

What Drives Liquidity on Decentralized Exchanges? Evidence from the Uniswap Protocol

Brian Zhu¹, Dingyue Liu², Xin Wan², Gordon Liao³,
Ciamac Moallemi¹, and Brad Bachu²

¹ Columbia University

bzz2101@columbia.edu, ciamac@gsb.columbia.edu

² Uniswap Labs

kite.liu@uniswap.org, xin@uniswap.org, brad.bachu@uniswap.org

³ Circle Internet Financial

gordon@circle.com

Abstract. We study liquidity on decentralized exchanges (DEXs), identifying factors at the platform, blockchain, token pair, and liquidity pool levels with predictive power for market depth metrics. Using the *counterfactual v2 spread* metric, we decompose the contribution of each factor on market depth into two channels: total value locked (TVL) and concentration. We further explore how external liquidity from competing DEXs and private inventory on DEX aggregators influence market depth. We find that (i) gas prices, token price returns, and the market share of trading volume affect liquidity through concentration, (ii) internalization of order flow by private market makers affects TVL but not the overall market depth, and (iii) token price volatility, fee revenue, and markout affect liquidity through both channels.

1 Introduction

Liquidity plays a fundamental role in financial markets, serving as a critical determinant of market efficiency and stability. This is particularly evident in traditional finance (TradFi), where liquidity impacts execution prices, price discovery, and overall market robustness. Extensive research has explored how factors such as asset volatility and investor behavior shape liquidity in TradFi. However, the evolving nature of decentralized finance (DeFi) introduces dynamics for liquidity provision that remain under-explored.

Decentralized exchanges (DEX) introduce novel paradigms for liquidity provision and trading, utilizing *liquidity pools* and *pricing functions* as opposed to limit order books. Understanding the dynamics of liquidity in DEXs under this new paradigm is not only important for traders and investors, but also for the design and development of DEXs. While a substantial body of literature exists in TradFi regarding liquidity, there is a pressing need for more research on the idiosyncratic elements of liquidity provision in DEXs. This paper addresses this gap by investigating the forces that drive liquidity and market depth in DEXs, contributing to both academic discourse and practical applications in DeFi.

One recent development in DeFi has been the rise of liquidity aggregators, which combine liquidity from on- and off-chain sources to deliver better execution prices for trades. While research has shown that these services improve prices for traders, their impact on liquidity provision in AMMs is less studied. We answer this question by analyzing if and how these services affect on-chain liquidity.

We focus on liquidity dynamics within the Ethereum ecosystem, examining pools on the Uniswap v3 protocol deployed on the Ethereum Mainnet (L1) and Layer 2 (L2) networks. As the primary blockchain for decentralized applications, Ethereum is host to a variety of DEXs, with Uniswap standing out as the leading platform in trading volume, total value locked (TVL), and user adoption. While our analysis focuses on Uniswap v3, we show that our framework is applicable to a broader class of AMMs, including those used on Uniswap v2 and v4.

Our Contributions. The results of our analysis offer valuable insights into the determinants of liquidity in AMMs. The key contributions of this paper are:

1. **Identifying on-chain predictive factors for liquidity:** We identify factors on period t that forecast various market depth metrics on period $t + 1$. We find that gas prices, token pair returns and volatilities, and in-pool fee revenue and markout have significant explanatory power on future market depth, consistent with prior theoretical results.
2. **Decomposing the contribution of each factor on liquidity into two channels:** With the *counterfactual v2 spread* metric, we decompose effective spreads into TVL and concentration components. This allows us to identify the channel(s) in which changes to market depth occur, whether through the *deployment* and/or *concentration* of liquidity.
3. **Understanding impacts of external liquidity:** We examine the impact of liquidity sources outside Uniswap v3 pools, focusing on competing DEX liquidity and off-chain liquidity used by aggregators. We find that a higher competitor market share negatively impacts liquidity, while more internalization by fillers has no significant impact on overall liquidity.

Related Literature. Our paper contributes to the literature on liquidity provision in DEXs. Some studies focus on incentives for/against liquidity providers (LPs). Lehar and Parlour [15] as well as Capponi and Jia [5] study equilibrium in liquidity pools, showing that volatility arbitrage risk causes LPs to exit pools. Capponi, Jia, and Zhu [7] analyze the phenomenon of just-in-time liquidity [23], showing that it may lead to shallower pools by taking fees away from and leaving toxic order flow to passive LPs. An important factor behind these incentives are the losses incurred by LPs. Two popular loss metrics are *impermanent loss* [12,13,17,18] and *loss-versus-rebalancing* [20,21], which compare the profitability of providing liquidity against holding and rebalancing, respectively.

Other studies focus on liquidity in AMMs that use *concentrated liquidity*, e.g. Uniswap v3. Lehar, Parlour, and Zoican [16] find that large (small) LPs prefer low-fee (high-fee) pools on Uniswap v3 and adjust their positions (in)frequently. Lyandres and Zaidelson [19] examine capital allocation on Uniswap v3, finding

that market efficiency causally impacts capital efficiency. Cartea et al. [8] and Fan et al. [11] study strategies for liquidity provision on Uniswap v3.

We contribute to this literature by providing the first comprehensive empirical analysis regarding determinants of liquidity and market depth on DEXs. We consider the effects of multiple factors on liquidity simultaneously, with our sample spanning three years and across multiple blockchains.

Closest to our *counterfactual v2 spread (Cv2S)* metric for liquidity concentration is the *capital allocation efficiency (CAE)* metric of [19]. While CAE is dependent on the trades that occurred in the pool during the calibration period, Cv2S is a function of trade size and independent of other trades. Thus, two pools with the same “liquidity landscape” can have different CAE values, but always have the same Cv2S given trade size.

Our paper also contributes to the literature on informed trading taking place in DEXs. Capponi, Jia, and Yu [6] show that trades with higher priority fees contain more information and have a higher price impact. Klein et al. [14] analyze information contained in both trade and liquidity events on DEXs, finding evidence of heterogeneity in price impact across several dimensions. We contribute by showing that informed trading within a pool, proxied by markout, has a negative effect on market depth.

Another contribution of our paper is to the literature on liquidity in off-chain exchanges for cryptocurrencies, i.e. centralized exchanges (CEXs) and DEX aggregators. Brauneis et al. [4] study liquidity on CEXs, finding that returns and volume have predictive power on liquidity. Bachu, Wan, and Moallemi [3] provide empirical evidence of DEX aggregators improving prices for traders. Chitra et al. [9] analyze a model with theoretical implications that internalization of order flows by fillers negatively affects on-chain liquidity. We contribute by finding that in practice, there is no evidence of such an effect, which speaks positively to the coexistence of DEXs with order flow auction venues.⁴

2 Background

Automated Market Makers. AMMs use liquidity pools and algorithmic pricing functions to facilitate the on-chain exchange of tokens. When LPs deposit tokens into a pool, they receive pool tokens that indicate their stake of the pool and determine the amount they withdraw. Traders typically have to pay a fee proportional to the trade size; this fee is distributed pro-rata among LPs by their stake and incentivizes them to stay in the pool.

Many AMMs are constant function market makers (CFMMs), which requires post-trade pool reserves to be on the same level set of the pricing function as pre-trade reserves. For example, in a pool with X tokens X and Y tokens Y , a trade of Δ_X tokens X for Δ_Y tokens Y must satisfy the relation $F(X + \Delta_X, Y + \Delta_Y) = F(X, Y)$, where $F : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ is the pricing function. AMM protocols using this design include Uniswap v2, Curve, and Balancer.

⁴ Our finding does not necessarily contradict those of [9], since their assumption that filler inventory must be in-wallet does not always hold (cf. UniswapX, CoW Swap).

The “concentrated liquidity” (CL) mechanism for AMMs [2], first pioneered by Uniswap v3, allows LPs to choose the price range in which their liquidity is active, with narrower ranges yielding higher “virtual” liquidity. This allows for higher capital efficiency, but also introduces new complexities for LPs in terms of managing risk and exposure, since out-of-range LP positions do not earn fees and poorly concentrated pools may increase trading costs.

DEX Aggregators. The success of AMMs has led to a proliferation in DEXs, with there being over 100 DEXs at the time of writing. This growth has led to the fragmentation of liquidity across multiple DEXs. In response, protocols such as 1inch Fusion, CowSwap, and UniswapX have introduced new methods to handle order flow, leveraging liquidity from various on-chain sources and off-chain private market makers (PMMs), to optimize trading outcomes for users in a fragmented ecosystem.

DEX aggregators process order flow from their interfaces and allow specialized users, including PMMs, to determine the ordering and/or routing of trades to achieve better execution prices, most commonly implemented via order flow auctions (OFAs). These OFAs can have varying formats: for example, CowSwap uses batch auctions, whereas 1inch Fusion and UniswapX use Dutch auctions.

3 Data

We use publicly available Uniswap v3 data from May 5, 2021 to July 31, 2024. The liquidity pools in our sample, shown in Table 1, are selected as follows:

- Obtain the top 4 blockchains by average trading volume through the sample period. For each selected blockchain, obtain the top 100 pools by average trading volume through the sample period.
- Select the pools corresponding to the token pair and fee tier combinations appearing in all four top-100 lists.

Pair \ Network	Ethereum	Arbitrum (L2)	Optimism (L2)	Polygon (L2)
CRV–WETH	30 bps	30 bps	30 bps	30 bps
DAI–WETH	30 bps	30 bps	30 bps	30 bps
LDO–WETH	30 bps	30 bps	30 bps	30 bps
LINK–WETH	30 bps	30 bps	30 bps	30 bps
USDC–WETH	5, 30 bps	5, 30 bps	5, 30 bps	5, 30 bps
WBTC–WETH	5, 30 bps	5, 30 bps	5, 30 bps	5, 30 bps
WETH–USDT	5, 30 bps	5, 30 bps	5, 30 bps	5, 30 bps

Table 1: Liquidity Pools Included in Sample by Pair, Network, and Fee Tier.

3.1 Liquidity Metrics

Effective Spread. We compute the effective spread, which we define as the difference in quoted price between buying and selling a fixed amount of a given token in a liquidity pool, minus transaction fees. By using quoted prices rather than execution prices and subtracting out transaction fees, we ignore the impacts of MEV- and fee-related slippage (see [1] for example), which isolates the effect of liquidity on the trading costs. Effective spreads are a key measure of liquidity, representing the difference between buying and selling an asset in a market, with a smaller (larger) effective spread indicates a deeper (shallower) market.

We acquire quoted prices from Uniswap v3 quoter contracts.⁵ These contracts have functions to obtain quotes for buying or selling a token at historical blocks, allowing the user to specify blockchain, token pair, fee tier, and swap size. For a given trade size of Δ WETH, we compute the relative difference in quoted price between buying and selling Δ WETH minus twice the fee tier f . Using Uniswap v3’s quoter contract, the `ExactOutput` function yields the ask price, denoted A , and the `ExactInput` function yields the bid price, denoted B . For each day in our sample period, we obtain quotes every six hours (for four samples per day), and take the spread for the day as the average of these measurements.

Formally, the normalized effective spread (“v3 spread”), in basis points attributed to market depth at day t on a given pool with fee tier f is

$$\text{v3S}_t^{\text{pool}} = 10^4 \times \left(\frac{1}{4} \sum_{i \in [4]} \frac{A_{t,i}^{\text{pool}} - B_{t,i}^{\text{pool}}}{\frac{1}{2}(A_{t,i}^{\text{pool}} + B_{t,i}^{\text{pool}})} - 2f \right)$$

where i indexes the in-day samples. Since quoted prices contain fees, we subtract $2f$ from the average and multiply by 10^4 to obtain basis points. For our selected pools, f is 0.0005 or 0.003, corresponding to 5 or 30 basis points, respectively.

Total Value Locked. Total value locked (TVL), the US dollar value of a pool’s token reserves, is another important measure of liquidity in DEXs. In CFMMs, the TVL is a perfect signal of market depth as execution prices are computed directly based on the pool reserves. In AMMs using CL like Uniswap v3, however, TVL is a noisier signal as market depth also depends on how those reserves are concentrated around the pool’s current tick. Since fee revenue and risk are shared pro-rata on the Uniswap protocol, TVL is also a benchmark that other pool metrics can be normalized against.

We compute TVL by aggregating mint (deposit), burn (withdrawal), and swap events on liquidity pools, keeping track of how token quantities in each pool vary over time. At the end of each day t , we sum the number of each token in a liquidity pool weighted by the end-of-day token price in USD to arrive at the end-of-day TVL.

⁵ <https://docs.uniswap.org/contracts/v3/reference/periphery/lens/Quoter>

Summary Statistics. The summary statistics of our liquidity metrics are displayed in Table 2. We highlight the differences between liquidity in pools on L1 (Ethereum) and L2 networks. Effective spreads are lower, less dispersed, and more right-skewed on L1 pools relative to L2 pools. TVL is higher, less dispersed, and more left-skewed on L1 pools than L2 pools. Notably, both liquidity metrics are relatively normally distributed for L2 pools in our sample. Relevant data visualizations for these liquidity metrics are in the Online Appendix.

	log v3S (L1)	log TVL (L1)	log v3S (L2s)	log TVL (L2s)
<i>N</i>	11371	11371	27069	27069
Mean	-0.6944	17.976	2.6869	13.7407
S.D.	2.1329	1.9306	2.3198	1.9985
Skew	1.7493	-1.5520	0.1135	-0.1417
25%	-1.9277	17.5133	0.9339	12.3722
50%	-1.2828	18.3872	2.5252	13.7500
75%	-0.1236	19.5271	4.1776	15.2480

Table 2: Summary Statistics for Liquidity Metrics.

3.2 Independent Variables

The set of factors to regress our liquidity metrics is motivated by previous theoretical and empirical research on liquidity in DEXs, from which several factors consistently appear: gas price at the blockchain level, price returns and volatility at the token pair level, and adverse selection (informed trading) and fee revenue (noise trading) at the pool level. These are summarized in Table 3, and generally agree with each other on the direction effects of the variables on liquidity.

Variable	Prediction/Finding (Setting)
gas price	[16,18]: ↗ gas price ⇒ ↘ rebalancing frequency (v3)
returns	[8]: returns have ambiguous effects on price ranges (v3)
volatility	[15]: ↗ volatility ⇒ ↘ pool size (v2) [5,8]: ↗ volatility ⇒ wider ranges (v3)
fee revenue	[5,7,15]: ↗ fee revenue ⇒ ↗ pool size (v2) [8]: ↗ fee revenue ⇒ narrower ranges (v3)
adverse selection	[5,7]: ↗ adverse selection ⇒ ↘ pool size (v2) [5]: ↗ adverse selection ⇒ narrower ranges (v3)

Table 3: Factors Affecting Liquidity Provision Studied by Previous Works.

We enhance these studies by empirically testing theoretical predictions and verifying empirical findings in a setup with more variables to rule out confounding. Our counterfactual v2 spread metric also allows us to test implications for both Uniswap v2- and v3-like AMMs with only Uniswap v3 data. We thus select gas prices, returns, volatility, markout, and fee revenue as a baseline set of factors, with markout as a proxy for adverse selection. We compute these variables at a daily frequency, using publicly available data from Dune Analytics. We collect data on gas prices per transaction on each blockchain, token price data from centralized exchanges (CEXs), and data on liquidity and swap events occurring on each pool. More detailed descriptions of each variable follow.

Gas Prices. We compute the average gas price, in USD, of all transactions on a given chain for the current day t .

Log>Returns. Let $\{p_t\}$ be the price ratio of the token pair traded on a pool in our sample, in units of the other token per WETH and measured via CEX prices, at day t . Log-returns at time horizon h_r are then given by

$$\text{LogReturns}_t^{\text{pair}} = 100 \times \log(p_t/p_{t-h_r}).$$

Volatility. We compute the annualized volatility of the token price ratio during day t , using 15-minute intervals to obtain returns. This choice of time interval captures the fine-grained intra-day price variability while reducing the influence of microstructure noise present in shorter intervals.

Fee Revenue. The pro-rated fee revenue (or *pool APR*), computed by dividing total fees accrued to a pool from a day’s swaps by the end-of-day pool TVL, is

$$\text{FeeRevenue}_t^{\text{pool}} = \frac{1}{\text{TVL}_t} \times \frac{f}{1+f} \sum_{\text{swaps } s \text{ on pool in day } t} p_{\tau_s}^{\text{TI}}(s) \cdot q^{\text{TI}}(s)$$

where τ_s is the time of the swap and $p_{\tau_s}^{\text{TI}}(s)$ and $q^{\text{TI}}(s)$ are the dollar price and amount, respectively, of the token that swap s puts *into* the pool.⁶

Markout. Markouts capture the informativeness of trades on an exchange by comparing the price of a trade to a benchmark price sometime after the trade, in our case the pool’s mid-price, which indicates how favorable to the trader the swap was in hindsight. Commonly used in traditional market microstructure, markouts have also been used as a proxy for LVR [16],[21] in DEXs.

For each swap, we compare the swap price with the mid-price of the pool determined by the “current tick” at time $\tau_s + h_m$, where τ_s is the time of a swap s and h_m is the time horizon for computing markout. The resulting difference in

⁶ As fees on Uniswap are determined by the token-in amount, and our swap event tracker includes fees in q^{TI} , we multiply the sum across swaps by $f/(1+f)$.

price is then volume-weighted. We aggregate markouts for all swaps in a given day and normalize by the end-of-day pool TVL:

$$\text{Markout}_t^{\text{pool}} = \frac{1}{\text{TVL}_t} \times \sum_{\text{swaps } s \text{ on pool in day } t} D_s \cdot |q^{TI}(s)| \cdot \left(\left| \frac{q^{TO}(s)}{q^{TI}(s)} \right| - p_{\tau_s+h_m}^{\text{pool}} \right)$$

where $D_s = 1$ if the swapper is selling WETH and $D_s = -1$ otherwise (i.e. buying WETH), $q^{TO}(s)$ is the token-out amount for swap s , and $p_{\tau_s+h_m}^{\text{pool}}$ denotes the pool price, in units of the other token per WETH, at time h_m after the swap occurred. Under this definition, more positive (negative) values indicate better (worse) LP profitability and thus less (more) adverse selection costs from swappers.

4 Methodology

In AMMs with concentrated liquidity, market depth is not only influenced by the TVL in the pool, but is also by how concentrated that liquidity is across different price ranges. In this section, we introduce a novel method to measure concentration, which we use to distinguish between the effects of TVL and liquidity concentration on changes in effective spreads. This decomposition allows us to better understand the mechanics of liquidity provision in concentrated AMM pools and provide insights on LP behavior.

4.1 Decomposing Spread in Uniswap v3

The *counterfactual v2 spread (Cv2S)*, is computed by looking at the TVL in a v3 pool at some given time, counterfactually considering a v2 pool with the same TVL under no trading fees such that the spot price on the v2 pool aligns with the CEX price at that time, and computing the effective spread as described in Section 3 on the counterfactual pool. Note that this v2 pool is counterfactual and does not correspond to any actual pool on Uniswap v2 for the token pair. We align the counterfactual v2 pool’s reserves with the CEX price to simulate the effect of arbitrageurs. A quick derivation (with details in Appendix A) shows that the counterfactual v2 spread for a trade size of Δ WETH can be expressed as a function of token prices and TVL:

$$\text{Cv2S}_t^{\text{pool}} = 10^4 \times \frac{4p_t^{\text{ETH}}}{\text{TVL}_t^{\text{pool}}} \Delta.$$

The quotient between the actual v3 and counterfactual v2 spreads, which we call the *spread quotient (SQ)*, is defined as

$$\text{SQ} := \frac{\text{v3S}}{\text{Cv2S}},$$

is a proxy for how well-concentrated the pool is around its mid-price: as the spread quotient increases, the spread of the actual pool becomes higher relative

to that of the counterfactual pool, meaning that liquidity is not well-concentrated in the actual pool; conversely, as the spread quotient decreases, the spread of the actual pool becomes lower, suggesting a more efficient concentration of liquidity. Taking logarithms of the above equation reveals that

$$\log v3S = \log Cv2S + \log SQ.$$

This motivates the following three regression models:

$$\begin{aligned} \log v3S_{t+1}^{pool} = & \beta_0 + \beta_1 \log \text{GasPrice}_t^{chain} + \beta_2 \text{LogReturns}_t^{pair} + \beta_3 \text{Volatility}_t^{pair} \\ & + \beta_4 \log \text{FeeRevenue}_t^{pool} + \beta_5 \text{Markout}_t^{pool} + \gamma^{pool} + \delta_t + \varepsilon_{t+1}^{pool} \end{aligned} \quad (1)$$

$$\begin{aligned} \log Cv2S_{t+1}^{pool} = & \beta_0 + \beta_1 \log \text{GasPrice}_t^{chain} + \beta_2 \text{LogReturns}_t^{pair} + \beta_3 \text{Volatility}_t^{pair} \\ & + \beta_4 \log \text{FeeRevenue}_t^{pool} + \beta_5 \text{Markout}_t^{pool} + \gamma^{pool} + \delta_t + \varepsilon_{t+1}^{pool} \end{aligned} \quad (2)$$

$$\begin{aligned} \log SQ_{t+1}^{pool} = & \beta_0 + \beta_1 \log \text{GasPrice}_t^{chain} + \beta_2 \text{LogReturns}_t^{pair} + \beta_3 \text{Volatility}_t^{pair} \\ & + \beta_4 \log \text{FeeRevenue}_t^{pool} + \beta_5 \text{Markout}_t^{pool} + \gamma^{pool} + \delta_t + \varepsilon_{t+1}^{pool} \end{aligned} \quad (3)$$

Due to skewed data, we take logarithms of the spread, *GasPrice*, and fee revenue. The terms γ^{pool} and δ_t represent pool-level and day-level fixed effects, respectively, while ε_{t+1}^{pool} is the error term. The pool fixed effects capture time-invariant characteristics specific to each pool, such as whether the pool is included in the default Uniswap interface or other platform-specific settings that remain consistent. Day fixed effects account for factors that affect all pools on a given date, such as regulatory news or shifts in overall market sentiment. Standard errors are clustered at the pool level to account for heteroscedasticity and autocorrelation within pools over time.

4.2 Results and Discussion

We estimate the regression models with a trade size of $\Delta = 1$ WETH to compute spreads, a return horizon h_r of 1 day, and a markout horizon h_m of 5 minutes.⁷ Prior to estimation, we normalize each independent variable in the data matrix to have mean zero and standard deviation one, preserving the significance of the coefficients while allowing for interpretable effect sizes.⁸ Our main results are generally robust to modifications in Δ , h_r and h_m ; specifically, we also have considered $\Delta \in \{0.1, 10\}$, $h_r = 7$ days, and $h_m = 1$ hour. See the Online Appendix for the regression results under robustness checks.

⁷ The gas price and markout variable exhibits extreme values that could disproportionately influence the regression results. To address this, we exclude pool-days where the gas price markout exceeds 5 standard deviations from their respective means, removing 56 pool-days from the sample.

⁸ For example, raw gas prices and markouts may have values that differ in orders of magnitude, so normalizing makes the coefficients easier to interpret in measuring the strength of the relationship between the independent and dependent variables, while keeping t -statistics the same.

Our decomposition implies that for each control, the estimated coefficients from model (2) and (3) sum to the coefficient from estimating model (1), though their significance levels may vary:

$$\hat{\beta}_i^{\log v3S} = \hat{\beta}_i^{\log Cv2S} + \hat{\beta}_i^{\log SQ} \quad \forall i.$$

We also regress the log-TVL at time $t + 1$ on the independent variables at time t for completeness. Since the counterfactual $v2$ spread is a function of price and TVL, and fixed effects by pool and day are included, the coefficients from this regression will equal to those from regression (2) times minus one:

$$\hat{\beta}_i^{\log Cv2S} = -\hat{\beta}_i^{\log TVL} \quad \forall i.$$

Table 4: Baseline Regression Model (1) with Decomposition (2) + (3)

	(1) log v3S	(2) log Cv2S	(3) log SQ	(4) log TVL
log GasPrice	0.213 (0.132)	0.085 (0.126)	0.128*** (0.048)	-0.085 (0.126)
LogReturns	-0.033*** (0.008)	-0.009 (0.006)	-0.024*** (0.005)	0.009 (0.006)
Volatility	0.401*** (0.053)	0.101** (0.044)	0.300*** (0.027)	-0.101** (0.044)
log FeeRevenue	-0.928*** (0.117)	-0.237*** (0.074)	-0.690*** (0.086)	0.237*** (0.074)
Markout	-0.086*** (0.028)	-0.169*** (0.019)	0.083*** (0.021)	0.169*** (0.019)
Observations	38440	38440	38440	38440
N. of groups	40	40	40	40
R^2	0.313	0.078	0.435	0.078

Pool and day fixed effects are included; standard errors are clustered at the pool level.
Note: *p<0.1; **p<0.05; ***p<0.01

Table 4 presents the estimated effects of factors on each dependent variable specification. We omit β_0 estimates for brevity. All factors except for GasPrice significantly impact the overall effective spread. Specifically, $v3$ spreads are increasing in volatility and is decreasing in returns, fee revenue, and markout.

The estimated coefficients on volatility, fee revenue, and markout for regression (2) are significant, suggesting that these factors have predictive power on how TVL affects overall spreads. Higher fee revenue and better markouts against swappers indicate more profitability for LPs, incentivizing them to provide liquidity, thus reducing spreads. Since volatility is associated with LP losses (both impermanent loss and LVR), higher volatility lowers liquidity provision.

All factors are significant in regression (3), implying that they play an important role in liquidity concentration. Higher gas prices increase rebalancing costs, leading to “stale” positions that are not concentrated around current pool prices. More volatility and negative markout (indicating informed trading) lead to LPs widening their price ranges in order to, as explained in [5], create a more convex pricing function that reduces losses to informed traders and volatile prices. Conversely, more fee revenue means that LPs can increase their profits by targeting narrower price ranges with a larger concentration liquidity, according to [8]. Recall that putting a fixed amount of assets in a wider range lowers the amount of “virtual liquidity” (see [2]) in the AMM, resulting in lower spreads.

For returns, we note that in traditional markets, liquidity tends to dry up during market declines and periods of increased volatility [10,22]. This implies that higher returns should positively predict market depth, while higher volatility has a negative effect, which is consistent with our results, suggesting that this stylized fact also carries over to decentralized markets with WETH as numéraire.

5 Extension: External Liquidity

The success of AMMs has led to the proliferation of DEXs, with there being over one hundred DEXs at the time of writing. The increase in competition between DEXs introduces challenges such as the fragmentation of liquidity across multiple liquidity pools and DEXs. In addition, DEX aggregators have created new methods for order routing, using liquidity from various on-chain sources and off-chain private market makers (PMMs). While these services improve trading outcomes [3], PMMs take away fees from on-chain LPs.

5.1 Measuring External Liquidity

To better understand how these external liquidity sources might affect on-chain liquidity provision, we introduce variables that (i) capture the volume of swaps taking place on other DEXs and (ii) filled by private liquidity due to aggregator routing. Using data from Dune Analytics, we track swap volume on other DEXs and routed through aggregators, isolating swaps filled completely by private liquidity by comparing transaction hashes with on-chain events. A simple heuristic to identify these types of swaps is to take all swaps emitting events to aggregator trackers that did not emit a swap event to on-chain data trackers.

Competitor Market Share. For a set \mathcal{D} of DEXs (including the given DEX) and a chain-pair, we compute the fraction of trading volume for that chain-pair occurring outside of the given DEX on day t :⁹

$$\text{CompetitorShare}_t^{\text{chain,pair}} = 1 - \frac{v_t}{\sum_{D \in \mathcal{D}} v_t^D}$$

⁹ Typically, market share for DEXs is evaluated in terms of trading volume. As a single DEX can have several pools trading a pair on a chain, we require a platform-level metric to assess competition and trader’s sentiments towards a given DEX.

where v_t and v_t^D are the swap volumes on the given DEX and DEX D in USD, respectively, for a given chain-pair on day t .

Internalization Ratio. Given a set \mathcal{D} of DEXs, a set \mathcal{A} of aggregators, and a chain-pair, we compute the proportion of swap volume routed or internalized by private market makers active on \mathcal{A} , henceforth referred to as “private volume,” to the total on-chain plus private volume for that chain-pair on day t .¹⁰

$$\text{Internalization}_t^{\text{chain,pair}} = \frac{\sum_{A \in \mathcal{A}} v_t^A}{\sum_{D \in \mathcal{D}} v_t^D + \sum_{A \in \mathcal{A}} v_t^A}$$

where v_t^A is the private volume on aggregator A and v_t^D is the swap volume on DEX D , both in USD, for a given chain-pair on day t .

Extended Regression Model. We take \mathcal{D} as the 130 DEXs whose swap events are tracked on Dune Analytics and \mathcal{A} as the 13 aggregators whose swap events are tracked on Dune Analytics. We add the competitor market share and internalization ratio variables to the baseline model and estimate the model with the dependent variable specifications in models (1)–(3):

$$\begin{aligned} y_{t+1}^{\text{pool}} = & \beta_0 + \beta_1 \log \text{GasPrice}_t^{\text{chain}} + \beta_2 \text{LogReturns}_t^{\text{pair}} + \beta_3 \text{Volatility}_t^{\text{pair}} \\ & + \beta_4 \log \text{FeeRevenue}_t^{\text{pool}} + \beta_5 \text{Markout}_t^{\text{pool}} + \beta_6 \text{CompetitorShare}_t^{\text{chain,pair}} \\ & + \beta_7 \text{Internalization}_t^{\text{chain,pair}} + \gamma^{\text{pool}} + \delta_t + \varepsilon_{t+1}^{\text{pool}} \end{aligned} \quad (4)$$

where $y \in \{\log \text{v3S}, \log \text{Cv2S}, \log \text{SQ}\}$.¹¹

5.2 Results and Discussion

Table 5 displays the estimation results of each dependent variable using the extended model that includes the competitor trading volume share and internalization ratios. We find that a higher competitor share of the token pair predicts higher effective spreads, while there is no significant explanatory power from internalization. Interestingly, the channels in which competitor share and internalization affect market depth differ: the former affects liquidity via concentration while the latter affects liquidity through value locked.

One possible explanation for this difference is that while providing liquidity privately and to competing DEXs are both alternative opportunities for LPs to earn fee revenue, providing liquidity privately is more discretionary, as PMMs can choose which orders to fill, while providing liquidity to a competing DEX requires the LP to take the opposite position of all trades routed to the DEX. The greater risk involved in this option incentivizes LPs to widen price ranges, following the same intuition as the discussion on volatility and markout. Conversely, the lesser

¹⁰ Since an aggregator could route trades to a variety of DEXs, we need \mathcal{D} to include all DEX that the aggregators in \mathcal{A} may route to.

¹¹ We performed robustness checks similar to those in footnote 4.

risks involved in being a PMM mean that LPs choosing this option may not need to widen ranges on existing position, instead directly withdrawing liquidity from pools to serve as private liquidity.

Table 5: Model (4) with External Liquidity Variables

	(1) log v3S	(2) log Cv2S	(3) log SQ	(4) log TVL
log GasPrice	0.178 (0.126)	0.083 (0.120)	0.095** (0.047)	-0.083 (0.120)
LogReturns	-0.033*** (0.008)	-0.009 (0.006)	-0.024*** (0.006)	0.009 (0.006)
Volatility	0.379*** (0.052)	0.089** (0.045)	0.290*** (0.026)	-0.089** (0.045)
log FeeRevenue	-0.869*** (0.119)	-0.201** (0.079)	-0.668*** (0.079)	0.201** (0.079)
Markout	-0.088*** (0.026)	-0.169*** (0.019)	0.081*** (0.020)	0.169*** (0.019)
CompetitorShare	0.222*** (0.065)	0.088 (0.056)	0.134*** (0.031)	-0.088 (0.056)
Internalization	0.062 (0.082)	0.113*** (0.027)	-0.051 (0.070)	-0.113*** (0.027)
Observations	38440	38440	38440	38440
N. of groups	40	40	40	40
R^2	0.327	0.093	0.447	0.093

Pool and day fixed effects are included; standard errors are clustered at the pool level.
Note: *p<0.1; **p<0.05; ***p<0.01

6 Conclusion

This study provides a valuable understanding of what drives liquidity on DEXs, specifically within the Uniswap v3 protocol, having analyzed various factors and their explanatory power in predicting future market depth. We use the counterfactual v2 spread metric to decompose the drivers of overall effective spread, distinguishing between impacts through TVL and liquidity concentration. Our findings suggest that increased competition between DEXs and the presence of private liquidity sources are significant contributors to liquidity fragmentation on Uniswap v3, though they influence market depth via differing channels.

Our findings have significant implications for both LPs and DEX designers. Understanding these dynamics is essential for LPs looking to optimize liquidity provision strategies, and for DEX designers, these insights can guide the development of features that address the adverse effects of liquidity fragmentation. Our results on private liquidity are optimistic for the coexistence of DEX ag-

gregators and on-chain liquidity, as internalization shows no significant effect on *overall* market depth.

Future research could explore more elements of competition and internalization in multi-DEX ecosystems, especially as aggregator services evolve. Further studies on how alternative blockchain environments and emerging Layer 2 solutions support affect liquidity provision can provide a broader perspective on the scalability and sustainability of DEXs in a rapidly growing DeFi landscape.

Acknowledgements and Disclosures. The fifth author was supported by the Briger Family Digital Finance Lab at Columbia Business School and is an advisor to fintech companies.

References

1. Adams, A., Chan, B.Y., Markovich, S., Wan, X.: Don't Let MEV Slip: The Costs of Swapping on the Uniswap Protocol. arXiv Preprint arXiv:2309.13648 (2024)
2. Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D.: Uniswap v3 core. Tech. rep., Uniswap Labs (2021)
3. Bachu, B., Wan, X., Moallemi, C.C.: Quantifying Price Improvement in Order Flow Auctions. arXiv Preprint arXiv:2405.00537 (2024)
4. Brauneis, A., Mestel, R., Riordan, R., Theissen, E.: How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance* **124**, 106041 (2021)
5. Capponi, A., Jia, R.: Liquidity Provision on Blockchain-based Decentralized Exchanges. *Review of Financial Studies* **Forthcoming** (2025)
6. Capponi, A., Jia, R., Yu, S.: Price discovery on decentralized exchanges. Available at SSRN 4236993 (2023)
7. Capponi, A., Jia, R., Zhu, B.: The Paradox of Just-in-Time Liquidity in Decentralized Exchanges: More Providers Can Lead to Less Liquidity. Available at SSRN 4648055 (2023)
8. Cartea, A., Drissi, F., Monga, M.: Decentralised Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision. *SIAM Journal on Financial Mathematics* **15**(3), 931–959 (2024)
9. Chitra, T., Kulkarni, K., Pai, M., Diamandis, T.: An Analysis of Intent-Based Markets. arXiv preprint arXiv:2403.02525 (2024)
10. Chordia, T., Roll, R., Subrahmanyam, A.: Market Liquidity and Trading Activity. *Journal of Finance* **56**(2), 501–530 (2001)
11. Fan, Z., Marmolejo-Cossio, F., Moroz, D., Neuder, M., Rao, R., Parkes, D.C.: Strategic Liquidity Provision in Uniswap V3. In: 5th Conference on Advances in Financial Technologies. pp. 25:1–22 (2023)
12. Heimbach, L., Schertenleib, E., Wattenhofer, R.: Risks and Returns of Uniswap V3 Liquidity Providers. In: Proceedings of the 4th ACM Conference on Advances in Financial Technologies. pp. 89–101 (2022)
13. Kim, H.J., Lee, G.M., Lee, J., Kang, S., Chae, S.W., Park, J.S.: A Comparison of Impermanent Loss for Various CFMMs. In: 2024 IEEE International Conference on Blockchain (Blockchain). pp. 542–548 (2024)
14. Klein, O., Kozhan, R., Viswanath-Natraj, G., Wang, J.: Price Discovery in Cryptocurrencies: Trades versus Liquidity Provision. Available at SSRN 4642411 (2023)
15. Lehar, A., Parlour, C.A.: Decentralized Exchange: The Uniswap Automated Market Maker. *Journal of Finance* **79**(6) (2024)

16. Lehar, A., Parlour, C.A., Zoican, M.: Fragmentation and Optimal Liquidity Supply on Decentralized Exchanges. Available at SSRN 4267429 (2024)
17. Li, T., Naik, S., Papanicolaou, A., Schoenleber, L.: Implied Impermanent Loss: A Cross-Sectional Analysis of Decentralized Liquidity Pools. Available at SSRN 4811111 (2024)
18. Li, T., Naik, S., Papanicolaou, A., Schoenleber, L.: Yield Farming for Liquidity Provision. Available at SSRN 4422213 (2024)
19. Lyandres, E., Zaidelson, A.: Does Market Efficiency Impact Capital Allocation Efficiency? The Case of Decentralized Exchanges. Available at SSRN 4853306 (2024)
20. Milionis, J., Moallemi, C.C., Roughgarden, T.: Automated Market Making and Arbitrage Profits in the Presence of Fees . arXiv Preprint arXiv:2305.14604 (2023)
21. Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L.: Automated Market Making and Loss-Versus-Rebalancing. arXiv Preprint arXiv:2208.06046 (2022)
22. Pástor, L., Stambaugh, R.F.: Liquidity Risk and Expected Stock Returns. *Journal of Political Economy* **111**(3), 642–685 (2003)
23. Wan, X., Adams, A.: Just-in-time Liquidity on the Uniswap Protocol. Available at SSRN 4382303 (2022)

A Derivation of the Counterfactual v2 Spread

As a v2 pool between tokens X and Y aligned to CEX prices has equal values of each token, the pool reserves (X_t, Y_t) given TVL at time t should satisfy

$$(X_t^{pool}, Y_t^{pool}) = \left(\frac{\text{TVL}_t^{pool}}{2p_t^X}, \frac{\text{TVL}_t^{pool}}{2p_t^Y} \right)$$

where p_t^X and p_t^Y are the CEX prices of tokens X and Y, respectively. A swap buying Δ_X tokens X for Δ_Y tokens Y must satisfy $(X - \Delta_X)(Y + \Delta_Y) = XY$. Solving for Δ_Y and scaling by Δ_X yields the counterfactual ask price of

$$A_t^{pool} = \frac{\Delta_Y}{\Delta_X} = \frac{Y_t^{pool}}{X_t^{pool} - \Delta_X}.$$

A swap selling Δ_X tokens X for Δ_Y tokens Y must satisfy $(X + \Delta_X)(Y - \Delta_Y) = XY$. Solving for Δ_Y and scaling by Δ_X yields the counterfactual bid price of

$$B_t^{pool} = \frac{\Delta_Y}{\Delta_X} = \frac{Y_t^{pool}}{X_t^{pool} + \Delta_X}.$$

The counterfactual v2 spread for a trade size of Δ WETH is then

$$\text{Cv2S}_t^{pool} = 10^4 \times \frac{A_t^{pool} - B_t^{pool}}{\frac{1}{2}(A_t^{pool} + B_t^{pool})} = 10^4 \times \frac{4p_t^{ETH}}{\text{TVL}_t^{pool}} \Delta.$$