Liquidity Misallocation on Decentralized Exchanges

March 2023
click here for latest version

Abstract

I show that capital is misallocated across liquidity pools on blockchain-based decentralized exchanges. Many pools have persistent abnormal returns, either with respect to factor models or relative to options-implied liquidity premia. Pools with higher past returns continue to have significantly higher risk-adjusted returns in the future. This return predictability arises because liquidity flow is insensitive to net returns. Instead, liquidity flows chase fee revenues – the part of return that is prominently marketed as APY – but ignore adverse selection losses – the part of return that is implicit and rarely displayed. Aggregate liquidity on decentralized exchanges would have shrunk by one third if liquidity providers were equally sensitive to adverse selection losses.

Keywords: liquidity provision, automated market makers, investor sophistication

1 Introduction

Decentralized exchanges (DEXs) are a key building block of decentralized finance (DeFi). Figure 1 shows the aggregate monthly trading volume on DEXs, which has grown rapidly and surpassed \$150 billion in the "DeFi Summer" of 2021. Due to technical reasons such as transaction costs, most DEXs operate as automated market makers (AMMs), which depart significantly from the traditional limit order book exchanges (LOBs) that have been extensively studied.[1] Who provides liquidity on DEXs? Do they allocate liquidity efficiently across assets? These are the questions I seek to answer in this paper.

I show that liquidity is misallocated across asset pairs on Uniswap, one of the largest DEXs. I arrive at this conclusion through a detail empirical analysis of both the returns to liquidity provision (LP returns) and liquidity providers' actions to deposit or withdraw liquidity (LP flows). LP returns exhibit significant cross-sectional predictability that is hard to reconcile with risk-based explanations. Instead, LP return anomalies can be explained by LP flows, which appear to ignore un-marketed adverse selection losses. As a result, pool with higher realized net returns do not attract sufficient liquidity inflows and continue to yield higher net returns in the future.

I begin by describing in detail how liquidity provision works in a constant-product automated market maker (CPAMM), the most widely adopted design of DEX. I show that liquidity provision on CPAMM is subject to the usual trade-off between gains from noise traders (in the form of fee revenue) and adverse selection losses from informed traders (the so-called "impermanent losses"). In contrast to traditional limit order book exchanges, blockchain-based automated market makers necessitate passive liquidity and facilitate the entry of unsophisticated retail investors. Indeed, more than half of the liquidity positions live longer than a week from open to close.

[1] Section 3 provides details on why automated market markers dominate limit order books on the current blockchains.

In an efficient market, investors are expected to earn zero risk-adjusted returns. I find evidence of persistent return anomalies in the cross section of liquidity pools. First, on average, some pools have average risk-adjusted returns (Sharpe ratios) that are as high as 67.7% per annum (1.77), whereas others as low as -61.2% per annum (-0.71). Secondly, exploiting the fact that adverse selection losses on AMM are price-contingent, I compute options-implied liquidity compensation and find them to be persistently different from realized liquidity compensation. Lastly, I use a sorted portfolio approach to show that pools in the top quintile of past returns continue to have significantly higher risk-adjusted returns than pools in the bottom quintile of past returns, at 66.8% per annum.

Why are there persistent return differences across pools? Why doesn't liquidity adjust to equilibriate the return differences? I show that liquidity flows are insensitive to net returns. Instead, liquidity flows chase fee revenues – the part of return that is prominently marketed as APY – but ignore adverse selection losses – the part of return that is implicit and rarely displayed. As a result, pools with higher fee revenues attract more liquidity, even though they tend to have higher adverse selection losses and insignificant returns on net. I show that taking adverse selection risk into account would significantly increase the payoffs to liquidity providers.

I provide a simple derivation of what would be the counterfactual liquidity flows if liquidity providers care equally about adverse selection losses as fee revenues. There would have been 34 % less aggregate inflow from Uniswap V2 over its 10-month active period.

My results show that the average LPs on DEXs are unsophisticated investors that chase simple signals. From a policy standpoint, my results suggest more regulation – by either the government or industry participants themselves – be put in place to curb the asymmetric marketing of fee revenues vs adverse selection losses. In addition, there should be stricter guidance on the use of "APY" in marketing materials, especially for investment products with volatile returns such as liquidity pools. From a developer standpoint, there are opportunities in building platforms that more accurately describe the gains and losses to retail investors, which should succeed in the long run.

The rest of this paper is organized as follows. Section 2 summarizes this paper's contribution to the literature. Section 3 provides background on constant-product automated market makers. 4 describes the data. Section 5 shows evidence of persistently abnormal LP returns. Section 6 shows evidence of sub-optimal LP flows. Section 8 concludes.

2 Literature

A growing literature studies the returns to liquidity provision on decentralized exchanges, including Angeris et al. [2019], Evans [2020], Aigner and Dhaliwal [2021], Evans et al. [2021], Heimbach et al. [2021], Loesch et al. [2021], Adams and Liao [2022], Heimbach et al. [2022]. Most of the papers are theoretical. On the empirical side, Heimbach et al. [2021] show that Uniswap V2 LP returns are largely driven by adverse selection losses, Loesch et al. [2021] show that most Uniswap V3 LPs make negative returns on net, Adams and Liao [2022] show that Uniswap V3 provide higher fee returns than Uniswap V2 due to concentrated liquidity, and Heimbach et al. [2022] shows that Uniswap V3 LP returns depend heavily on the complexity of LP strategies. My contributions are two folds. First, I highlight that the right benchmark for LP net return is not zero, but rather the premium for bearing risks associated with adverse selection losses. I quantify the risk premium using data from the options market and show a large gap between options-implied vs realized LP returns. Secondly, I show that not only are LP returns abnormal, they are also highly persistent – a simple trading strategy that exploits the persistence of LP returns earn a high risk-adjusted return. These novel empirical findings suggest that misallocation is a key driver of liquidity returns on decentralized exchanges.

A few papers study how liquidity providers behave on decentralized exchanges, including Aoyagi [2020], Capponi and Jia [2021], Neuder et al. [2021], Aoyagi and Ito [2021], Heimbach et al. [2021], Lehar and Parlour [2021]. Most of the papers are theoretical and assume perfectly rational agents. In contrast, I find that LPs are quite unsophisticated – they chase fee revenues, the part of returns prominently marketed as APY, and ignore adverse selection losses, the part of returns that is never displayed. My findings differ from the previous literature due to the following reasons. First, I assemble a comprehensive dataset covering more pools and spanning longer time periods.[2] Secondly, I make improvement on the regression framework by incorporating the latest tools from the mutual fund literature, such as Fama-MacBeth regression, pool fixed effects to account for unobserved persistent factors, inclusion of lagged flows to account for flow persistence, clustering of standard errors for more robust inference, and a battery of robustness checks. Lastly, I build a new dataset on liquidity mining programs – not just those conducted by Uniswap but also those carried out by native protocols – which have been an important omitted variable in previous analyses.

In analyzing LP flows, My methodologies borrow heavily from the mutual fund literature, including Berk and van Binsbergen [2016], Barber et al. [2016], Ben-David et al. [2021], Song [2020]. Song [2020] shows that mutual funds with high factor-related returns accumulate too much assets that their future returns significantly under-perform. In a similar spirit, I show that pools with high net returns do not experience higher inflows and continue to out-perform in the future.

More generally, this paper contributes to the rapidly growing literature on DeFi. I refer to Harvey et al. [2021] for an excellent and comprehensive survey.

My results speak to the dark side of cryptocurrencies [Foley et al., 2019, Cong et al., 2021]

[2] I study all Uniswap V2 pools from their inceptions to May 2021, focusing on before Uniswap V3 was launched in May 2021. Lehar and Parlour [2021] study all Uniswap V1 and V2 pools from their inceptions to May 2021. Capponi and Jia [2021] study 6 Uniswap V2 pools and 6 SushiSwap pools from April to December 2021.

and financial innovations in general [Célérier and Vallée, 2017]. Most liquidity providers on decentralized exchanges seem to be unsophisticated retail investors that do not fully understand the risks involved. More discipline is needed, either by market participants themselves or through regulators, so that efficiency of this market can be improved. However, there is also bright side [Calvet et al., 2020]. Decentralized exchanges democratize market making and allow passive investors to potentially earn high premia from liquidity provision, and I show simple strategies on how to identify and capture these lucrative investment opportunities.

3 Constant-Product Automatic Market Maker

In traditional limit order book exchanges, liquidity is usually supplied by high-frequency trading firms that post and update limit orders at millisecond interval. Such high-frequency operations have been infeasible on blockchains such as Ethereum mainnet due to technical limits on block size and block speed.[3] As a result, limit order book (e.g. Airswap) has not been very successful. Instead, automated market makers (AMMs) have proven to be the successful application. I focus on one particular class of AMMs, the constant-product automatic market maker (CPAMM). As will be shown in Section 4, CPAMM accounts for most of the trading volume and most of the liquidity among all DEXs during my sample period.

3.1 Trading on CPAMM

Each pair of assets forms a pool. For example, suppose that the first asset is BTC and the second asset USD. Denote their amounts in the pool by X and Y, respectively. The

[3]New developments such as Polygon has overcome these technical limits and led to revival of limit order book on blokchains.

exchange of the two assets – e.g. swapping in BTC and swapping out USD, or vice versa – follows the simple rule:

$$XY = K \tag{1}$$

where K stays invariant to all swaps. Figure 2 visualizes the rule. Suppose there are initial X_0 amount of BTC and Y_0 amount of USD. When a trader swaps in ΔX of BTC, the amount of BTC in the pool becomes X_1. The smart calculates calculates what would be the new amount of USD Y_1 so that the pool stays on the the curve. The difference ΔY is what the trader swaps out in return.

Note that the more that one trades, the worse the term is. This is because the curve is convex. If the convexity is zero, i.e. the curve is a straight line, one additional unit of ΔX always gets the same change in ΔY. However, because the curve is convex, one additional unit of ΔX always gets less change in ΔY, compare to the last unit change of ΔX. This design allows price to change in response to trading, and moreover change in the direction that trades are supposed to reveal new information. **?** discuss more on the properties implied by the convexity of the curve.

3.2 Arbitraging on CPAMM

Arbitrageurs play a key role in CPAMM. The pool price of BTC in terms of USD is the amount of USD swapped out per unit of BTC swapped in, for an infinitesimal amount:

$$P = -\frac{\partial Y}{\partial X} = \frac{Y}{X} \tag{2}$$

where the second equality follows from Equation 1.

Since the pool price is entirely determined by relative reserves of the two assets, it can deviate from the true price, i.e. what is observed in the market. However, large deviation represents an arbitrage opportunity, and anyone can make a profit by buying at the low price and then selling at the high price. As arbitrageurs buy (sell) the lower (higher) pool price relative to the market price, they increase (decrease) the reserve of BTC, which then push up (down) the pool price towards the market price.

We assume that no arbitrage holds, i.e. arbitrageurs do a good job and pool price does not deviate much from market price. This is a mild assumption on market efficiency that is very common in the finance literature. Angeris and Chitra [2020] and and Lehar and Parlour [2021] show theoretical and empirical support of this assumption.

3.3 Liquidity provision on CPAMM

In order for traders or arbitrageurs to swap one asset for the other at short notice, there need to be positive reserves of both assets, and that is where liquidity providers (LPs) come in. In CPAMM, liquidity provision means depositing or withdrawing equal value of both assets, so that pool price remains unchanged. To see why, first observe that the two assets always have equal value: $PX = Y$, which follows directly from Equation 2. Therefore, any addition or subtraction of reserves need to happen to both assets in equal value.

When an LP deposits liquidity, she gets LP tokens that represent their share of the pool. The number of LP tokens outstanding equals to $\sqrt{K}$.

One useful measure of liquidity is total value locked (TVL). Using the second token as the numeraire, TVL is defined as:

$$TVL = PX + Y = 2PX = 2Y \tag{3}$$

where the second and third equality follows from Equation 2. At times when I need to compare liquidity across pools, I express TVL in terms of U.S. dollar:

$$TVL^{\$} = P_X^{\$} X + P_Y^{\$} Y \tag{4}$$

3.3.1 Returns on liquidity provision (LP returns)

Why would anyone engage in liquidity provision, instead of just holding on to the assets? The answer is that they earn returns on liquidity provision (LP returns). I define LP returns as relative to the holding returns (i.e., the opportunity cost of providing liquidity), and it consist of two components: fee revenues and adverse selection losses.

Fee revenues. As the price of consuming liquidity, each trade needs to pay a fee equal to a constant fraction of the amount of asset swapped in. This fee is automatically charged and deposited into the pool along with the asset swapped in.

Expressing fee revenues as returns at the pool level requires some work. For example, suppose that the pool originally has certain amount of liquidity, some trades occur, then a large amount of new liquidity is deposited to the pool, and then a large number of trades occur. I cannot attribute all of the trading fees to the original liquidity. The original liquidity only gets a share of the fees associated with the second wave of trading, depending on its size relative to the new liquidity. At the same time, the original liquidity does get all of the fees associated with the first wave of trading.

To address this, I focus on the fee revenues to a hypothetical marginal LP. Specifically, I calculate the hypothetical return of an LP that deposits a small fraction of the pool (e.g. 1%) at the beginning and withdraw that liquidity at the end. If there are inflows of new liquidity, the marginal LP's share of the pool shrinks, and her share of the fee revenues is proportionally reduced.

Formally, given fee rate f (e.g. 0.3% for all Uniswap V2 pools), amount of fee revenues for a token in a given time period t can be easily calculated from amount of that token swapped in:

$$fSwapIn_t$$

For a marginal LP that entered the pool at the end of t, the amount of fee revenues accrued to her in a future time period $t+s$ is given by:

$$AccruedFee_{t,t+s} = fSwapIn_t \frac{LP^{\#}_{t+s-1}}{LP^{\#}_t}$$

where $LP^{\#}_t$ denotes number of LP tokens outstanding at the end of t. Finally, the return on fee revenues from t to $t+h$ is:

$$R^{fee}_{t,t+h} = \frac{(\sum_{s=1}^{h} AccruedFee^{A}_{t,t+s})P_{t+h} + (\sum_{s=1}^{h} AccruedFee^{B}_{t,t+s})}{TVL_t} \tag{5}$$

Note that fee revenues are never negative – the worst that can happen is that there is no trade and hence zero fee revenue.

Adverse selection losses. Absent fees, LPs incur adverse selection losses whenever price changes. To see why, suppose, without loss of generality, that the pool price equates the market price initially, and then the market price increases. Pool price does not adjust automatically – arbitrageurs need buy from the pool (and then sell to the market for a profit) to push up the pool price. However, in this process, LPs incur losses by selling at the stale and low pool price instead of the fair and high market price. Another way to understand this is that, since this is a zero-sum game, any gains to the arbitrageurs are losses to the LPs.

These adverse selection losses are commonly referred to by practitioners as “impermanent losses”, since the losses go back to zero if prices revert, hence impermanent. In reality, cryptocurrency prices rarely mean-revert and losses are mostly permanent.

Assuming no arbitrage – that is, pool price equal to market price – adverse selection losses become a deterministic function of price changes. To see this, first note that given any pool price P', pool reserves X and Y and TVL are completely pinned down, absent any deposit or withdrawal of liquidity and hence no change in K:

$$X(P') = \sqrt{K/P'},\ Y(P') = \sqrt{KP'},\ TVL(P') = 2\sqrt{KP'}$$

In this world where fee revenues are assumed to be zero, LP returns are:

$$R^{AS} = \frac{TVL(P') - (P'X(P) + Y(P))}{TVL(P)} = \frac{2\sqrt{P'/P} - (P'/P + 1)}{2} \tag{6}$$

where the first term in the numerator is the new TVL and the second term is the holding value (i.e. the opportunity cost of providing liquidity). It is easy to show that R^{AS} is never positive. I call $-R^{AS}$ adverse selection losses, which are never negative.

Figure 3 visualize the mechanical relationship between adverse selection losses and prices changes. adverse selection losses are bigger for bigger price changes. Therefore, pools where token prices are more volatile incur higher risk of adverse selection losses to the liquidity provider.[4]

Net returns. From an LP's perspective, what matters is the net effect of fee revenues and adverse selection losses:

$$R^{net} = R^{fee} + R^{AS} \tag{7}$$

In what follows, all LP returns are subtracted by the risk-free rates using data on U.S. Treasury bills.

[4]See Aigner and Dhaliwal [2021] for a more in-depth discussion.

4 Data and Factor Models

4.1 Data sources

I study Uniswap V2, the largest CPAMM and also the largest DEX as of May 2021 (before Uniswap V3 was launched). Figure 1 shows aggregate dollar trading volume and aggregate dollar liquidity over time for Uniswap V2 vs other major DEXs. Uniswap V2 had $5.0 billions of total value locked as of May 2021 and supported $261 billions of trading volume in the preceding year.

I obtain transaction-level data on Uniswap V2 from BigQuery and Dune Analytics. I then aggregate transactions to daily, weekly, or monthly level to perform my empirical analyses.

Daily prices on cryptocurrencies are from CoinMarketCap and obtained through Yahoo Finance API.

Risk-free rates are estimated using data on U.S. Treasury bills from the CRSP US Treasury Database.

4.2 Factor models

High returns can be due to either high alpha or high risk. To capture risk, I use the three-factor model from LIU et al. [2022], who show that market, size and momentum capture the cross section of cryptocurrency returns. Specifically, the three-factor model is:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i^{CMKT} CMKT_t + \beta_i^{CSMB} CSMB_t + \beta_i^{CMOM} CMOM_t + \epsilon_{i,t}$$

where $R_{i,t}$ is return on the asset of interest (e.g., net return on a liquidity pool), $R_{f,t}$ risk-free return, $alpha_i$ the level of abnormal return not explained by the risk factors, $CMKT$ the size-weighted aggregate return on all cryptocurrencies, $CSMB$ the return difference between cryptocurrencies with small vs large market capitalization, $CMOM$ the return difference between past winners vs losers, and β exposure to the risk factors.

We are interested in $alpha_i$, the level of abnormal return not explained by the risk factors, which can be estimated through a OLS regression.

5 Persistent Abnormal LP Returns

In this section, I show persistent abnormal return differences across liquidity pools. First, I show that while some pools have highly positive risk-adjusted average returns, others have highly negative risk-adjusted average returns. Then, focusing on a subset of pools, I calculate their options-implied liquidity premia, which deviate persistently from actually realized LP revenues. Lastly, using a sorted portfolio approach, I show that pools with higher (lower) past returns continue to have significantly higher risk-adjusted future returns.

5.1 Average returns

I first examine average returns across pools. I focus on raw averages, risk-adjusted alphas with respect to common risk factors, as well as Sharpe ratios that adjust for idiosyncratic volatility. To minimize the influence of outliers, I restrict to pools with at least 40 weeks of returns and winsorize returns at the 99th percentile for each cross section. I also exclude small pools and restrict to pools with at least \$100K in total value locked.

Figure 4 shows the results. Most pools have negative average returns, either in raw or factor-adjusted, consistent with the findings in [Loesch et al., 2021]. More importantly, many pools have large positive or negative returns that are hard to reconcile with risk-based explanations. Specifically, some pools have average returns as large as (negative) 50% per annum, even after adjusting for common cryptocurrency factors, and the Sharpe ratios are as large as (negative) 0.50. For reference, the Sharpe ratio of weekly Bitcoin (US equity) return is 0.23 (0.13) over the same period.

The cross-sectional return differences indicate misallocation: pools with large positive returns are supposed to see higher inflows of liquidity that will dilute the fee revenues and decrease overall returns, while pools with large negative returns should see the opposite. As opposed to liquidity misallocation, one alternative explanation is risk-based: some pools require abnormally high or low expected returns, due to the risks involved. This argument seems weak in the presence of negative average returns (risk premia should generally be positive) and as I account for common risk factors in the cryptocurrency space. In the next section, I use options data to provide explicit estimates of risk premia and show that large abnormal returns remain.

5.2 Options-implied liquidity premium

As shown in Section 1, adverse selection losses are a deterministic function of the relative price of the two tokens, assuming no arbitrage. This means that I can derive the present value of adverse selection losses in a a Black-Scholes option-pricing framework. Indeed, as shown in Figure 3, adverse selection losses resemble the payoff of a straddle, i.e. the combination of a call and a put.[5]

Formally, the relative price of the two tokens P_t follows a geometric Brownian motion under the risk-neutral measure:

$$dP_t = (r - \frac{1}{2}\sigma^2)P_t dt + \sigma P_t dW_t^Q.$$

[5]See Evans [2020] for more.

Terminal price can be written as:

$$P_T = P_0 \exp[(r - \frac{1}{2}\sigma^2)T + (\sigma\sqrt{T})\epsilon^Q].$$

The present value of an instrument whose payoff is solely a function of terminal price $f(P_T)$ is then:

$$V = e^{rT} E^Q[f(P_T)].$$

In particular, the present value of adverse selection losses is:

$$V^{AS} = e^{rT} E^Q[R_T^{AS}], \tag{8}$$

where R^{AS} is given in Equation 6.

One key input to this framework is volatility. I use the options-implied volatilities calculated by T3 Index. Options-implied volatility is the volatility that matches theoretical Black-Scholes option prices with the ones actually observed in the market. In other words, for call (put) options, whose payoffs are given by $f^{call}(P_T) = \max(0, P_T - S)$ ($f^{put}(P_T) = \max(0, S - P_T)$), where S stands for strike price, options-implied volatility is the value of σ such that V in My framework above matches the actually observed market value of these call (put) options. What is special about implied volatility is that it has market premium baked in. Using option-implied volatility enables to answer the following question: what would be the present value of options that replicate adverse selection losses?

The other inputs to this framework are as follows. Since the options-implied volatilities from T3 Index are 30-day, I choose T to be 30-day. r is one-month U.S. Treasury Bill rate.

We derive V^{AS} using Monte Carlo. Specifically, I simulate a large sample of ϵ^Q, which give us a large sample of P_T and in turn R_T^{AS}. I take the average of the sample of R_T^{AS} to obtain $E^Q[R_T^{AS}]$ and discount it to get V^{AS}.

Figure 5 shows the results. The graphs plot realized and options-implied LP returns for the BTC-USD pools and the ETH-USD pools, averaged across the three stablecoins: USDC, USDT and DAI.[6] Realized fee revenues are noticeably bigger than options-implied present value of adverse selection losses. In other words, it is significantly more profitable to provide liquidity to BTC-USD (ETH-USD) pools than selling options on BTC (ETH), even though the two investments have identical exposure adverse selection losses. Moreover, there is a visible positive correlation between realized fee revenues and realized adverse selection losses, suggesting that fee revenues are a good hedge against adverse selection losses and therefore should require a lower premium – in other words, the positive gap between actual vs options-implied LP returns is bigger than it looks. The results here are stronger evidence of misallocation: there has been insufficient liquidity on BTC-USD and ETH-USD pools, since liquidity provision would have earned much higher returns than what is implied by the options market.

5.3 Sorted portfolio approach

I use a sorted portfolio approach to further demonstrate the persistence of cross-sectional return differences. At the beginning of each week, liquidity pools are sorted into five quintiles based on their realized net returns over the past week. A TVL-weighted portfolio is formed for pools in each quintile. The portfolios are held for one week and then dissolved.

To avoid outlier effects, I restrict to pool-week observations with at least $100K TVL. Furthermore, I require at least 50 such pools to be available, or otherwise discard the time period. I winsorize returns at the 99th percentile for each cross section.

Table 1 shows the results. Pools with high (low) past returns continue to earn high (low) returns. Specifically, pools in the highest (lowest) quintile of past returns earn a risk-adjusted

[6]Technically, the pools are WBTC and WETH, which have near identical prices as BTC and ETH.

return of 0.42% (-0.70%) in the week that follows, or 25.64% (-35.51%) per annum. The future return difference between the top quintile and the bottom quintile of past return is 1.12% per week on a risk-adjusted basis, or 93.75% per annum.

6 Unsophisticated LP Flows

The previous section shows predictably persistent cross-pool return differences. Why doesn't liquidity adjust? In this section, I directly examine LPs' actions to deposit or withdraw liquidity. I find that LPs are insensitive to net returns. Instead, LPs chase fee revenues, which are prominently marketed as APY on the online platforms, but ignore adverse selection losses, which are not. As a result, pools with higher fee revenues attract more liquidity, even though they tend to have higher adverse selection losses and insignificant returns on net.

6.1 Measuring LP flows

I measure LP flows to a pool p in from time t to time $t+1$ as:

$$Flow_{p,t+1} = \frac{1}{2}\frac{Deposit^{A}_{p,t+1} - Withdrawal^{A}_{p,t+1}}{Reserve^{A}_{p,t}} + \frac{1}{2}\frac{Deposit^{B}_{p,t+1} - Withdrawal^{B}_{p,t+1}}{Reserve^{B}_{p,t}} \quad (9)$$

where $Deposit_{t+1}$ denotes total amount of token deposited into the pool from time t to $t+1$, $Withdrawal_{t+1}$ total amount of token withdrawn from the pool from time t to $t+1$, and $Reserve_t$ reserve of token at the end of time t.

6.2 Determinants of LP flows

What determines LP flows? I borrow the regression framework that is typical in the large literature on mutual fund flows [e.g., Barber et al., 2016]:

$$Flow_{p,t+1} = \beta_1 R^{fee}_{p,t} + \beta_2(-R^{AS}_{p,t}) + \gamma Controls_{p,t} + FE + \epsilon_{p,t+1} \tag{10}$$

where $Flow$, R^{fee}, and R^{AS} are previously defined in Equations 9, 5, and 6, controls include past LP flows, log TVL and indicators for liquidity mining rewards, and FE includes time fixed effects and/or pool fixed effects. I change the sign of R^{AS} so that it represents adverse selection losses.

I construct My sample as follows. I start with the universe of all Uniswap V2 pools. As small pools can have abnormally large flows that bias the results, I exclude pools that never reach $1 million in TVL and exclude observations where TVL is less than $10,000. I focus on the time period between May 2020 (after Uniswap V2 was launched) and April 2021 (before Uniswap V3 was launched that renders V2 obsolete). I focus on the weekly frequency but show robustness with daily and monthly frequency. I winsorize all variables at the 99th percentile to mitigate the influence of outliers.

Table 2 shows the baseline results. Column 1 and 2 examine the relationship between future liquidity flow and past net return. The coefficient on past net return is zero, meaning that future liquidity flow does not respond to past net return at all. This provides a consistent explanation with my previous results on persistent return differences – pools with high past net returns do not experience higher inflows of liquidity and therefore continue to provide high future net returns.

Column 3 and 4 breaks down net return into its two components, fee revenue and adverse selection loss. The coefficient on past fee revenues is positive and significant. According to column 4, 1 p.p. increase in fee revenues over the past week is associated with 1.57 p.p. higher flow of liquidity into the pool over the subsequent week. This is economically significant, since the standard deviation of weekly fee revenues and weekly LP flows are respectively 0.69 p.p. and 8.29 p.p.

The coefficient on past adverse selection losses is small and not significant from zero. Without fixed effects, LPs do appear to respond to past adverse selection losses – this is consistent with the findings in Lehar and Parlour [2021], Capponi and Jia [2021]. However, there could be persistent unobserved factors that affect both flows and adverse selection losses. For example, tokens that are popular can have both lower volatility (and hence lower adverse selection losses) and high flows, on average. Therefore, I argue that it is important to include pool fixed effects, which can absorb these persistent unobserved factors. With pool fixed effects, I effectively focus on news to fee revenues and news to adverse selection losses, and I find that LPs only respond to the former.

Among the control variables, the coefficient on past flows is highly significant. This is consistent with the literature on mutual fund flows that retail investors flows are highly persistent and affirms the importance of controlling for past flows.[7].

I conduct a battery of robustness checks in Table 3. Panel B and C conduct the same analysis but at daily or monthly frequency. The key results remain the same. Panel D and E perform the same analysis but expands to smaller pools or restricts to bigger pools – the coefficient on past fee revenues becomes bigger as I restrict to bigger pools, implying that there is more chasing of past fee revenues in more popular pools and assets.

Instead of LP returns, I can instead focus on the primitives that determine LP returns, as is done in Capponi and Jia [2021]. Since fee revenues are simply a fraction of trading volumes,

[7] This is another important difference between my regressions and the one in Lehar and Parlour [2021] and Capponi and Jia [2021]

I can alternatively focus on trading volumes, defined as:

$$Volume_{p,t} = \frac{1}{2}\frac{SwapIn^A_{p,t}}{Reserve^A_{p,t-1}} + \frac{1}{2}\frac{SwapIn^B_{p,t}}{Reserve^B_{p,t-1}}$$

As shown in Section 3, assuming no arbitrage, adverse selection losses are a deterministic function of price change. Therefore, I can alternatively focus on volatility, defined as:

$$Volatility_{p,t} = StandardDeviation^t_{\tau=t-1}(\log P_\tau)$$

This measurement is invariant with respect to the choice of base currency (i.e. whether token A or token B is the numeraire) and to scalar multiplication of token values. Panel A of Table 3 shows that my main results remain unchanged.

Instead of the pooled regression approach, I can use the Fama-MacBeth approach [Fama and MacBeth, 1973] that is more resistant to the time-varying weighting problem [Ben-David et al., 2021]. The procedure is to run the same regression in Equation 10 but only using the cross section for each time period. Figure 6 shows the results. Flows respond positively to past fee revenues consistently over time, up until the launch of Uniswap V3 (when V2 became obsolete). In contrast, flows do not respond to past adverse selection losses in a consistent manner.

So far, I have focused on realized LP returns. However, if LPs are sophisticated, they should instead focus on expected LP returns. Since My baseline specification is with pool fixed effects and focuses on news to LP returns, realized LP returns are indeed the best predictors of future LP returns. However, for volatilities, options-implied volatilities are known to have forecasting power beyond what realized volatilities can do, so I test whether LP flows respond to changes in implied volatilities. Table **??** shows that LP flows actually respond positively to increase in implied volatility. I hypothesize the following explanation: implied volatility is high when sentiment is high, and when sentiment is high flow is high.

6.3 Why do LPs ignore adverse selection losses?

Why do LP flows chase fee revenues but not adverse selection losses? Past fee revenues are prominently marketed as "APY" on these online platforms. Indeed, Figure 7 shows what regular users would see on the websites. APY (annual percentage yield) is usually associated with deposit-like investment products to imply safe guaranteed returns. However, fee revenues from liquidity pools are highly volatile, what was realized in the past does not forecast exactly what is the future return. More importantly, adverse selection losses are in the calculation of APY and not displayed anywhere. I hypothesize that LPs are unsophisticated investors that are heavily influenced by the marketing.

6.4 The cost of ignoring adverse selection losses

Given the results above, one can still think of ways to justify LP rationality. In particular, LPs might chase past fee revenues more than past adverse selection losses, because the former is more persistent and hence more predictable than the latter. These arguments can indeed find justification in the data: the auto-correlations of fee revenues and adverse selection losses are 0.5 and 0.3 at weekly frequency, 0.4 and 0.2 at monthly frequency. However, one can already raise doubt on this line of reasoning by observing that fee revenues are not more persistent than adverse selection losses by as wide a margin as their regression coefficients are in Table 2.

In what follows, I show that LPs can derive significant gains by paying attention to adverse selection losses instead of chasing past fee revenues alone. First, I simulate the payoff of naively chasing past fee revenues without considering adverse selection losses, using the same sorted portfolio approach from before. Specifically, at the beginning of each week, I sort all Uniswap V2 pools into five quintiles based on their past fee revenues, form a TVL-weighted portfolio for pools in each quintile, hold the portfolios for one week, and dissolve the portfolios.

Table 4 shows the results. Column 1 confirms that the sorting variable goes up monotonically from the bottom quintile to the top quintile. Column 4 shows the future fee revenue, which also goes up monotonically, but considerably less magnitude. To understand this dampened monotonicity, I use the fact that fee revenue is fee amount divided by liquidity amount, and examine the numerator and the denominator separately. Column 2 shows reversion in the numerator – pools with low past fee revenue see high growth in fee amount, whereas pools with high past fee revenue see large decline in fee amount. Column 3 shows that the denominator behaves as expected – pools with high past fee revenue experience large inflows of liquidity, whereas pools with low past fee revenue experience large outflows of liquidity. These two effects together lead to significant dampening of the monotonicity of fee revenues.

More importantly, Column 5 shows that pools with higher fee revenue also have higher adverse selection losses. When an asset is being traded a lot, that is usually when there is large price to the asset, either because of sentiment or new information. Column 6 shows that future net return is not significantly different from zero for all five quintiles, and there is no significant difference between the top quintile and the bottom quintile.

I now demonstrate the power of adding adverse selection losses into the decision-making process. After sorting pools into five quintiles by past fee revenues, I further sort the pools in each quintile into five sub-quintiles by past price volatility. This gives us a total of 25 portfolios. Table 5 shows the results. There is significant return difference between high-past-volatility pools and low-past-volatility pools, conditional on having the same past fee revenues. For each past-fee-revenue quintile, the low-past-volatility pools predictably earn significantly higher net returns than high-past-volatility pools. Even restricting to pools with the highest past fee revenues, future net returns vary from 0.42% on average for the low-past-volatility pools to -1.28% on average for the high-past volatility pools. Within the high past fee quintile, the net return difference between the low-volatility pools and the high-volatility pools is 1.70% per week, or 140.26% per annum.

6.5 Counterfactual liquidity

What would happen if LPs are equally sensitive to adverse selection losses as they are to fee revenue? Since adverse selection losses are always positive, the counterfactual liquidity flows should be much lower. Lower liquidity leads to lower trading activities and hence lower fee revenue, which leads to further declines in liquidity provision. In equilibrium, there should be substantially lower liquidity, due to this multiplier effect.

In this current version of the paper, I provide a naive derivation of counterfactual liquidity without the multiplier effect. Specifically, I calculate counterfactual outflows due to adverse selection losses as $\beta_1(-R^{AS}_{p,t})$, where β_1 is LPs' sensitivity to fee revenue (1.57 according to Table 2) and $(-R^{AS}_{p,t})$ is realized adverse selection losses. I calculate counterfactual outflows for each pool in each period. The results show that there would have been 34% cumulative outflows (or 34% less inflows) to Uniswap V2 pools if LPs were as sensitive to adverse selection losses as they are to fee revenue.

7 Efficiency with Sophisticated LPs

What would happen to LP returns and LP flows, if LPs become more sophisticated? I get a glimpse from the launch of Uniswap V3. After its launch in May 2021, Uniswap V3 dethroned Uniswap V2 and became the default dapp on Uniswap's website. As a result, it attracts most of the liquidity since May 2021. It is reasonable to assume that remaining active liquidity providers on Uniswap V2 are more sophisticated.

Figure 5 shows that realized fee revenues are much closer to options-implied LP premia after the launch of Uniswap V3.

Column 2 of Table 1 shows that return predictability decreases significantly after the launch of Uniswap V3.

Table 7 runs the following regression:

$$\begin{aligned} Flow_{p,t+1} = \beta_1 R^{fee}_{p,t} + \beta_1' V3_t \times R^{fee}_{p,t} + \beta_2(-R^{AS}_{p,t}) + \beta_2' V3_t \times (-R^{AS}_{p,t}) \\ + \gamma Controls_{p,t} + FE + \epsilon_{p,t+1} \end{aligned} \tag{11}$$

The results in Table 7 that flows are much less responsive to past fee revenues and much more sensitive to adverse selection losses after the launch of Uniswap V3.

8 Conclusion

In this paper, I provide an in-depth empirical analysis of LP returns and LP flows on one of the largest decentralized exchanges. My main finding is that DEX LPs behave in a very unsophisticated manner. They chase fee revenues – the part of return that is prominently marketed as APY – but ignore adverse selection losses – the part of return that is implicit and rarely displayed. As a result, liquidity is misallocated. Pools with higher realized net returns do not attract sufficient liquidity and continue to yield higher net returns in the future. On the other hand, pools with higher fee revenues attract significantly more liquidity, even though these pools tend to also have higher adverse selection losses and zero expected returns on net.

References

A. Adams and G. Liao. When uniswap v3 returns more fees for passive lps. May 2022. URL `https://uniswap.org/SuperiorReturnsForLiquidityProviders.pdf`.

A. A. Aigner and G. Dhaliwal. Uniswap: Impermanent loss and risk profile of a liquidity provider. 2021. doi: 10.13140/RG.2.2.32419.58400/6. URL `http://rgdoi.net/10.13140/RG.2.2.32419.58400/6`.

G. Angeris and T. Chitra. Improved price oracles. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. ACM, oct 2020. doi: 10.1145/3419614.3423251.

G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of uniswap markets, 2019. URL `https://arxiv.org/abs/1911.03380`.

J. Aoyagi. Liquidity provision by automated market makers. *ERN: Other IO: Theory (Topic)*, 2020.

J. Aoyagi and Y. Ito. Coexisting exchange platforms: Limit order books and automated market makers. *Entrepreneurship & Economics eJournal*, 2021.

B. M. Barber, X. Huang, and T. Odean. Which Factors Matter to Investors? Evidence from Mutual Fund Flows. *The Review of Financial Studies*, 29(10):2600–2642, 06 2016. ISSN 0893-9454. doi: 10.1093/rfs/hhw054. URL `https://doi.org/10.1093/rfs/hhw054`.

I. Ben-David, J. Li, A. Rossi, and Y. Song. What Do Mutual Fund Investors Really Care About? *The Review of Financial Studies*, 35(4):1723–1774, 07 2021. ISSN 0893-9454. doi: 10.1093/rfs/hhab081. URL `https://doi.org/10.1093/rfs/hhab081`.

J. B. Berk and J. H. van Binsbergen. Assessing asset pricing models using revealed preference. *Journal of Financial Economics*, 119(1):1–23, 2016. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2015.08.010. URL `https://www.sciencedirect.com/science/article/pii/S0304405X1500149X`.

L. Calvet, C. Célérier, P. Sodini, and B. Vallee. Can Security Design Foster Household Risk-Taking? CEPR Discussion Papers 14955, C.E.P.R. Discussion Papers, June 2020. URL `https://ideas.repec.org/p/cpr/ceprdp/14955.html`.

A. Capponi and R. Jia. The adoption of blockchain-based decentralized exchanges, 2021. URL `https://arxiv.org/abs/2103.08842`.

L. W. Cong, X. Li, K. Tang, and Y. Yang. Crypto Wash Trading. Papers 2108.10984, arXiv.org, Aug. 2021. URL `https://ideas.repec.org/p/arx/papers/2108.10984.html`.

C. Célérier and B. Vallée. Catering to Investors Through Security Design: Headline Rate and Complexity*. *The Quarterly Journal of Economics*, 132(3):1469–1508, 02 2017. ISSN 0033-5533. doi: 10.1093/qje/qjx007. URL `https://doi.org/10.1093/qje/qjx007`.

A. Evans. Liquidity provider returns in geometric mean markets, 2020. URL https://arxiv.org/abs/2006.08806.

A. Evans, G. Angeris, and T. Chitra. Optimal fees for geometric mean market makers, 2021. URL https://arxiv.org/abs/2104.00446.

E. F. Fama and J. D. MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, 1973. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1831028.

S. Foley, J. R. Karlsen, and T. J. Putniņš. Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies? *The Review of Financial Studies*, 32(5): 1798–1853, 04 2019. ISSN 0893-9454. doi: 10.1093/rfs/hhz015. URL https://doi.org/10.1093/rfs/hhz015.

C. Harvey, A. Ramachandran, and J. Santoro. *DeFi and the Future of Finance*. Wiley, 2021. ISBN 9781119836018. URL https://books.google.com/books?id=CGM4EAAAQBAJ.

L. Heimbach, Y. Wang, and R. Wattenhofer. Behavior of liquidity providers in decentralized exchanges, 2021. URL https://arxiv.org/abs/2105.13822.

L. Heimbach, E. Schertenleib, and R. Wattenhofer. Risks and returns of uniswap v3 liquidity providers, 2022. URL https://arxiv.org/abs/2205.08904.

A. Lehar and C. A. Parlour. Decentralized exchanges. *Investments eJournal*, 2021.

Y. LIU, A. TSYVINSKI, and X. WU. Common risk factors in cryptocurrency. *The Journal of Finance*, 77(2):1133–1177, 2022. doi: https://doi.org/10.1111/jofi.13119. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.13119.

S. Loesch, N. Hindman, M. B. Richardson, and N. Welch. Impermanent loss in uniswap v3, 2021. URL https://arxiv.org/abs/2111.09192.

M. Neuder, R. Rao, D. J. Moroz, and D. C. Parkes. Strategic liquidity provision in uniswap v3, 2021. URL https://arxiv.org/abs/2106.12033.

Y. Song. The mismatch between mutual fund scale and skill. *The Journal of Finance*, 75(5): 2555–2589, 2020. doi: https://doi.org/10.1111/jofi.12950. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12950.

Figures

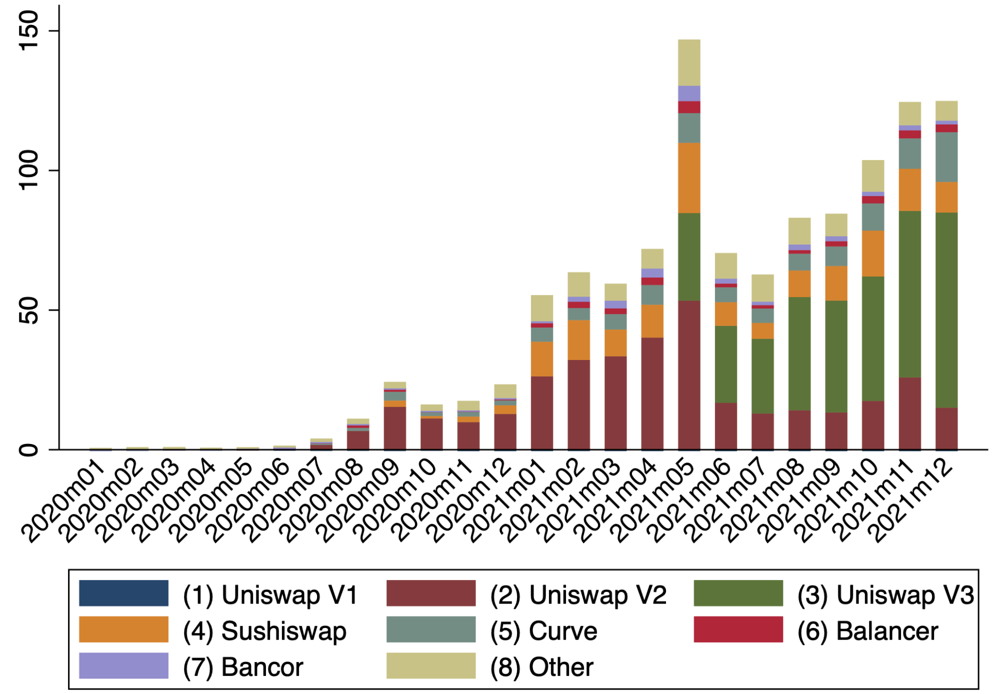

Panel A: Aggregate Trading Volume

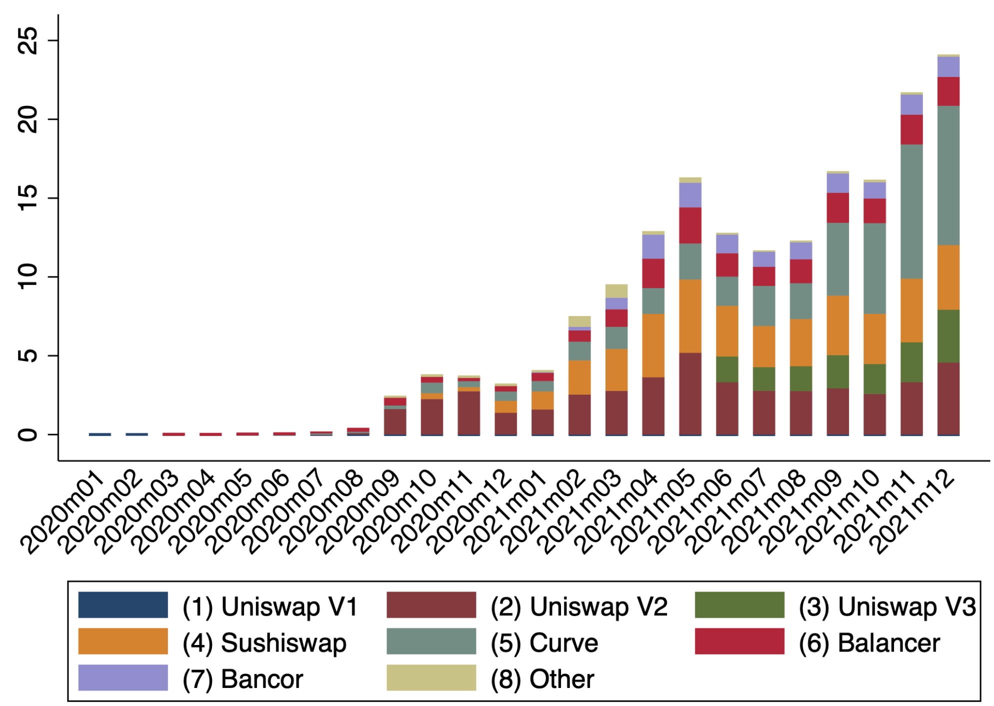

Panel B: Aggregate Liquidity

Figure 1: **Aggregate Trading Volume and Liquidity on Decentralized Exchanges**. Data are from Dune Analytics (dex.trades and dex.liquidity).

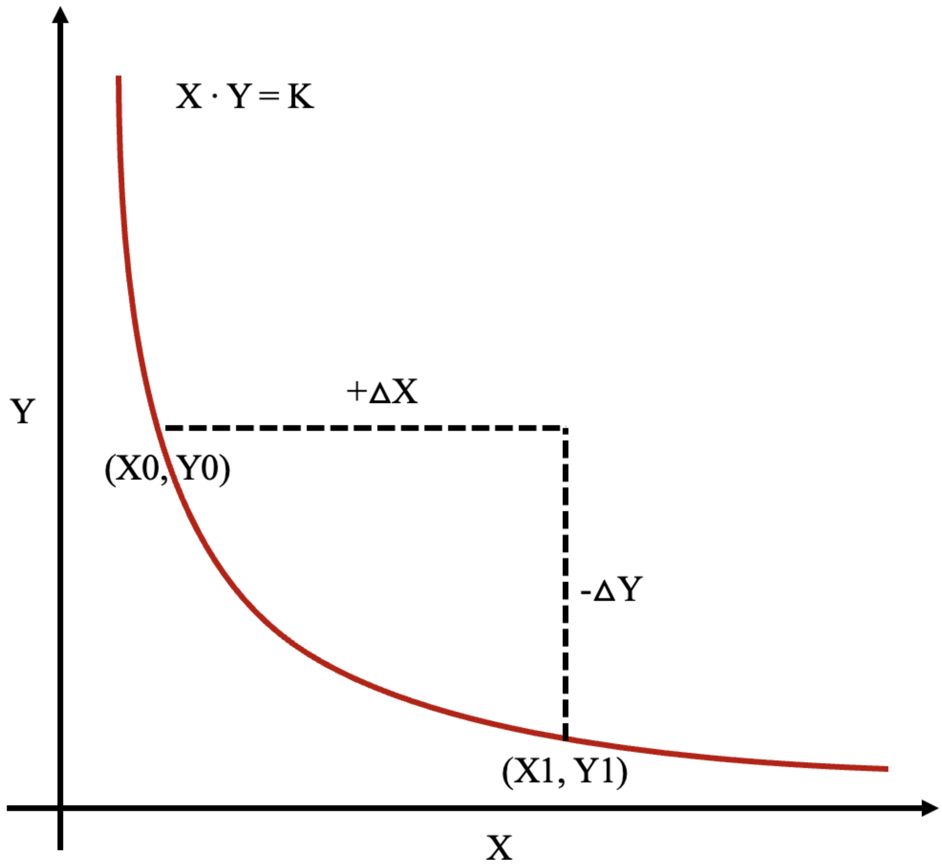

Figure 2: **Constant-Product Automatic Market Maker**. X-axis is reserve of one token (e.g. BTC) and y-axis is reserve of the other token (e.g. USD). The solid line plots Equation 1. The dash lines show what would be the amount of the second token swapped out (ΔY) if certain amount of the first token is swapped in (ΔX).

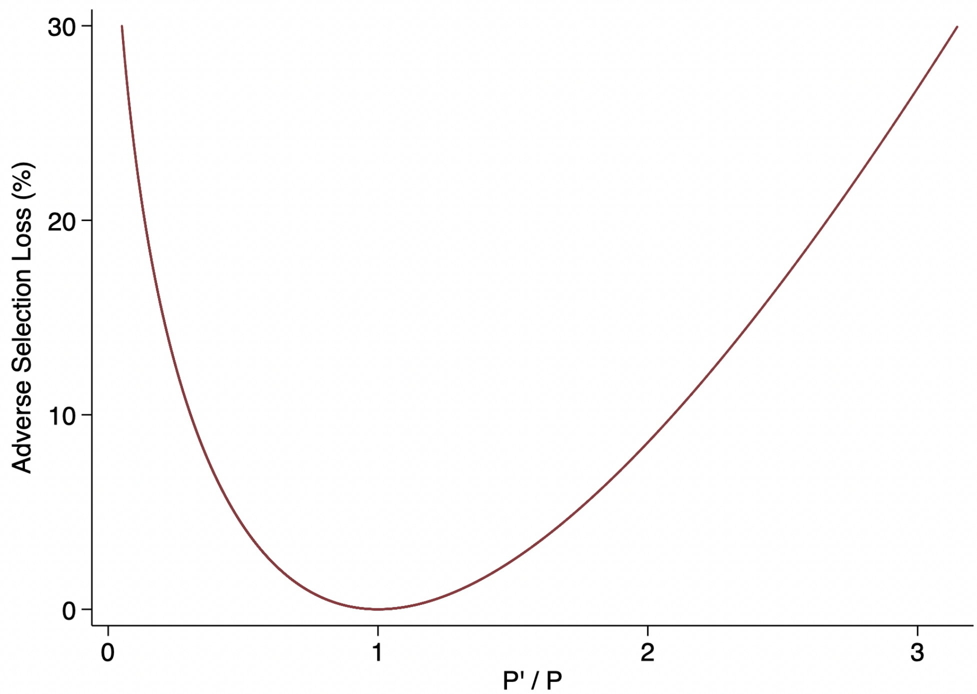

Figure 3: **Adverse Selection Loss**. X-axis is gross change in relative price of the two tokens in a CPAMM pool. Y-axis is adverse selection loss, i.e. the opportunity cost of providing liquidity relative to holding on to the two tokens, assuming that fee revenue is zero. The red line plots the mechanical relationship between adverse selection loss and price change, given by Equation 6.

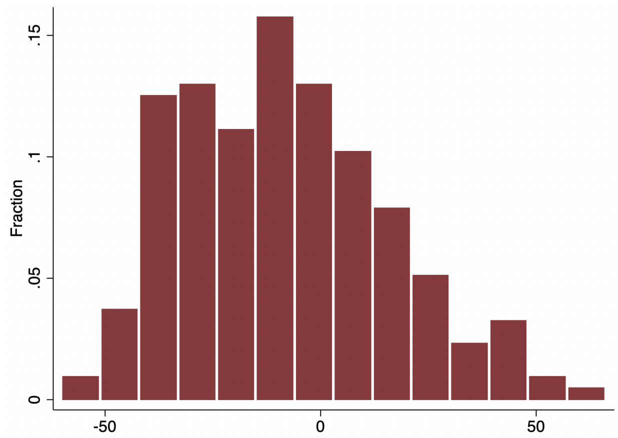

Panel A: Raw Average (%)

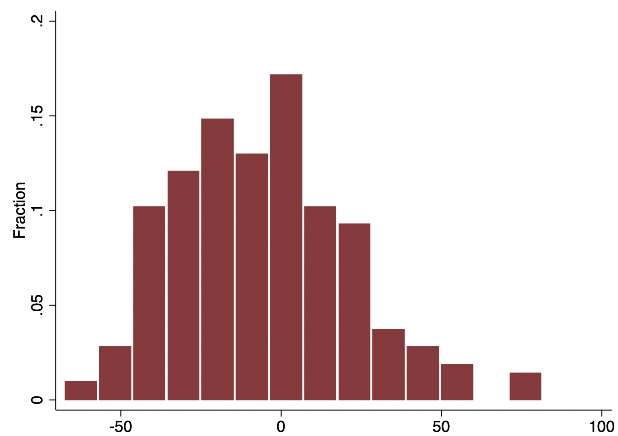

Panel B: Risk-Adjusted Alpha (%)

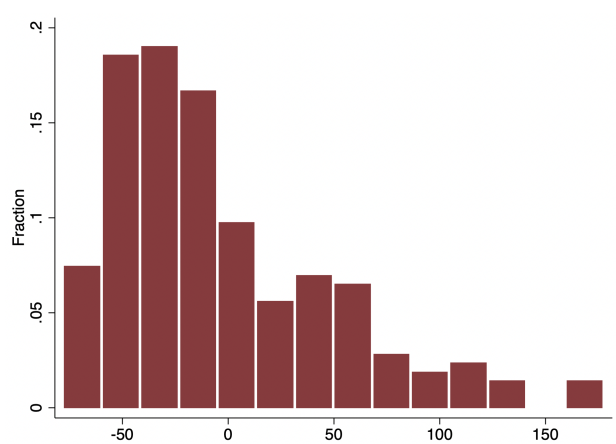

Panel C: Sharpe Ratio (%)

Figure 4: **Cross-Sectional Distribution of LP Net Returns**. The figures plot distributions of raw averages (Panel A), risk-adjusted alphas (Panel B), and Sharpe ratios (Panel C) of annualized weekly LP net returns across Uniswap V2 pools with at least \$100K total value locked.

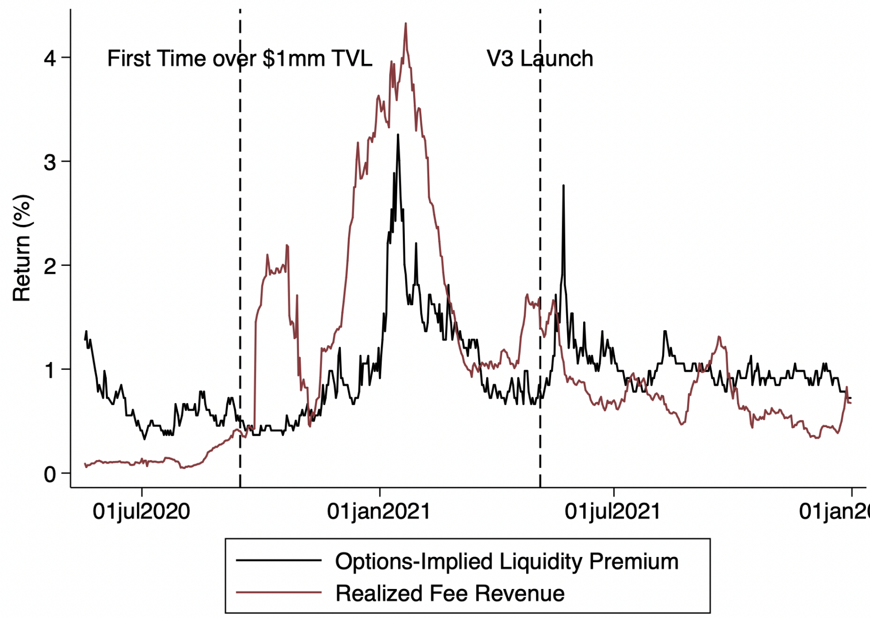

BTC-USD

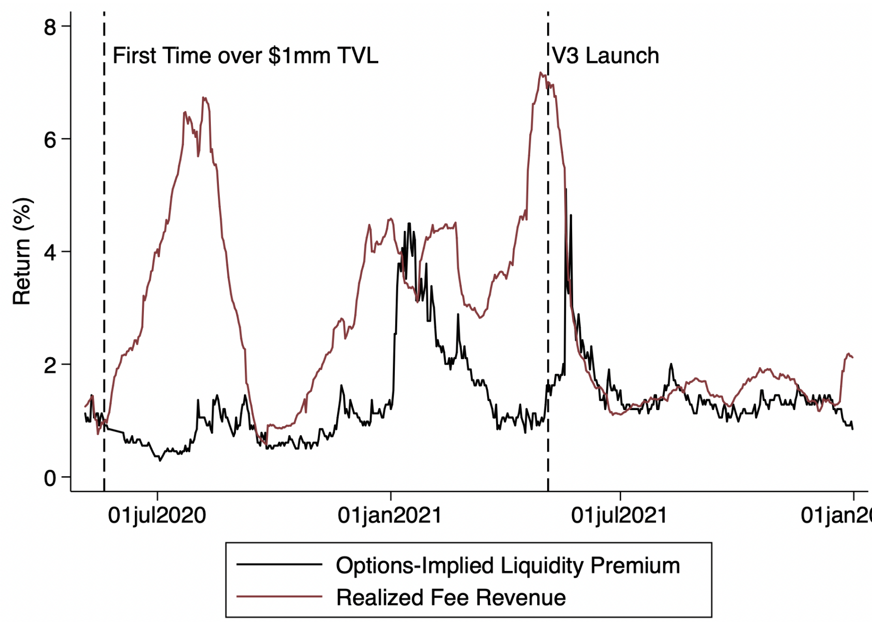

ETH-USD

Figure 5: **Options-Implied Liquidity Premium**. The lines show options-implied vs realized LP compensation for the BTC-USD pools and the ETH-USD pools. Options-implied premia are derived based on Equation 8 and Monte Carlo simulation. Start time corresponds to when the pool was created, the first dash vertical line corresponds to when there is at least $1 million in total value locked, and the second dash vertical line corresponds to when Uniswap V3 was launched.

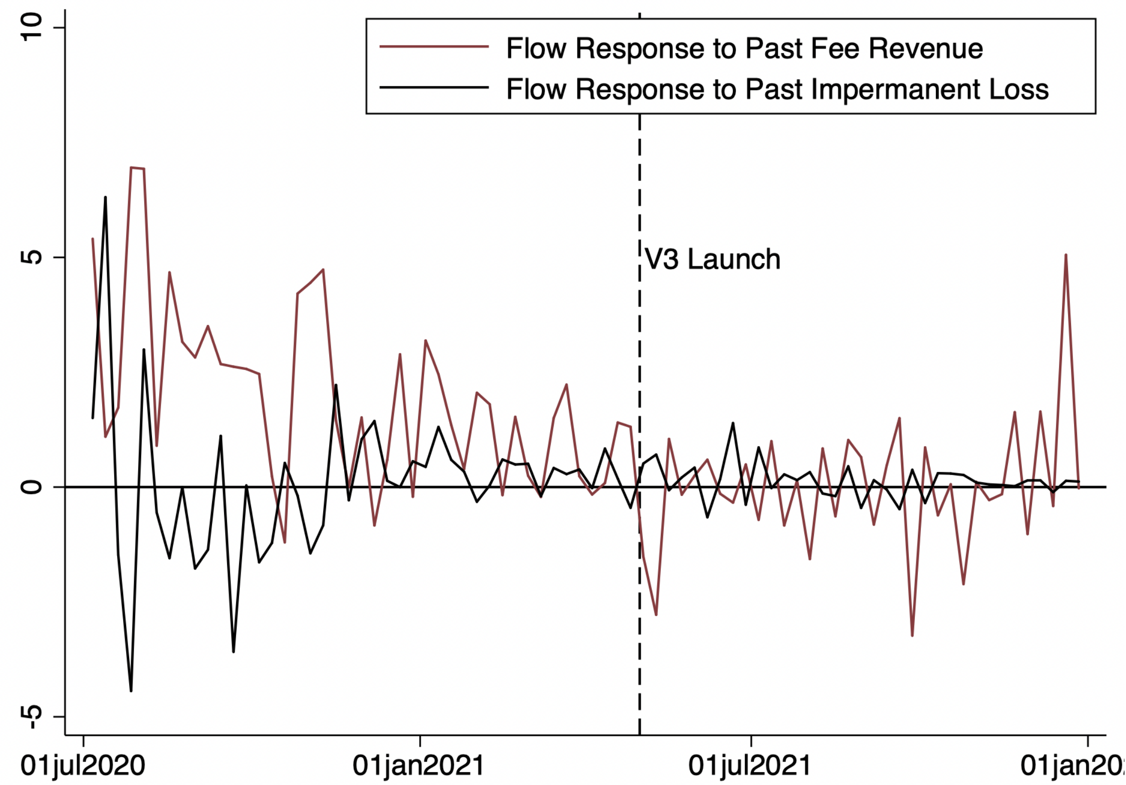

Figure 6: **Determinants of LP Flows, Fama-MacBeth**. The lines show estimated β_1 and β_2 for each week t from the flow regressions in Equation 10.

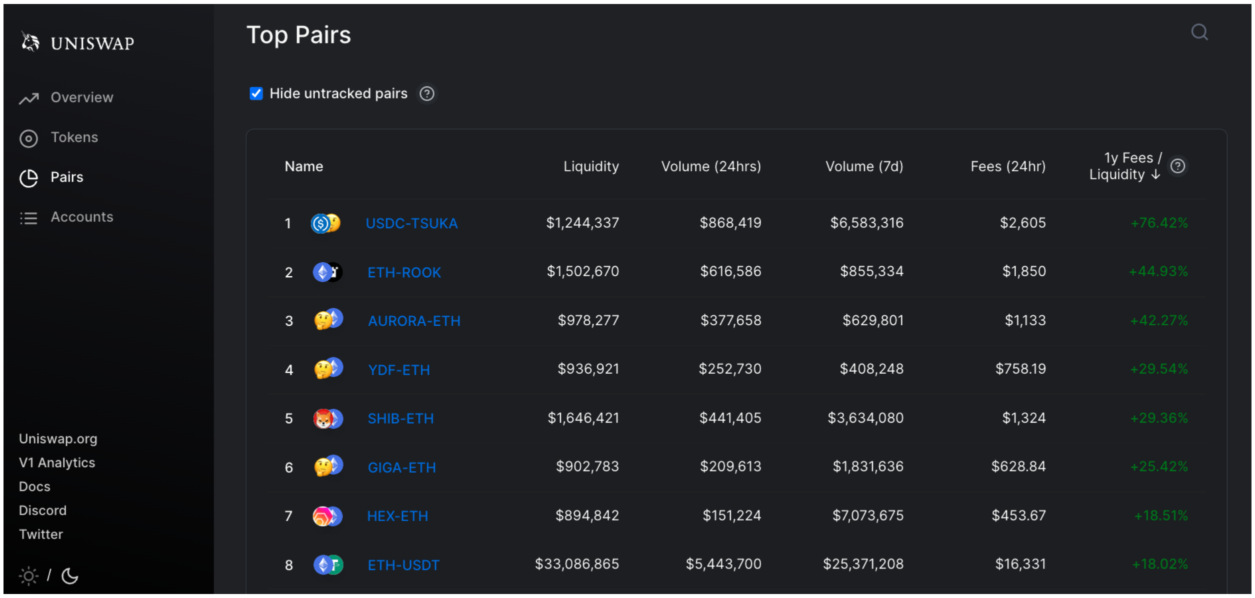

Uniswap V2

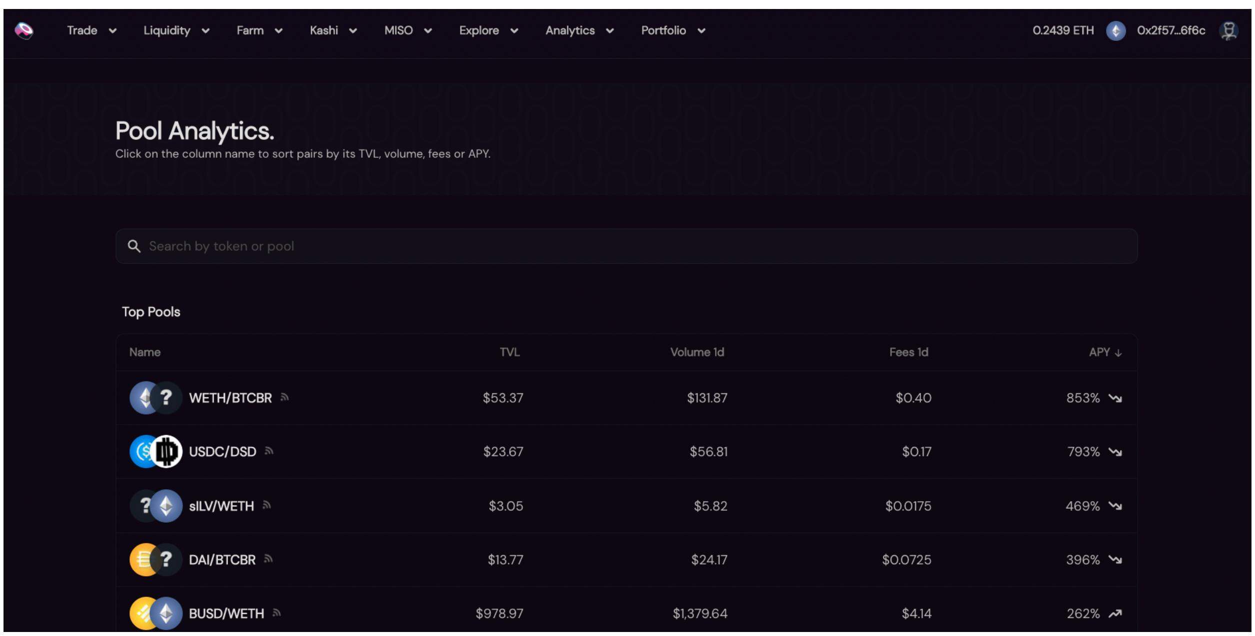

SushiSwap

Figure 7: **The Display of Fee Revenue as APY**. The top graph shows how fee revenue is displayed on Uniswap V2's website. The bottom graph shows how fee revenue is displayed on SushiSwap's website. These screenshots were taken on October 1, 2022.

Tables

Table 1: **Future Net Returns Sorted by Past Net Returns**. At the beginning of each week, I sort Uniswap V2 pools into five quintiles based on their realized net returns over the past week, form a TVL-weighted portfolio for pools in each quintile, hold the portfolios for one week, and then dissolve the portfolios. The table shows the mean (and its t-statistic in parentheses) of weekly returns for each portfolio, as well as of the difference between the top quintile and the bottom quintile.

	Future Net Return (week t to t+1, %)	
	Raw Average	Three-Factor Alpha
Quintile of Past Net Return (week t-1 to t)		
Low	-0.61***	-0.70***
	(-4.73)	(-3.94)
2	-0.17*	-0.22*
	(-1.80)	(-1.68)
3	0.01	0.05
	(0.28)	(0.93)
4	0.08*	0.12**
	(1.72)	(2.19)
High	0.37***	0.42***
	(6.34)	(5.89)
High – Low	0.98***	1.12***
	(10.26)	(9.01)

Table 2: **Determinants of LP Flows, Baseline**. The table shows regression results of Equation 10. The sample includes weekly observations of all Uniswap V2 pools with more than $100K in TVL from June 2020 (after Uniswap V2 was launched) to April 2021 (before Uniswap V3 was launched). t-statistics are reported in parentheses. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: -0.65, SD: 8.29			
Past Net Return (%, t-1 to t)	0.00 (0.02)	-0.01 (-0.17)		
Past Fee Revenue (%, t-1 to t)			2.21*** (6.26)	1.57*** (5.33)
Past AS Loss (%, t-1 to t)			-0.14 (-1.53)	-0.08 (-1.01)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Week FE	Y	Y	Y	Y
Pool FE		Y		Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	50888	50878	50888	50878
R2	0.06	0.09	0.06	0.10

Table 3: **Determinants of LP Flows, Robustness**. The tables show various robustness checks of the regression 10.

Panel A: Volume and Volatility

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: -0.65, SD: 8.29			
Past Volume (standardized, t-1 to t)	0.60***	0.33**	1.10***	0.75***
	(3.25)	(2.57)	(5.46)	(4.96)
Past Volatility (standardized, t-1 to t)	-0.22	-0.09	-0.23*	-0.12
	(-1.56)	(-0.76)	(-1.88)	(-1.20)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Week FE		Y		Y
Pool FE			Y	Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	51024	51024	51013	51013
R2	0.03	0.06	0.06	0.10

Panel B: Daily Flows

Dependent Variable	Daily Liquidity Flow (%, t to t+1) Mean: 0.16, SD: 3.25			
Past Net Return (%, t-1 to t)	-0.00	-0.01		
	(-0.33)	(-0.50)		
Past Fee Revenue (%, t-1 to t)			0.41***	0.30***
			(5.32)	(3.74)
Past AS Loss (%, t-1 to t)			0.00	0.01
			(0.27)	(0.62)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Day FE	Y	Y	Y	Y
Pool FE		Y		Y
Standard Error	Twoway Clustered by Pool and by Day			
Observations	98672	98664	98664	98664
R2	0.05	0.07	0.03	0.07

Panel C: Monthly Flows

Dependent Variable	Monthly Liquidity Flow (%, t to t+1) Mean: 2.80, SD: 33.95			
Past Net Return (%, t-1 to t)	0.00	-0.03		
	(0.06)	(-0.42)		
Past Fee Revenue (%, t-1 to t)			0.43**	0.90*
			(2.43)	(1.87)
Past AS Loss (%, t-1 to t)			-0.01	0.01
			(-0.13)	(0.11)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Month FE	Y	Y	Y	Y
Pool FE		Y		Y
Standard Error	Twoway Clustered by Pool and by Month			
Observations	2603	2509	2603	2509
R2	0.03	0.27	0.03	0.28

Panel D: TVL > $10K

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: -0.35, SD: 9.81			
Past Net Return (%, t-1 to t)	0.03	0.03		
	(0.27)	(0.27)		
Past Fee Revenue (%, t-1 to t)			2.08***	1.73***
			(3.45)	(5.13)
Past AS Loss (%, t-1 to t)			-0.10	-0.12
			(-0.81)	(-1.22)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Week FE	Y	Y	Y	Y
Pool FE		Y		Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	56974	56966	56974	56966
R2	0.07	0.09	0.08	0.10

Panel E: TVL > $1M

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: -1.16, SD: 8.64			
Past Net Return (%, t-1 to t)	-0.10	-0.12		
	(-1.06)	(-1.06)		
Past Fee Revenue (%, t-1 to t)			1.35*	1.58***
			(1.94)	(3.22)
Past AS Loss (%, t-1 to t)			0.05	0.03
			(0.54)	(0.27)
Controls	Lagged Flow, Log TVL, Liquidity Mining Indicator			
Week FE	Y	Y	Y	Y
Pool FE		Y		Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	25160	25085	25160	25085
R2	0.07	0.14	0.07	0.14

Table 4: **LP Returns Sorted by Past Fee Revenues Only**. At the beginning of each week, I sort Uniswap V2 pools into five quintiles based on their fee revenues over the past week, form a TVL-weighted portfolio for pools in each quintile, hold the portfolios for one week, and then dissolve the portfolios. The table shows the mean (and its t-statistic in parentheses) of different metrics for each portfolio, as well as of the difference between the top quintile and the bottom quintile.

	Past Fee Revenue (%)	Fee Growth (%)	Liquidity Flow (%)	Future Fee Revenue (%)	Future AS Loss (%)	Future Net Return (%)	
						Raw	Alpha
Quintile of Past Fee Revenue (week t-1 to t):							
Low	0.11	34.06***	-4.18*	0.14	0.25	-0.11	-0.10
		(3.25)	(-1.92)			(-1.36)	(-1.39)
2	0.28	18.67*	-1.20	0.31	0.40	-0.09	-0.04
		(1.72)	(-0.39)			(-1.51)	(-0.76)
3	0.48	10.13	8.47	0.49	0.59	-0.10	-0.11
		(1.25)	(1.03)			(-0.85)	(-1.22)
4	0.83	-0.33	13.69*	0.70	0.55	0.15	0.17
		(-0.05)	(1.69)			(1.52)	(1.53)
High	2.01	-15.73***	22.68**	1.09	0.96	0.13	0.17
		(-3.54)	(2.31)			(0.92)	(0.89)
High – Low	1.90***	-49.80***	26.86***	0.95***	0.71***	0.24	0.27*
	(5.78)	(-6.91)	(3.71)	(4.26)	3.03	(1.61)	(1.75)

Table 5: **LP Returns Sorted by Both Past Fee Revenues and Past Volatilities**. At the beginning of each week, I sort Uniswap V2 pools into five quintiles based on their fee revenues over the past week, further sort each quintile into five sub-quintiles based on their price volatilities over the past week, form a TVL-weighted portfolio for pools in each sub-quintile, hold the portfolios for one week, and then dissolve the portfolios. Panel A (B) shows the raw average (three-factor alpha) – and its t-statistic in parentheses – of weekly returns of each portfolio, as well as of the difference between the sub-quintiles with the lowest volatility vs with the highest volatility.

Panel A: Raw Average

	Future Net Return, Raw Average (week t to t+1, %)				
Qunitile of Past Volatility \ Fee (week t-1 to t)	Low	2	3	4	High
Low	-0.04	-0.07*	0.11***	0.23***	0.39***
	(-1.66)	(-1.75)	(2.74)	(5.57)	(4.78)
2	-0.25***	-0.24	-0.02	-0.06	0.30***
	(-3.68)	(-1.10)	(-0.23)	(-0.19)	(2.77)
3	-0.46***	-0.19**	-0.34*	0.20*	0.50***
	(-3.14)	(-2.00)	(-1.86)	(1.74)	(4.49)
4	-0.59***	-0.62***	-0.23*	-0.11	0.31*
	(-4.82)	(-3.12)	(-1.94)	(-0.75)	(1.77)
High	-1.26***	-0.95**	-0.89***	-0.81	-0.46
	(-3.64)	(-2.64)	(-3.33)	(-1.14)	(-1.29)
Low Vol – High Vol	1.22***	0.88**	1.00***	1.04***	0.85***
	(2.92)	(2.04)	(6.03)	(6.94)	(5.63)

Panel B: Three-Factor Alpha

	Future Net Return, Three-Factor Alpha (week t to t+1, %)				
Qunitile of Past Volatility \ Fee (week t-1 to t)	Low	2	3	4	High
Low	-0.03	-0.08*	0.16***	0.26***	0.46***
	(-1.14)	(-1.85)	(3.36)	(5.76)	(4.89)
2	-0.23***	-0.23	0.06	-0.01	0.37***
	(-2.86)	(-0.87)	(0.60)	(-0.01)	(2.95)
3	-0.48***	-0.28**	-0.34**	0.31**	0.59***
	(-2.77)	(-2.64)	(-2.01)	(2.41)	(4.69)
4	-0.55***	-0.74***	-0.19	-0.23	0.31
	(-3.88)	(-3.24)	(-1.41)	(-1.38)	(1.53)
High	-1.57***	-1.11**	-0.93***	-0.81	-0.50
	(-3.64)	(-2.62)	(-2.99)	(-0.98)	(-1.21)
Low Vol – High Vol	1.54***	1.03**	1.09***	1.07***	0.96***
	(3.60)	(2.10)	(3.90)	(4.73)	(3.94)

Table 6: **LP Returns Before vs After V3 Launch**. This table repeats Table 1 for before vs after the launch of Uniswap V3. t-statistics are reported in parentheses. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

	Future Net Return (t to t+1, %)	
	Before V3 Launch	After V3 Launch
Quintile of Past Net Return (t-1 to t):		
Low	-3.17***	-2.15
	(1.39)	(1.49)
2	-0.78	-0.61
	(0.47)	(0.48)
3	-0.13	-0.13
	(0.24)	(0.12)
4	0.09	1.21
	(0.67)	(1.04)
High	1.07***	0.69
	(0.36)	(0.44)
High – Low	4.24***	2.84
	(1.45)	(1.85)

Table 7: **LP Flows Before vs After V3 Launch**. The table shows regression results of Equation 11. The sample includes weekly observations of Uniswap V2 pools with more than $100K TVL from June 2020 to March 2022. t-statistics are reported in parentheses. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Dependent Variable:	Weekly Liquidity Flow (%, t to t+1)			
Past Fee Revenue (%, t-1 to t)	5.00***	3.42***	5.47***	3.93***
	(4.75)	(4.21)	(5.43)	(4.84)
× Post V3 Launch	-4.37***	-2.97***	-3.99***	-2.45***
	(-4.00)	(-3.25)	(-3.62)	(-2.71)
Past Impermanent Loss (%, t-1 to t)	-0.85*	-0.67	-0.79*	-0.69
	(-1.90)	(-1.55)	(-1.73)	(-1.53)
× Post V3 Launch	-0.96	-0.94*	-0.90	-0.87*
	(1.32)	(1.87)	(1.25)	(1.79)
Lagged Liquidity Flow (%, t-1 to t)	0.12***	0.08***	0.10***	0.06**
	(4.63)	(2.83)	(4.05)	(2.22)
Week FE		Y		Y
Pool FE			Y	Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	7607	7607	7607	7607
R2	0.05	0.10	0.08	0.12

Dependent Variable:	Weekly Liquidity Flow (%, t to t+1)			
Past Volume (standardized, t-1 to t)	4.33***	2.83***	4.89***	3.48***
	(4.50)	(3.60)	(5.21)	(4.29)
× Post V3 Launch	-3.95***	-2.61***	-3.72***	-2.24***
	(-4.04)	(-3.11)	(-3.79)	(-2.68)
Past Volatility (standardized, t-1 to t)	-1.72	-1.36	-1.56	-1.50
	(-1.30)	(-1.61)	(-1.06)	(-1.54)
× Post V3 Launch	-1.67*	-1.42	-1.78*	-1.57*
	(-1.86)	(-1.59)	(-1.95)	(-1.74)
Lagged Liquidity Flow (%, t-1 to t)	0.13***	0.08***	0.12***	0.07***
	(5.63)	(3.32)	(5.17)	(2.81)
Week FE		Y		Y
Pool FE			Y	Y
Standard Error	Twoway Clustered by Pool and by Week			
Observations	7674	7673	7674	7673
R2	0.06	0.10	0.08	0.12