tick above the current price (equivalent to the “best ask”). To remove the ETH, he needs to deposit Δy_1 USDT, where Δy_1 solves:

$$\left(3.33 - 3.33 + \frac{L_A}{\sqrt{1509.65}}\right) \left(0 + \Delta y_1 + L_A \sqrt{1500.62}\right) = L_A^2, \quad (\text{A.6})$$

which leads to $\Delta y_1 = 5026.19$ USDT. Trader **C** pays an average price of $5026.19/3.33=1507.86$ USDT for each unit of ETH purchased. Further, he pays a fee of 50.26 USDT to liquidity provider **A** (the only liquidity provider at this tick).

The new current price is given by the ratio of virtual reserves,

$$p' = \frac{\Delta y_1 + L_A \sqrt{1500.62}}{3.33 - 3.33 + \frac{L_A}{\sqrt{1509.65}}} = 1509.65, \quad (\text{A.7})$$

that is the next price on the tick grid since **C** exhausts the entire liquidity on [1500.62, 1509.65).

2. **Tick 2:** [1509.65, 1518.73). Trader **C** still needs to purchase 6.67 ETH at the next tick level (where the depth is 10 ETH). The liquidity level at this tick is $L_A + L_B$, that is the sum of liquidity provided by **A** and **B**. To remove the 6.67 ETH from the pool, he needs to deposit Δy_2 , where

$$\left(10 - 6.67 + \frac{L_A + L_B}{\sqrt{1518.73}}\right) \left(0 + \Delta y_2 + (L_A + L_B) \sqrt{1509.65}\right) = (L_A + L_B)^2. \quad (\text{A.8})$$

It follows that trader **C** purchases 6.67 ETH by depositing $\Delta y_2 = 10089.12$ USDT, at an average price of 1512.61. The pool price is updated as the ratio of virtual reserves:

$$p'' = \frac{\Delta y_2 + (L_A + L_B) \sqrt{1509.65}}{10 - 6.67 + \frac{L_A + L_B}{\sqrt{1518.73}}} = 1515.7. \quad (\text{A.9})$$

The updated price is in between the two liquidity ticks, since not all depth on this tick level was exhausted in the trade. Following the swap, liquidity on the tick range [1509.65, 1518.73) is composed of both assets: that is 10089.12 USDT and $10-6.67=3.33$ ETH.



Finally, trader **C** pays 100.89 USDT as liquidity fees (1% of the trade size), which are distributed to **A** and **B** proportionally to their liquidity share. That is, **A** receives a fraction

$\frac{L_A}{L_A+L_B}$ of the total fee (33.57 USDT), whereas **B** receives 67.32 USDT.

Figure A.3 illustrates the impact of the swap. Within tick [1500.62, 1509.65), **A** sells 3.33 ETH and buys 5026 USDT. Unlike on limit order books, the execution does not remove liquidity from the book. Rather, **A**'s capital is converted from one token to another and remains available to trade. This feature underscores the passive nature of liquidity supply on decentralized exchanges. Mapping the concepts to traditional limit order book, this mechanism would imply that every time a market maker's sell order is executed at the ask, a buy order would automatically be placed on the bid side of the market.

The final price of 1517.70 lies within the tick [1509.65, 1518.73), rather than on its boundary. Trader **C** only purchases 6.66 ETH out of 10 ETH available within this price interval. The implication is that liquidity on [1509.65, 1518.73) contains both tokens: 3.33 ETH (the amount that was not swapped by **C**) as well as 10089.12 USDT that **C** deposited in the pool.

Figure A.3: Swap execution and price impact



The bottom panel of Figure A.3 shows the price impact of the swap. From equation (6.15) in the Uniswap V3 white paper, we can solve for the price within tick $[p_{\min}, p_{\max})$ with liquidity L , following the execution of a buy order of size x :

$$p(x) = \frac{p_{\min} L^2}{(L - \sqrt{p_{\min}} x)^2}. \quad (\text{A.10})$$

As expected, the price impact of a swap decreases in the liquidity available L – each ETH unit purchased by **C** has a smaller impact on the price once tick 1509.65 is crossed and the market becomes deeper.

B Notation summary

Variable Subscripts	
Subscript	Definition
T and N	Pertaining to the token and numeraire assets, respectively.
L and H	Pertaining to the low- and high-fee pool, respectively.
LP	Pertaining to liquidity providers.
LT	Pertaining to liquidity traders.
Exogenous Parameters	
Parameters	Definition
v, Δ	Common and private values of the token.
ℓ, h	Liquidity fee on the low- and high-fee pool.
q_i	Token endowment of liquidity provider i .
$\varphi(q, Q)$	Distribution density of LP endowment q .
Q	Maximum token endowment for LPs .
Γ	Gas price on the blockchain.
θ	Linear trading rate of small LTs .
λ	Poisson arrival rate of large LTs .
Endogenous Quantities	
Variable	Definition
\mathcal{L}_k	Equilibrium liquidity supply on exchange $k \in \{\mathbf{L}, \mathbf{H}\}$.
q_t^*	Token endowment of the LP who is indifferent between pools.
\underline{q}	Lowest token endowment deposited on the market (from break-even condition).
d_k	Duration of a liquidity cycle on exchange k .
π_k	Expected liquidity provider profit on exchange k .

C Proofs

Proposition 1

Proof. We rewrite equation (10) as

$$\pi_L - \pi_H = \frac{1}{\lambda d_H d_L} \left[q_i \left(h \times \left(\frac{q_t}{Q} \right)^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} - (h-l) \right) - \Gamma \left(\frac{q_t}{Q} \right)^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \right]. \quad (\text{C.1})$$

First off, if the coefficient of q_i in (C.1) is negative, that is $h \times \left(\frac{q_t}{Q} \right)^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} - (h-l) \leq 0$, then $\pi_L - \pi_H < 0$ and all **LPs** optimally choose to post liquidity on pool H . A necessary condition for a fragmented equilibrium is therefore that

$$f_1(q_t) := h \times \left(\frac{q_t}{Q} \right)^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} - (h-l) > 0 \quad (\text{C.2})$$

The function $f_1(q_t)$ increases in q_t , and therefore it has at most one root q_r in $[1, Q]$. Since $f_1(Q) = \ell > 0$, it follows that $q_r \in [1, Q]$ if and only if

$$f_1(1) < 0 \Rightarrow \frac{h-l}{h} Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} > 1 \quad (\text{C.3})$$

Second, in any fragmented market equilibrium, the marginal **LP** is indifferent between providing liquidity on the L or H pool. Equivalently, the marginal trader solves $\pi_L - \pi_H = 0$, which from equation (11) translates to a solution of

$$f_2(q_t) := \Gamma \frac{q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}}}{h \left[q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}} \right] - (h-l)} - q_t = 0. \quad (\text{C.4})$$

We show that $f_2(\cdot)$ monotonically decreases in q_t since

$$\frac{\partial f_2(q_t)}{\partial q_t} = -\Gamma \frac{\lambda(h-l)Q^{\frac{\lambda Q}{\theta(Q-1)}+1} q_t^{\frac{\lambda Q}{\theta(Q-1)}-1}}{\theta(Q-1) \left((l-h)Q^{\frac{\lambda Q}{\theta(Q-1)}} + h q_t^{\frac{\lambda Q}{\theta(Q-1)}} \right)^2} - 1 < 0, \quad (\text{C.5})$$

and therefore it has at most one root. That is, if a fragmented market equilibrium exists, then it is unique.

To complete the proof, we consider two cases. First, let $\frac{h-l}{h} Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} > 1$ such that $q_r > 1$. For a belief $q_t < q_r < Q$, liquidity providers optimally choose pool H – which is equivalent to having an outcome in which the marginal trader is $q_t = Q$. Therefore, there is no fulfilled expectation equilibrium on $[1, q_r]$ in which beliefs match outcomes. From the intermediate value theorem and monotonicity of $f_2(q_t)$, there is exactly one root of $f_2(q_t)$ on $(q_r, Q]$ since

$$\underbrace{\frac{\Gamma}{\ell} - Q}_{<0} = f_2(Q) < f_2(q_t) < \lim_{q_t \downarrow q_r} f_2(q_t) = \infty. \quad (\text{C.6})$$

Therefore, for $\frac{h-l}{h} Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} > 1$ there is one unique equilibrium with fragmented markets.

Second, consider the case where $\frac{h-l}{h}Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \leq 1$ such that $q_t < 1$. In this case the reaction function $f_2(q_t)$ is decreasing and has no discontinuities. If $\frac{\Gamma}{h} > 1$ such that the marginal **LP** earns zero expected profit, it follows from the intermediate value theorem and monotonicity of f_2 that there is a unique fragmented equilibrium since $f_2\left(\frac{\Gamma}{h}\right) > 0$ and $f_2(Q) < 0$.

Otherwise – if $\frac{\Gamma}{h} < 1$ – a necessary and sufficient condition for a unique interior root is that $f_2(1) > 0$ or, equivalently, that

$$\frac{\Gamma}{h} > 1 - \frac{h-l}{h}Q^{\frac{\lambda}{\theta} \frac{Q}{Q-1}}. \quad (\text{C.7})$$

If (C.7) does not hold, then $f_2(1) < 0$ which, from monotonicity, implies that $f_2(q_t) < 0$ for all $q_t \in [1, Q]$. In this case, the unique equilibrium is that all liquidity providers deposit their endowments on pool L . \square

Lemma C.1. *In equilibrium, the endowment marginal trader q_t^* :*

- (i) *increases in the gas cost (Γ), the fee on pool H (h), and the arrival rate of large trades (λ).*
- (ii) *decreases in the fee on pool L (ℓ) and the arrival rate of small trades (θ).*

Lemma C.1

Proof. From the implicit function theorem,

$$\frac{\partial q_t^*}{\partial x} = - \left(\frac{\partial f_2}{\partial q_t} \right)^{-1} \frac{\partial f_2}{\partial x}. \quad (\text{C.8})$$

Since from equation (C.5) $f_2(\cdot)$ decreases in q_t , it follows that the sign of $\frac{\partial q_t^*}{\partial x}$ is the same as the sign of $\frac{\partial f_2}{\partial x}$, which significantly simplifies the task.

First, we note that

$$\frac{\partial f_2}{\partial \Gamma} = \frac{q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}}}{h \left[q_t^{\frac{\lambda}{\theta} \frac{Q}{Q-1}} \times Q^{-\frac{\lambda}{\theta} \frac{Q}{Q-1}} \right] - (h-l)} > 0, \quad (\text{C.9})$$

and therefore q_t^* increases in Γ . Second,

$$\frac{\partial f_2}{\partial \lambda} = \frac{\Gamma(h-l)Q^{\frac{\lambda Q}{\theta(Q-1)}+1}(\log(Q) - \log(q_t))q_t^{\frac{\lambda Q}{\theta(Q-1)}}}{\theta(Q-1) \left((l-h)Q^{\frac{\lambda Q}{\theta(Q-1)}} + hq_t^{\frac{\lambda Q}{\theta(Q-1)}} \right)^2} > 0, \quad (\text{C.10})$$

since $q_t < Q$. Therefore, q_t^* increases in λ . Third,

$$\frac{\partial f_2}{\partial \theta} = - \frac{\Gamma \lambda (h-l) Q^{\frac{\lambda Q}{\theta(Q-1)}+1} (\log(Q) - \log(q_t)) q_t^{\frac{\lambda Q}{\theta(Q-1)}}}{(Q-1) \left(\theta (h-l) Q^{\frac{\lambda Q}{\theta(Q-1)}} - h \theta q_t^{\frac{\lambda Q}{\theta(Q-1)}} \right)^2} < 0, \quad (\text{C.11})$$

which implies that q_t^* decreases in θ .

Finally, we compute partial derivatives with respect to pool fees h and l . That is

$$\frac{\partial f_2}{\partial h} = \frac{\Gamma Q^{-\frac{\lambda Q}{(Q-1)\theta}} \left(q_t^{-\frac{\lambda Q}{(Q-1)\theta}} - Q^{-\frac{\lambda Q}{(Q-1)\theta}} \right)}{\left((l-h)q_t^{\frac{\lambda Q}{\theta-\theta Q}} + hQ^{\frac{\lambda Q}{\theta-\theta Q}} \right)^2} > 0 \quad (\text{C.12})$$

since $q_t < Q$ and $\frac{\lambda Q}{(Q-1)\theta} > 0$; further,

$$\frac{\partial f_2}{\partial l} = -\frac{\Gamma Q^{\frac{\lambda Q}{\theta-\theta Q}} \text{qm}^{\frac{\lambda Q}{\theta(Q-1)}}}{\left(h \left(Q^{\frac{\lambda Q}{\theta-\theta Q}} \text{qm}^{\frac{\lambda Q}{\theta(Q-1)}} - 1 \right) + l \right)^2} < 0. \quad (\text{C.13})$$

□

Proposition 2

Proof. The comparative statics with respect to λ , θ , and ℓ follow immediately from Lemma C.1 and the fact that w_L decreases in q_t .

Comparative statics for h . If $\frac{\Gamma}{h} > 1$ such that $\underline{q} = \frac{\Gamma}{h}$, then the partial derivative of w_L with respect to h is

$$\frac{\partial w_L}{\partial h} = -\frac{(\log Q - \log q_t) q_t + h (\log q_t - \log \frac{\Gamma}{h}) \frac{\partial q_t}{\partial h}}{h (\log Q - \log q_t)^2 q_t}, \quad (\text{C.14})$$

which is positive since from Lemma C.1 we have that $\frac{\partial q_t}{\partial h} > 0$. Otherwise, if $\frac{\Gamma}{h} > 1$ such that $\underline{q} = 1$, then

$$\frac{\partial w_L}{\partial h} = -\frac{\partial q_t}{\partial h} \frac{1}{\log Q q_t} < 0. \quad (\text{C.15})$$

Comparative statics for Γ . If $\underline{q} = \frac{\Gamma}{h}$ (i.e., if $\frac{\Gamma}{h} > 1$), then the partial derivative of w_L with respect to Γ is

$$\frac{\partial w_L}{\partial \Gamma} = \frac{\Gamma \frac{\partial q_t}{\partial \Gamma} (\log(\frac{\Gamma}{h}) - \log(Q)) + q_t (\log(Q) - \log(q_t))}{\Gamma q_t (\log(Q) - \log(\frac{\Gamma}{h}))^2}. \quad (\text{C.16})$$

From equation (11), q_t is linear in Γ and therefore $q_t = \Gamma \frac{\partial q_t}{\partial \Gamma}$. We can rewrite (C.16) as

$$\frac{\partial w_L}{\partial \Gamma} = -\frac{\log(q_t) - \log(\frac{\Gamma}{h})}{\Gamma (\log(Q) - \log(\frac{\Gamma}{h}))^2} < 0. \quad (\text{C.17})$$

Further, if $\underline{q} = 1$, then it is easy to see that w_L decreases in Γ since q_t increases with Γ and the denominator is not a function of Γ .

□

D Just-in-time liquidity

Just-in-time (JIT) liquidity is a strategy that leverages the transparency of orders on the public blockchains. If a liquidity provider observes an incoming large order that has not been processed by miners and it deems uninformed in the public mempool, it can conveniently re-arrange transactions and propose a sequence of actions to sandwich this trade as follows:

1. Add a large liquidity deposit at block position k , at the smallest tick around the current pool price.
2. Let the trade at block position $k + 1$ execute and receive liquidity fees.
3. Remove or burn any residual un-executed liquidity at block position $k + 2$.

The mint size is optimally very large (i.e., of the order of hundred of millions USD for liquid pairs), such that the JIT liquidity provider effectively crowds out the existing liquidity supply and collects most fees for the trade. That is, the strategy is made possible by pro-rata matching on decentralized exchanges because with time priority, the JIT provider cannot queue-jump existing liquidity providers. Since the JIT liquidity provider does not want to passively provide capital, it removes any residual deposit immediately after the trade.

We identify JIT liquidity events by the following algorithm as in [Wan and Adams \(2022\)](#):

1. Search for mints and burns in the same block, liquidity pool, and initiated by the same wallet address. The mint needs to occur exactly two block positions ahead of the burn (at positions k and $k + 2$).
2. Classify the mint and the burn as a JIT event if the transaction in between (at position $k + 1$) is a trade in the same liquidity pool.

JIT events are rare in our sample, and account for less than 1% of the traded volume on Uniswap v3. Further, more than half of them occur in a single pair - USDC-WETH, and in low-fee pools. The Uniswap Labs provides further discussions on the aggregate impact of JIT liquidity provision [here](#). Regarding the economic effects, JIT liquidity reduces price impact for incoming trades, but dilutes existing liquidity providers in the pro-rata markets, and can discourage liquidity supply in the long run.

E Impermanent loss measure

We build our measure of impermanent loss in line with the definition of token reserves within a price range in the Uniswap V3 white paper (Adams, Zinsmeister, Salem, Keefer, and Robinson, 2021) and Section 4.1 in Heimbach, Schertenleib, and Wattenhofer (2022).

Consider a liquidity provider who supplies L units of liquidity into a pool trading a token x for a token y . The chosen price range is $[p_\ell, p_u]$ with $p_\ell < p_u$. Further, the current price of the pool is p_0 . We are interested in computing the impermanent loss at a future point in time, when the price updates to p_1 .

From Adams, Zinsmeister, Salem, Keefer, and Robinson (2021), the actual amount of tokens x and y (“real reserves”) deposited on a Uniswap v3 liquidity pool with a price range $[p_\ell, p_u]$ to yield liquidity L are functions of the current pool price p :

$$x(p) = \begin{cases} L \times \left(\frac{1}{\sqrt{p_\ell}} - \frac{1}{\sqrt{p_u}} \right) & \text{if } p \leq p_\ell \\ L \times \left(\frac{1}{\sqrt{p}} - \frac{1}{\sqrt{p_u}} \right) & \text{if } p_\ell < p \leq p_u \\ 0 & \text{if } p > p_u \end{cases} \quad \text{and } y(p) = \begin{cases} 0 & \text{if } p \leq p_\ell \\ L \times (\sqrt{p} - \sqrt{p_\ell}) & \text{if } p_\ell < p \leq p_u \\ L \times (\sqrt{p_u} - \sqrt{p_\ell}) & \text{if } p > p_u. \end{cases} \quad (\text{E.1})$$

From equation (E.1), the value of the liquidity position at $t = 1$ is therefore

$$V_{\text{position}} = p_1 x(p_1) + y(p_1) = \begin{cases} L p_1 \times \left(\frac{1}{\sqrt{p_\ell}} - \frac{1}{\sqrt{p_u}} \right) & \text{if } p_1 \leq p_\ell \\ L \times \left(2\sqrt{p_1} - \frac{p_1}{\sqrt{p_u}} - \sqrt{p_\ell} \right) & \text{if } p_\ell < p_1 \leq p_u \\ L \times (\sqrt{p_u} - \sqrt{p_\ell}) & \text{if } p_1 > p_u. \end{cases} \quad (\text{E.2})$$

Conversely, the value of a strategy where the liquidity provider holds the original token quantities and marks them to market at the updated price is

$$V_{\text{hold}} = p_1 x(p_0) + y(p_0) = \begin{cases} L p_1 \times \left(\frac{1}{\sqrt{p_\ell}} - \frac{1}{\sqrt{p_u}} \right) & \text{if } p_0 \leq p_\ell \\ L \times \left(\frac{p_1 + p_0}{\sqrt{p_0}} - \frac{p_1}{\sqrt{p_u}} - \sqrt{p_\ell} \right) & \text{if } p_\ell < p_0 \leq p_u \\ L \times (\sqrt{p_u} - \sqrt{p_\ell}) & \text{if } p_0 > p_u. \end{cases} \quad (\text{E.3})$$

The impermanent loss is then defined as the excess return from holding the assets versus providing liquidity on the decentralized exchange:

$$\text{ImpermanentLoss} = \frac{V_{\text{hold}} - V_{\text{position}}}{V_{\text{hold}}}. \quad (\text{E.4})$$

Empirically, we follow Heimbach, Schertenleib, and Wattenhofer (2022) and compute impermanent loss for “symmetric” positions around the current pool price, that is $p_\ell = p_0 \alpha^{-1}$ and $p_u = p_0 \alpha$, with $\alpha > 1$. We allow for a time lag of 200 blocks between p_0 and p_1 (approximately 45 minutes).