Learning Financial Innovations: Experience from Automated Market Makers*

Chuck Fang†

December 2024

Abstract

I show the evolution of capital allocation efficiency on blockchain-based automated market makers. There was initially large misallocation of capital in the form of persistent abnormal returns, which arose because capital flows were asymmetrically sensitive to returns that are salient versus shrouded. Capital allocation improves over time, and this improvement is concentrated in liquidity providers that experience actual realization of large losses. The findings demonstrate how behavioral biases can impede the adoption of financial innovations and how agents overcome these biases by learning from personal experiences.

Keywords: financial innovation, automated market maker, liquidity provision, behavioral bias, personal experience

*This paper was previously circulated as "Liquidity Misallocation on Decentralized Exchanges".

†Drexel University. Email: chuck.fang@drexel.edu. I thank comments from Agostino Capponi, Itamar Drechsler, Itay Goldstein, Urban Jermann, Weitong Han, Ruizhe Jia, Alfred Lehar, Xinwei Li, Daniel Nathan, Thomas Rivera, Samuel Rosen, Wanyi Wang, Yan Xiong, Yang You, Anthony Lee Zhang and participants at Waseda Workshop, Columbia TLDR, EasternFA, MFA, Tsinghua PBCSF, FMA, JFDS Conference, UCSB-ECON DeFi Seminar, Wharton-INSEAD PhD Consortium, Cardiff Fintech Conference and Wharton PhD Seminar.

1 Introduction

New financial innovations can be initially difficult to understand.[1] Mistakes are often made, especially when involving unsophisticated investors that are prone to behavioral biases. At the same time, investors should be able to learn from their mistakes, correct their biases and behave more optimally over time. On the macro level, we should observe a dynamic evolution of market equilibrium, where efficiency might be low initially but should improve over time.

In this paper, I study the adoption of automated market makers (AMMs), a new design of liquidity market for the trading of cryptocurrencies directly on blockchain. Launched in the summer of 2020, AMMs have intermediated a total trading volume of $3.2 trillion by the end of 2023, with many calling it the future of finance (Harvey et al., 2021). As the name suggests, AMMs automate liquidity provision, which facilitates the entry of unsophisticated retail investors that would have been excluded under traditional market structures such as the limit order book. Large retail participation[2], together with the availability of granular transaction-level data on blockchain, makes AMMs an ideal laboratory to shed light on the adoption process of financial innovations.

I have three findings. First, there was initially significant misallocation of liquidity in the form of large predictable abnormal return differences across AMM liquidity pools, which averaged at 33% per annum. Second, this return predictability can be traced to suboptimal liquidity flows, which were asymmetrically sensitive to the part of liquidity return that is salient vs the part that is shrouded. Lastly, over time, there is substantial improvement in the optimality of liquidity flows and the overall market efficiency, and this improvement is concentrated in liquidity providers (LPs) who experience actual realization of large losses.

[1] For example, we are still learning today about the full implications of exchange-traded funds (e.g. see survey by Ben-David et al., 2016), the first of which (SPY) was developed decades ago in 1993.

[2] Liao and Robinson (2022) shows that the participation of passive retail investors are essential for the competitiveness of AMM liquidity.

I begin by describing in detail how liquidity provision works in a constant-product automated market maker (CPAMM), the most common design of blockchain-native exchanges. Liquidity provision on CPAMM is subject to the usual trade-off between gains from noise trades (in the form of fee revenue) and losses from informed trades that take advantage of stale prices (a.k.a. "impermanent loss"). In contrast to traditional market structures such as the limit order book, CPAMM automates the adjustment of prices as simple deterministic functions of order flows. On one hand, it makes liquidity provision easy, which facilitates the entry of unsophisticated investors. On the other hand, this automation necessitates losses due to stale prices during the price adjustment process.

In an efficient market, liquidity provision is expected to earn zero risk-adjusted returns. Instead, I find significant abnormal returns to liquidity provision on Uniswap, the largest CPAMM. Using a sorted portfolio approach, I show that pools in the top quintile of *past* returns continue to have significantly higher *future* returns than the bottom quintile pools, at 0.64% per week, or 33% annualized. These return differences cannot be explained by common risk factors, idiosyncratic risks or transaction costs. More importantly, these abnormal liquidity returns shrink substantially over time, reflecting the improvement in the efficiency of this market. Exploiting the fact that impermanent losses on AMM are price-contingent, I compute options-implied liquidity compensation and find similar evolution of liquidity return anomaly.

Why were there large predictable return differences across pools? Why didn't liquidity adjust to equalize pool returns in equilibrium? I show that liquidity flows were insensitive to net returns. Instead, liquidity flows reacted strongly to fee revenues – the part of return that is saliently marketed as APY – but ignored impermanent losses – the part of return that is often shrouded. As a result, pools with higher fee revenues attracted more liquidity, even though they tend to have higher impermanent losses and insignificant returns on net. I verify that focusing on fee revenue alone is suboptimal, and taking impermanent loss risk into account

significantly boosts the payoffs to LPs. In a simple counterfactual exercise, I show that, had liquidity flows responded symetrically to impermanent losses, return anomalies would have shrunk by half.

Similar to the evolution of liquidity return predictability, I observe substantial improvement in liquidity flow optimality over time – that is, liquidity flows become more sensitive to impermanent losses. Using data at the individual LP address level, I show that the improvement in liquidity allocation primarily comes from those that have experienced actual realization of large impermanent losses. Using LP fixed effects, I show that for the same LP, realization of large losses (worse than -10%) upon liquidity withdrawal are associated with significantly higher sensitivity to impermanent losses in *future* liquidity provision.

The rest of this paper is organized as follows. Section 1.1 summarizes this paper's contributions to the literature. Section 2 provides detailed descriptions of automated market makers, in particular its liquidity provision. 3 describes the data. Section 4 shows liquidity return anomaly. Section 5 shows suboptimal liquidity flows that reconcile with liquidity return anomaly. Section 6 concludes.

1.1 Literature

This paper contributes to the understanding of financial innovations, more specifically their dark side due to the presence of retail behavioral biases (Henderson and Pearson, 2011; Gennaioli et al., 2012; Célérier and Vallée, 2017; Vokata, 2021; BARBER et al., 2022; Ben-David et al., 2022). My paper adds to this literature with detailed analyses of a new financial innovation with trillions of dollars of economic activities and granular transaction-level data. On one hand, AMMs democratize market making, and I show simple strategies to identify profitable investment opportunities. On the other hand, many AMM liquidity providers appear to be unsophisticated investors that make suboptimal decisions due to the salience effect (Au-

gustin et al., 2022), resulting in market inefficiencies on the largest AMM (Uniswap). More importantly, my paper emphasizes the role of personal experiences in mitigating behavioral biases (Malmendier and Nagel, 2011, 2015). From a policy standpoint, my results suggest more regulation – either by the government or by industry participants themselves – be put in place to curb the shrouding of impermanent losses (Gabaix and Laibson, 2006). More specifically, there should be stricter guidance on the use of "APY" in marketing materials, especially for investment products with volatile returns such as liquidity pools.

This paper adds to the growing literature on automated market makers, 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); Milionis et al. (2022); Augustin et al. (2022). Most of the papers are theoretical and assuming perfectly optimizing agents. My contributions are twofold. On the liquidity return side, I show that liquidity returns exhibit anomaly, suggesting that misallocation is an important driver of liquidity returns on AMMs. I also emphasize that the right benchmark for liquidity return is not zero, but rather the premium for bearing risks associated with impermanent losses, which I quantify using factor models as well as an option pricing framework. On the liquidity flow side, I find that LPs make suboptimal decisions – they are highly sensitive to fee revenues, the part of returns saliently marketed as APY, but not to impermanent losses, the part of returns that is often shrouded.

In analyzing liquidity flows, My methodologies borrow heavily from the literature on mutual funds (Berk and van Binsbergen, 2016; Barber et al., 2016; Song, 2020; Ben-David et al., 2021), which bear significant resemblance to AMM liquidity pools. Similar to the mutual fund market, AMM investors make suboptimal decisions in capital allocation, leading to abnormal returns. The most related paper is Song (2020), who shows that mutual funds with high factor-related returns accumulate too much assets that their future returns significantly under-perform. While mutual fund analyses are typically limited to the fund level, the

availability of individual address-level data on blockchain allows me to conduct more granular analyses and establish the role of learning in capital allocation decisions.

2 Constant-Product Automated Market Maker

In traditional limit order book exchanges, liquidity is usually supplied by high-frequency traders that post and update limit orders at millisecond intervals. Such high-frequency operations have been infeasible on blockchains such as the Ethereum mainnet due to technical limits on block size and block speed. As a result, the adoption of limit order book on blockchain (e.g. AirSwap) has been limited.[3] Instead, automated market makers (AMMs) have proven to be a successful design for decentralized exchanges (DEXs), i.e. exchanges that facilitate the trading of cryptocurrencies directly on blockchain.

I focus on one particular class of AMMs: the constant-product automated market maker (CPAMM). As will be shown in Section 3, CPAMM accounts for most of the trading volume and liquidity among DEXs during my sample period. On CPAMM, each pair of assets forms a pool – e.g. the first asset is BTC and the second asset is USD. Denote their amounts in the pool (i.e. reserves) by X and Y, respectively. The 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 measures the pool's liquidity and stays invariant to all swaps. Panel A of Figure 2 visualizes this rule. Suppose there are initially 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 contract calculates what would be the new amount of USD Y_1 so that the pool stays

[3] Layer-2 developments such as Polygon has overcome these technical limits and led to revival of limit order book on blockchains.

on the the curve. The difference ΔY is what the trader swaps out in return.

This simple design achieves a number of properties that are essential for market making. First, the more that one trades, the worse price she gets. As illustrated in Panel B of Figure 2, if a trader doubles or triples the amount of BTC swapped in ($\Delta X' = 2\Delta X$ or $\Delta X'' = 3\Delta X$), the amount of USD swapped out ($\Delta Y'$ or $\Delta Y''$) is much less than doubled or tripled. This is because the curve is convex. If convexity is zero, i.e. the curve is a straight line, one additional unit of ΔX always gets the same change in ΔY.[4]

Secondly, price reflects the relative scarcity of the two assets. The price of BTC in terms of USD can defined as the amount of USD swapped out (swapped in) per unit of BTC swapped in (swapped out), for an infinitesimal amount:

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

where the second equality follows easily from Equation 1. Therefore, if the reserve of BTC increases and/or the reserve of USD decreases, the price of BTC relative to USD would increase. This can also be seen from Panel B of Figure 2, where $P \equiv -\frac{\partial Y}{\partial X}$ is the slope of the tangent line to the constant-product curve. More BTC (less USD) leads to a flatter tangent line.

Lastly and most importantly, pool price automatically adjusts to market price through the force of arbitrage.[5] Suppose that the market price of BTC goes up – for example, you can buy or sell Bitcoin at a higher price on other exchanges (e.g. Binance). Since pool price is still low, there is an arbitrage opportunity to buy BTC from the pool and sell it on other exchanges. However, this arbitrage will decrease the reserve of BTC and increase the reserve of USD, which would lead to higher pool price according to Equation 2. Arbitrage

[4] Angeris et al. (2020) discuss more on the properties implied by the convexity of the curve.

[5] The language here assumes that price discovery occurs on other exchanges, which is likely given that AMMs only account for 1% of the total crypto trading volume. Alternatively, price discovery could happen directly on the AMM through informed trades.

opportunity exists until pool price equals to market price. Angeris and Chitra (2020) and Lehar and Parlour (2021) show that prices on CPAMMs are mostly aligned with prices on major exchanges – i.e. no arbitrage holds.

2.1 Liquidity provision on CPAMM

In order for traders 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 providers deposit or withdraw equal value of both assets, so that pool price remains unchanged.[6]

When an LP deposits liquidity, she gets LP tokens that represent her share of the pool. The number of LP tokens outstanding equals to $\sqrt{K}$. Another useful measure of liquidity is total value locked (TVL). Using the second token as the numeraire, TVL is defined as:

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

P denotes price in units of the second token and the second and third equality follows from Equation 2. When I need to compare liquidity across pools, I express TVL in terms of U.S. dollar:

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

where $P_X^{\$}$ and $P_Y^{\$}$ are prices in units of U.S. dollar.

[6]To see why, first observe that the two assets always have equal value: $PX = Y$, which follows directly from Equation 2. Therefore, in order for P to remain constant, any addition or subtraction of reserves need to happen to both assets in equal value.

2.2 Returns on liquidity provision (LP returns)

Why would anyone deposit their assets into a liquidity pool? The answer is that there can be large returns to liquidity provision (LP returns). I define LP returns as relative to holding returns. For example, suppose that an LP deposits some BTC and some USD into a liquidity pool, I focus on her return relative to holding on to the assets.[7] LP returns consist of two components: fee revenues and impermanent losses.

Fee revenue. As the price of consuming liquidity, each trade pays a fee equal to a constant fraction of the amount of asset swapped in, e.g. 0.3% for all Uniswap V2 pools. This fee is automatically charged and deposited into the pool along with the asset swapped in. For example, to sell 1 BTC, the trader actually needs to deposit 1.003 BTC, where the 0.003 BTC represents the fee. Fee revenue is shared equally by all liquidity providers. Therefore, there are natural decreasing returns to scale. If more liquidity providers are attracted to a pool, liquidity increases, and fee revenue per liquidity provider decreases.

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 and fees are paid. Now, a large amount of new liquidity is deposited to the pool, and then more trades occur and more fees are paid. 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 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,

[7] LPs can neutralize their exposure to holding returns through borrowing. For example, one can borrow BTC through a decentralized lending platform such as Aave and then provide liquidity to a BTC-USD pool.

the marginal LP's share of the pool shrinks, and her share of subsequent fee revenues is proportionally reduced. Conversely, if there are liquidity outflows, the marginal LP's share of the pool expands, and her share of subsequent fee revenues is proportionally multiplied. Formally, given fee rate f, total fee revenues for a token during block t can be easily calculated from amount of that token swapped in:

$$fSwapIn_t$$

For a marginal LP that entered the pool during block t, the amount of fee revenues accrued to her during $t+s$ is given by:

$$AccruedFee_{t,t+s} = fSwapIn_{t+s}\frac{\sqrt{K_{t+s-1}}}{\sqrt{K_t}}$$

where $\sqrt{K_t}$ denotes number of LP tokens outstanding at the end of block t. Finally, this marginal LP's fee revenue from t to $t+h$ is:

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

where the scaling by TVL expresses fee revenue as a percentage of liquidity. Note that fee revenues are never negative – the worst that can happen is that there is no trade and hence zero fee revenue.

Impermanent loss. Recall that CPAMM is able to achieve automatic price adjustment through the help of arbitrageurs. If market price of BTC falls, arbitrageurs will swap in BTC (swap out USD) until pool price equals market price, and this way subsequent trades occur at the "right" price. This automatic price adjustment is not cost free, and LPs incur losses to arbitrageurs whenever price changes. To see why, suppose that pool price equates market price initially, and then market price increases. Pool price does not adjust automatically – arbitrageurs need to 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 low pool price instead of the higher market price. Figure A1 give simple numerical examples. This loss is commonly referred to as "impermanent loss", where "impermanent" refers to the fact that losses go back to zero if price completely reverts, which may or may not occur.[8]

Assuming that no arbitrage condition holds – that is, pool price equals to market price – impermanent loss is a deterministic function of price change. To see this, first note that given pool price P, pool reserves X and Y and TVL are pinned down, absent any deposit or withdrawal of liquidity (i.e., K constant):

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

Suppose that price changes from P to P'. Assuming that there is no fee revenue, impermanent loss is defined as:

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

where the first term in the numerator is the holding value (i.e. the opportunity cost of providing liquidity) and the second term the new TVL. It is easy to show that impermanent loss is never negative. Intuitively, whenever price changes, LPs trade with arbitrageurs at the stale bad price and lose.

Figure 3 visualize the deterministic relationship between impermanent losses and prices changes. Impermanent losses are bigger for bigger price changes. Therefore, pools where token prices are more volatile incur higher risk of impermanent losses to the liquidity provider.

Net return. From a LP's perspective, what matters is the net effect of fee revenues and

[8]Instead of loss-vs-holding, Milionis et al. (2022) show that LP returns can be characterized by loss-vs-rebalancing (LVR). My main results are qualitatively the same when using LVR.

impermanent losses:

$$NetReturn = FeeRevenue - ImpermanentLoss \tag{7}$$

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

2.3 Comparison with limit order book (LOB)

In a limit order book (LOB), LPs earn bid-ask spreads. A market buy order and a market sell order that occur at the same time would pay different prices, and LPs earn the difference. There is only one price on AMM, so there is no bid-ask spread. When there are one buy and one sell at the same time, they trade at effectively the same price. Instead, CPAMM LPs earn fee revenue, which are more like the exchange operators themselves, such as Binance, which charge a constant fee for using their platforms.

LOB LPs need to specify both price and quantity. For example, they can post a bid of 1 BTC for 40,000 USD, and they need to actively monitor the price so that they do not overpay (e.g. when BTC price declines due to low sentiment). In contrast, CPAMM LPs only specify quantities and are designed to be passive. Prices are automatically adjusted by arbitrageurs. Due to this simplicity, CPAMM facilitates the entry of unsophisticated investors, as I will discuss in Section 5.

LPs can be picked off by arbitrageurs if price changes, and this is true on both LOB and CPAMM. On LOB, LPs need to actively monitor and adjust prices, and they can control how they want to change prices in respond to order flows. In contrast, CPAMM automates the response of price to order flows. There is one less degree of freedom that LPs can control. As a result, this stale price loss is a deterministic function of price change.

2.4 Characteristics of liquidity providers

Who provides liquidity on AMMs? While blockchain is by definition anonymous, we can obtain characteristics of the wallets that provide liquidity on AMMs. This is shown in Figure A2. Panel A plots distribution of size of liquidity positions on Uniswap V2, where each liquidity position refers to a wallet with positive ownership of LP tokens to any Uniswap V2 pools. It shows that most of the liquidity positions are small, at less than $20K.

Panel B plots distribution of duration of liquidity positions on Uniswap V2, defined as the number of days for an LP (i.e. one with non-zero LP token) to make any adjustment to its LP token holdings. It shows the average LP is quite passive, since less than half of the LPs update their positions on a daily frequency.

3 Data

I study Uniswap V2, arguably the first successful application of CPAMM and still one of the largest blockchain-native exchanges to date. As shown in Figure 1, Uniswap V2 had $1.1 billion ($5.0 billion) of total value locked in December 2022 (December 2021) and supported $114 billion ($261 billion) of trading in 2021 (2022). I focus on V2 of Uniswap because subsequent versions (e.g. V3) make the liquidity provision mechanism significantly less automated and therefore harder to access for retail investors.[9]

I obtain transaction-level data on Uniswap V2 from Dune Analytics. For the pool-level analyses, I aggregate transactions for each pool at daily, weekly, or monthly frequency. I focus on the weekly frequency because this appears to be the modal horizon for liquidity adjustment. According to Table A1, the standard deviation of weekly liquidity flows (i.e.

[9]Uniswap V3 introduces concentrated liquidity, which essentially allows liquidity providers to specify the price with which to provide liquidity. This makes liquidity provision much more similar to traditional limit order book models.

how much does liquidity change on a weekly basis) is 15%. According to Figure A2, over 50% of the LPs take between 3 to 30 days to make adjustments to their liquidity positions.

I hand collect airdrop programs for all Uniswap V2 pools. For each token in the pools, I go through the blog posts on the protocol's official website and twitter account, and record the dates that airdrops are given.

Daily prices on cryptocurrencies are from CoinMarketCap and obtained through Yahoo Finance API. Risk-free rates are estimated using yields on U.S. Treasury bills, sourced from the CRSP U.S. Treasury Database. Data on BTC and ETH options are from T3 Index.

4 Liquidity Return

I take an asset pricing approach to study AMM liquidity returns. In an efficient market, risk-adjusted returns should be zero in expectation (Fama, 1970). Suppose that some liquidity pools yield abnormally high returns, they should attract more liquidity providers, this would dilute fee revenue for each liquidity provider, and liquidity returns would go down. A similar argument can be made for pools yielding abnormally low returns. Therefore, persistent abnormal returns are signs of misallocation – pools with large positive (negative) abnormal returns do not get sufficient inflows (outflows) of liquidity.

I show that returns on Uniswap liquidity pools initially exhibit return anomalies. I use two complementary approaches. First, I show that liquidity pools with higher (lower) past returns continue to have significantly higher (lower) future returns. Then, I calculate risk-implied liquidity compensation on a small subset of liquidity pools with options data and find persistent deviation from what is actually realized. I show that these abnormal returns diminish significantly over time.

4.1 Sorted portfolio approach

I start by showing large differences in returns across liquidity pools. As shown in Figure A3, in a typical week, some pools can have large positive returns of over 1% (67% annualized) while some pools can have large negative returns lower than -1%. I will show that these large return differences are persistent and predictable.

I use a sorted portfolio approach. At the beginning of each week, liquidity pools are sorted into five quintiles based on their realized returns over the past week. A value-weighted portfolio is formed for pools in each quintile. The portfolios are held for one week and then dissolved. Small pools can occasionally have very large trades and hence fee revenues relative to their sizes, so I restrict to pool-week observations with at least $100K total value locked. I exclude pools with airdrops, because reward tokens are usually newly issued and do not have price information. I focus on the initial adoption phase, from the 18th week of 2020 (when V2 was launched) to the last week of 2021. I require at least 50 pools in the cross section, which drops the first three weeks since V2 launch.

Panel A of Table 1 shows the main results. In the first year of Uniswap V2 launch, pools with high (low) past returns continue to earn high (low) returns the week after. According to Column 3, pools in the highest (lowest) quintile of past returns earn on average 0.29% (-0.35%) in excess of risk-free rate in the week that follows, or 16% (-17%) annualized. The return difference between the top quintile and the bottom quintile of past return is 0.64% per week, or 33% annualized.

Column 1 and 2 break down liquidity return into its two components: fee revenue and impermanent loss. The low-return pools actually have high fee revenues. However, they have even higher impermanent losses, resulting in low returns on net. As I will show later, these low-net-return pools do not experience adequate outflows precisely because the average liquidity provider pays (excessive) attention to fee revenue but inadequate attention to impermanent

loss.

Do these return differences simply reflect different levels of risk? In theory, diversified investors should only care about common risk factors. To account for common risk factors, I use the three-factor model from Liu et al. (2022), who show that market, size and momentum capture a large portion of the cross-sectional variation 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 sorted portfolio strategy, $R_{f,t}$ risk-free rate, α_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 cryptocurrencies that are past winners vs losers, and $\{\beta_i\}$ exposure to the risk factors. Column 4 shows that the return differences between high-past-return pools vs low-past-return pools are not explained by the common crypto risk factors and remain largely the same. Column 5 shows that the results remain the same if I additionally control for Fama-French three factors of equity market (Fama and French, 1993).

Panel B of Table 1 explores whether the return patterns can be explained by other rational reasons, such as idiosyncratic risks, underlying returns or transaction costs. AMM LPs may not be well-diversified and are therefore sensitive to idiosyncratic risks. Column 1 examines whether the return differences can be explained by idiosyncratic volatility – namely, whether higher (lower) returns are associated with higher (lower) volatility. The results show that the relationship is not monotonic: both high-return pools and low-return pools have higher volatility.

Liquidity providers are also likely exposed to idiosyncratic risks of the underlying assets.[10] Can liquidity returns be explained by returns on the underlying assets – for example, some pools require higher liquidity returns because their underlying assets have lower returns? Column 2 shows that returns of the underlying assets are similar across the five quintiles.

Column 3 shows that the underlying assets of the low-liquidity-return pools have higher return volatilities, which is consistent with the fact that these pools have higher impermanent losses (see Equation 6). The patterns are inconsistent with a risk-based explanation: higher volatilities should provide higher not lower returns.

How would transaction costs affect the results? In order to provide liquidity on AMMs, LPs need to pay gas fees and costs of acquiring (either buying or borrowing) the underlying assets. Note that these costs are size contingent – the bigger amount of liquidity provided, the less the gas fees in relative terms (since gas fees are independent to transaction amount), and the harder it is to acquire the underlying assets (e.g. more slippage). I estimate transaction costs for an LP that deposits $10K of new liquidity, using the prevailing gas price and assuming that the LP buys the underlying assets using USDC through Uniswap V2 pools. Column 4 shows that transaction costs are large, at around 30 bps on average. However, transaction costs are not monotonic in liquidity returns. Conditional on deciding to provide liquidity, an LP would still earn substantially more returns from the high-past-return pools than from the low-past-return pools after accounting for transaction costs.

4.1.1 Liquidity return over time

Panel A of Figure 5 plots the 10-week rolling average of the LP net return difference between the top vs bottom quintile of pools sorted by past net returns. It shows that abnormal return differences across liquidity pools have declined significantly over time. Perhaps surprisingly,

[10] In theory, liquidity providers can borrow the underlying assets (e.g. from a decentralized lending platform such as Aave) and be insulated from holding returns.

there is still predictable abnormal return at year-end 2022 that is statistically significant, at around 0.20% per week, or 11% annualized, which is on the same magnitude as equity premium.

4.2 Option pricing approach

The abnormal liquidity returns I found can reflect compensation for risk that is not adequately captured by the factor models. In this section, I use an option pricing approach to further corroborate the anomaly of liquidity returns. As shown in Section 1, impermanent losses are a deterministic function of the relative price of the two tokens, assuming that no arbitrage holds. This means that I can derive the present value of impermanent losses in a Black-Scholes option-pricing framework. Indeed, as shown in Figure 3, impermanent losses resemble the payoff of a straddle, i.e. the combination of a call and a put. Evans (2020) and Milionis et al. (2022) provide in-depth derivations of this option-pricing approach.

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

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

Terminal price can be then 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 impermanent loss is:

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

where $ImpermanentLoss$ is given in Equation 6.

One key input to this framework is volatility.[11] 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 risk premium baked in. These option-implied volatilities allow me to answer the question: what would be the present value of options that replicate impermanent losses?

I 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 $ImpermanentLoss_T$. I take the average to obtain $E^Q[ImpermanentLoss_T]$ and discount it by risk-free rate to get V^{AS}.

Figure 6 shows the results. The graphs plot realized and options-implied LP compensation for the BTC-USD pools and the ETH-USD pools, averaged across the three stablecoins: USDC, USDT and DAI.[12] One can easily observe the persistent difference between realized fee revenues and options-implied compensation for impermanent losses. For example, during 2020-2021, 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 to impermanent losses. Because there is a positive correlation between fee revenue

[11] 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.

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

and impermanent loss, fee revenues are a natural hedge against impermanent losses and should therefore require a lower premium, which makes the positive gap between actual vs options-implied LP compensation even more puzzling. The results here indicate misallocation, since liquidity provision would have earned very different returns compared to the risk-based compensation implied by the options market.

4.3 Interpreting abnormal returns

Why were there large abnormal return differences across liquidity pools during the initial adoption phase? Theoretically, pools with large positive abnormal returns should see inflows of liquidity that would dilute fee revenue and hence lower net return for each LP. Pools with large negative abnormal returns should see the opposite.

One natural explanation is transaction costs. The fact that, in Column 4 of Panel B of Table 1, pools with bigger abnormal returns (both positive and negative) have higher transaction costs confirms that transaction costs play an important role in obstructing the free flow of capital. However, if transaction costs were the only explanation, we should observe a lack of liquidity flows, but in reality there were frequent liquidity adjustments at weekly frequency, as I have shown in Section 2.4. I will show in the next section that LPs did take actions frequently despite high transaction costs, but their actions were suboptimal, which perpetuate the misallocation of liquidity.

What is the role of arbitrageurs in this context? For the cross-sectional return difference anomaly, an arbitrageur could form a zero-position portfolio that longs the pools with high past returns and short the pools with low past returns. However, it is difficult to short a liquidity pool, since there is no active market for shorting LP tokens. The return anomaly here is akin to the mutual fund manager return anomaly (e.g., Song, 2020), where it is difficult to short a mutual fund share in practice. The barrier for the options-implied arbitrage is

arguably lower. To exploit the gap between options-implied liquidity compensation and realized fee revenue, arbitrageurs can provide liquidity to the BTC-USD pool and then short a basket of Bitcoin options. In other words, the arbitrageur can delta hedge its impermanent loss (Cong et al., 2022). One conceivable constraint here is that the arbitrageur might be averse to returns on the underlying, and it can be costly to borrow the underlying asset.

5 Liquidity Flow

The previous section shows that, during the initial adoption phase, there are persistent net return differences across liquidity pools. Why doesn't liquidity adjust? If some pools are earning abnormally high net returns, more liquidity should flows in, diluting the fee revenue and lowering the net return. In this section, I directly examine LPs' actions to deposit or withdraw liquidity. I find that LPs are initially insensitive to net returns. Instead, LPs are highly sensitive to fee revenues but not to impermanent losses. As a result, pools with higher fee revenues attracted more liquidity, even though they tended to have higher impermanent losses and insignificant returns on net.

Figure 4 shows what regular LPs would see on the websites of Uniswap and SushiSwap, two of the most popular DEXs that both work as CPAMM. Fee revenues are saliently displayed as APY (annual percentage yield). The rationale seems to be that fee revenues are never negative and therefore resemble deposit-like investment products with guaranteed minimum returns (0%). However, fee revenues are highly volatile (with annualized standard deviation of 6.71% according to Table A1) and moderately persistent (with a weekly autocorrelation of 0.4). More importantly, negative returns due to impermanent losses are not in the calculation of APY. Unsophisticated LPs therefore can be influenced by the display and place inadequate emphasis on impermanent loss relative to fee revenue.

5.1 Pool-level analysis

I first study liquidity flows at the pool level. I run the following regression using a panel of liquidity pools from the 18th week of 2020 (when V2 was launched) to the last week of 2021:

$$Flow_{p,t+1} = \beta_1 FeeRevenue_{p,t} + \beta_2 ImpermanentLoss_{p,t} + \gamma Controls_{p,t} + FE + \epsilon_{p,t+1} \quad (9)$$

where $Flow_{p,t+1}$ is defined as the net change in pool p's liquidity (i.e. number of LP tokens) from week-end t to week-end $t+1$:

$$Flow_{p,t+1} = \frac{\sqrt{K_{t+1}}}{\sqrt{K_t}} - 1 \quad (10)$$

$FeeRevenue$ and $ImpermanentLoss$ are fee revenue and impermanent loss (defined in 5 and 6) from week-end $t-1$ to week-end t, controls include past liquidity flows, log TVL and indicator of airdrop, and FE includes time fixed effects and pool fixed effects. As small pools are likely to have large flows (in percentage points) that may bias the results, I restrict to pool-week observations where TVL is greater than \$100K. I also winsorize all variables at the 5th percentile in each cross section to mitigate the influence of outliers.

In addition to the pooled regression above, I also run the Fama-MacBeth regression (Fama and MacBeth, 1973). Specifically, I run the following regression for each cross section of pools in each week:

$$Flow_{p,t+1} = \beta_{1,t} FeeRevenue_{p,t} + \beta_{2,t} ImpermanentLoss_{p,t} + \gamma_t Controls_{p,t} + \epsilon_{p,t+1} \quad (11)$$

and then conduct inference on the time series of regression coefficients $\{\beta_{1,t}, \beta_{2,t}\}$. As discussed in Ben-David et al. (2021), Fama-MacBeth is the preferred regression specification when different flow determinants are compared to one another and their influences on flows are expected to vary over time.

Table 2 shows results on liquidity flow determinants. Column 1 and 4 examine the relationship between future liquidity flow and past net return. The coefficient on past net return is not significantly different from zero, suggesting that future liquidity flow does not respond to past net return. This resonates with my previous results on persistent return differences: pools with high (low) past net returns do not experience enough inflows (outflows) of liquidity and therefore continue to provide high (low) net returns going forward.

One explanation is that, since liquidity can be passive, LPs simply do not take actions and liquidity flows are simply zero. I show that LPs *actively* make mistakes. Column 2 and 5 breaks down net returns into its two components: fee revenues and impermanent losses. The coefficient on past fee revenues is positive and significant. According to column 5, 1 p.p. increase in fee revenues over the past week is associated with 3.81 p.p. higher flow of liquidity into the pool over the subsequent week. This is economically significant, since the mean (standard deviation) of weekly liquidity flows is -0.05 p.p. (17.00 p.p.) In contrast, the coefficient on past impermanent losses is small and not significantly different from zero. Taking together, the results suggest that LPs are highly sensitive to past fee revenue but not to past impermanent losses.

Instead of LP returns, I can instead focus on the primitives that determine liquidity returns (Capponi and Jia, 2021). Specifically, I substitute fee revenues with trading volumes:

$$Volume_{p,t} = (\frac{SwapIn^A_{p,t}}{Reserve^A_{p,t-1}} \times \frac{SwapIn^B_{p,t}}{Reserve^B_{p,t-1}})^{1/2}$$

and impermanent losses with price volatility:

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

Column 3 and 6 show the results and the main message remain unchanged.

Among the control variables, the coefficient on past flows is highly significant, consistent with the large literature on mutual fund flows. The coefficient on airdrop indicator is also highly significant, consistent with the anecdotal evidence that airdrops are an effective tool to boost liquidity in the short run. These results affirm the importance of controlling for these variables.

I conduct a battery of robustness checks in Table A2. Panel A and B perform the same analysis but restricting to bigger pools where the underlying assets are more liquid. My main results hold. Panel C and D conduct the same analysis but at daily or monthly frequency. Net return and impermanent loss (or volatility) seem to have some explanatory power at the monthly frequency, suggesting that liquidity misallocation is corrected (partially) in the long run.

My finding that LPs are insensitive to impermanent losses may seem inconsistent with Table 2 of Lehar and Parlour (2021), who show that pools with higher past volatility have lower dollar TVL. I am able to replicate their results using my data in Table A3. In contrast to their focus on pool size, I focus on flows, which exclude asset appreciation or depreciation as a potential confounder.[13] More importantly, I will show in Figure 5 that there is an evolution of flow sensitivity to volatility, which converges to the theoretical prediction of Lehar and Parlour (2021) over time. My result that LPs chase fee revenue and ignore impermanent loss is similar to the findings in Augustin et al. (2022), who study a sample of PancakeSwap pools.

5.1.1 Evidence of suboptimality

How can we be sure that the flow patterns are indeed suboptimal? LPs might chase past fee revenues more than past impermanent losses, because the former is more persistent and

[13] In the mutual fund literature, fund size is rarely the left-hand side variable because of obvious endogeneity. Similarly, it is possible that smaller pool size (e.g. meme coins) reverse causes higher volatility.

hence a more reliable signal than the latter. Indeed, the auto-correlations of fee revenues and impermanent losses are 0.5 and 0.3 at weekly frequency, 0.4 and 0.2 at monthly frequency. I show below that LPs can derive significant gains by expanding their attention to impermanent losses instead of chasing past fee revenues alone.

First, I simulate the payoff of naively chasing past fee revenues without considering impermanent losses, using the same sorted portfolio approach from Section 4.[14] Table 3 shows the results. Column 1 shows future fee revenue, which goes up monotonically, confirming the persistence of fee revenue. Column 2 shows that pools with higher fee revenue also tend to have higher impermanent losses. How does this happen? High fee revenue means that the assets of the pools are traded a lot. Trading can lead to negligible or temporary price impacts if it is driven by liquidity reasons. However, trading can lead to long-lasting price impacts and hence impermanent losses if it reflects new information or sentiment. Column 3, 4 and 5 show that, after risk adjustment, 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, sorted by past fee revenue.

It is easy to take impermanent loss into account and significantly increase the payoff of liquidity provision. 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 4 shows the importance of controlling for impermanent loss risk. There is significant return difference between high-past-volatility pools and low-past-volatility pools, conditional on having the same past fee revenues. Pools with the highest fee revenue can have the worst net returns because of impermanent loss, given by those in the bottom right corner of the table. Pools at the top right corner, i.e. those with the highest fee and the lowest volatility – presumably those with the highest amount of liquidity trades

[14] Specifically, at the beginning of each week, I sort all Uniswap V2 pools into five quintiles based on their past fee revenues only, form a size-weighted portfolio for pools in each quintile, hold the portfolios for one week and dissolve them.

relative to the amount of informed or sentiment-driven trades – earn a risk-adjust return of 0.51% per week, or 30.28% annualized.

To further link the flow patterns above to the abnormal returns in Section 4, I conduct the following counterfactual exercise. Specifically, I simulate what would liquidity allocation look like if LPs were equally sensitive to impermanent losses as they were to fee revenues. Since impermanent losses are always positive, the counterfactual liquidity would have been lower. Lower liquidity would in turn lead to lower trading volume and lower fee revenue, which would further depresses liquidity provision, but I provide a naive derivation of counterfactual liquidity without this multiplier effect. For each pool p in each week t, there would have been additional an outflow equal to:

$$\tilde{Outflow}_t = (|\beta_1| - |\beta_2|)ImpermanentLoss_{p,t-1} \tag{12}$$

where β_1 and β_2 are, respectively, flow sensitivity to fee revenue and impermanent loss (which are 3.81 and -0.35, respectively, according to Column 5 of Table 2), and $ImpermanentLoss_{p,t-1}$ is impermanent loss realized in the preceding week. The counterfactual reserve for each token at week-end t is calculated recursively as:

$$\tilde{Reserve}_t = \tilde{Reserve}_{t-1}(1 - \tilde{Outflow}_t) + \Delta Reserve_t \tag{13}$$

where the last term represents actual net change in reserve.

Figure A4 shows the aggregate amount of counterfactual liquidity. The blue line shows the actual amount of liquidity aggregated across all Uniswap V2 pools, which peaked at $8.9 billion in May 2021 before the launch of Uniswap V3. The red line shows the counterfactual amount of aggregate liquidity if LPs were as sensitive to impermanent losses as to fee revenues, and the number turns out to be $5.7 billion as of May 2021, or 64% of actual liquidity.

Table 5 repeats the sorted portfolio exercise in Section 4.1 but with counterfactual liquidity. The results suggest that persistent abnormal return differences would would have declined by half to about 30 bps per week, from about 60 bps in Table 1. The effect is driven by the low-return pools: these pools have large impermanent losses, and IL-sensitive LPs would have pulled liquidity from these pools, increasing their fee revenues per unit of liquidity and overall net returns.

5.1.2 Liquidity flow over time

Panel B of 5 plots the 10-week rolling average of estimated liquidity flow sensitivity to past fee revenue and past impermanent loss, $\beta_{1,t}$ and $\beta_{2,t}$, from the Fama-MacBeth regression (11). It shows that liquidity flows have become less sensitive to past past fee revenue. Maybe surprisingly, flow sensitivity to impermanent loss seems to be always smaller than flow sensitivity to fee revenue, even at year-end 2022.

5.2 Address-level analysis

To further understand the improvement of liquidity flow optimality over time, this subsection conducts additional analysis at the individual LP level, proxied by blockchain addresses.[15] In selecting LP addresses, I focus on the ones that provide liquidity to at least 10 pools from 2020 to 2022. I then draw a random sample of 10,000 addresses to limit computation. I run the following regression using a three-dimensional panel of these 10,000 LP addresses and

[15] I assume that each address represents one LP. In reality, a liquidity provider can use multiple addresses, each holding different LP tokens.

their liquidity flows across the top 100 liquidity pools from May 2020 to December 2022:

$$Flow_{i,p,t+1} = \tilde{\beta}_1 FeeRevenue_{p,t} \times PriorExperience_{i,t} \tag{14}$$

$$+ \tilde{\beta}_2 ImpermanentLoss_{p,t} \times PriorExperience_{i,t} \tag{15}$$

$$+ \tilde{\gamma} Controls_{p,t} + FE + \epsilon_{i,p,t+1} \tag{16}$$

where $Flow_{i,p,t+1}$ is defined as LP i's net dollar deposit to pool p from week-end t to week-end $t+1$. $PriorExperience_{i,t}$ captures LP i's prior experience up to week-end t, such as whether the LP has withdrawn liquidity and realized large losses. I control for observable pool characteristics (e.g. total value locked). I include pool fixed effects to purge out persistent level differences in liquidity flows across pools. I include LP address-by-week fixed effects to focus on the choice across pools for each individual LP at each point in time.

Why and how would prior experiences affect LP behavior? An LP may incur impermanent losses on a pool position but may not be aware of it, due to the shrouded display of this information (Figure 4). However, when exiting its position from the pool (e.g. due to idiosyncratic liquidity needs), the LP would be shown with the exact amount of assets returned, with all impermanent losses realized. This experience of impermanent loss, especially if large, would Bayesian update the LP's belief on the accounting of impermanent loss for return (Malmendier and Nagel, 2011, 2015).

Table 6 shows the results. Column 1 shows that, similar to pool-level results, liquidity flows are strongly affected by past fee revenues but much less so by past impermanent losses. Column 2 shows that LP actions are heterogeneous based on their prior experiences. In particular, LPs that have withdrawn liquidity and realized losses at worse than -10% in the past pay much more attention to impermanent losses when allocating their capital in the future. Relative to the average LPs, whose liquidity flows decrease by \$4.95 in response to 1 p.p. higher impermanent loss, the sensitivity of LPs that have realized large losses in the

past is $19.43 more. The results suggest that LPs learn from their experiences and modify their future actions accordingly.

Interestingly, this learning behavior is quite asymmetric. For LPs that have realized large gains (+10% or more) upon withdrawal in the past, they continue to ignore impermanent losses going forward, shown by the small insignificant coefficient on the interaction term (-0.93). Instead, when LPs realize large returns, their sensitivity to past fee revenue is amplified by $8.49 more. This further confirms that an investor's learning is based on her actual experience. The results also suggest that the learning process is somewhat naive: the investor will simply do more what has given her positive rewards and less what what given her negative ones.

6 Conclusion

In this paper, I provide an in-depth empirical analysis of liquidity returns and liquidity flows on AMMs, one of the largest financial innovations that have facilitated trillions of economic activities. I show suboptimal capital allocation in the first year of AMMs: LPs are highly sensitive to fee revenues – the part of return that is saliently marketed as APY – but not to impermanent losses – the part of return that is often shrouded. As a result, liquidity returns exhibit large anomalies: pools with higher realized net returns do not attract sufficient liquidity and continue to yield higher net returns in the future, whereas pools with higher fee revenues attract significantly more liquidity, even though these pools tend to also have higher impermanent losses and zero expected returns on net.

Importantly, inefficiencies decay over time. There is large reduction in both liquidity return predictability and liquidity flow suboptimality. The improvement concentrates in LPs who have experienced actual realization of large losses. This dynamic evolution of market equi-

librium on AMMs highlights the role of behavioral biases in adopting financial innovations, and how agents can overcome these biases by learning from personal experiences.

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.

G. Angeris, A. Evans, and T. Chitra. When does the tail wag the dog? curvature and market making, 2020. URL https://arxiv.org/abs/2012.08040.

P. Augustin, R. Chen-Zhang, and D. Shin. Reaching for yield in decentralized financial markets. 2022.

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.

B. M. BARBER, X. HUANG, T. ODEAN, and C. SCHWARZ. Attention-induced trading and returns: Evidence from robinhood users. *The Journal of Finance*, 77(6):3141–3190, 2022. doi: https://doi.org/10.1111/jofi.13183. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.13183.

I. Ben-David, F. Franzoni, and R. Moussawi. Exchange Traded Funds (ETFs). NBER Working Papers 22829, National Bureau of Economic Research, Inc, Nov. 2016. URL https://ideas.repec.org/p/nbr/nberwo/22829.html.

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.

I. Ben-David, F. Franzoni, B. Kim, and R. Moussawi. Competition for attention in the etf space. *The Review of Financial Studies*, 36(3):987–1042, 08 2022. ISSN 0893-9454. doi: 10.1093/rfs/hhac048. URL https://doi.org/10.1093/rfs/hhac048.

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.

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

L. W. Cong, X. Hui, C. E. Tucker, and L. Zhou. Scaling smart contracts via layer-2 technologies: Some empirical evidence. 2022.

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. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2325486.

E. F. Fama and K. R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. ISSN 0304-405X. doi: https://doi.org/10.1016/0304-405X(93)90023-5. URL https://www.sciencedirect.com/science/article/pii/0304405X93900235.

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.

X. Gabaix and D. Laibson. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*. *The Quarterly Journal of Economics*, 121(2):505–540, 05 2006. doi: 10.1162/qjec.2006.121.2.505. URL https://doi.org/10.1162/qjec.2006.121.2.505.

N. Gennaioli, A. Shleifer, and R. Vishny. Neglected risks, financial innovation, and financial fragility. *Journal of Financial Economics*, 104(3):452–468, 2012. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2011.05.005. URL https://www.sciencedirect.com/science/article/pii/S0304405X11001176. Market Institutions, Financial Market Risks and Financial Crisis.

C. R. Harvey, A. Ramachandran, and J. Santoro. *DeFi and the Future of Finance*. John Wiley & Sons, 2021.

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.

B. J. Henderson and N. D. Pearson. The dark side of financial innovation: A case study of the pricing of a retail financial product. *Journal of Financial Economics*, 100(2):227–247, 2011. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2010.12.006. URL https://www.sciencedirect.com/science/article/pii/S0304405X10002953.

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

G. Liao and D. Robinson. The dominance of uniswap v3 liquidity. 2022.

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.

U. Malmendier and S. Nagel. Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1):373–416, 02 2011. ISSN 0033-5533. doi: 10.1093/qje/qjq004. URL https://doi.org/10.1093/qje/qjq004.

U. Malmendier and S. Nagel. Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87, 10 2015. ISSN 0033-5533. doi: 10.1093/qje/qjv037. URL https://doi.org/10.1093/qje/qjv037.

J. Milionis, C. C. Moallemi, T. Roughgarden, and A. L. Zhang. Automated market making and loss-versus-rebalancing. 2022.

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.

P. Vokata. Engineering lemons. *Journal of Financial Economics*, 142(2):737–755, 2021. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2021.04.035. URL https://www.sciencedirect.com/science/article/pii/S0304405X21001653.

Figures and Tables

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

Panel A: Aggregate Trading Volume

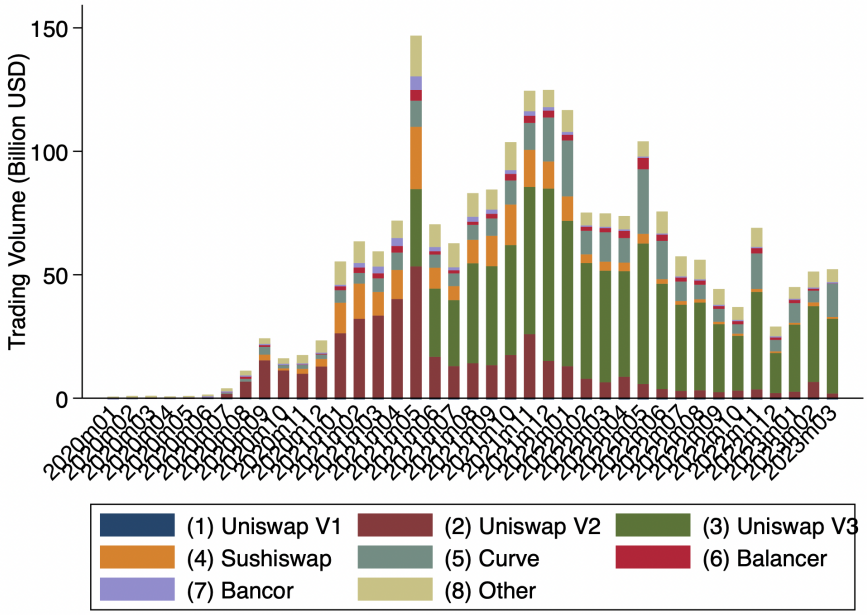

Panel B: Aggregate Liquidity

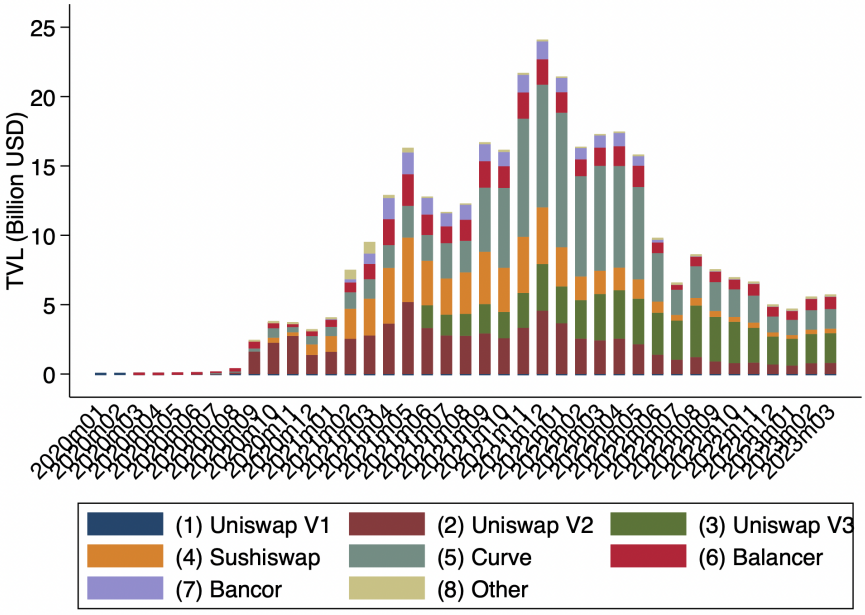

Figure 2: **Constant-Product Automated Market Maker**. X-axis is reserve of one token (e.g. BTC) and y-axis is reserve of the other token (e.g. USD). The solid red line plots the constant-product function, given by Equation 1. The black dashed lines show how amount of one token swapped in $(+\Delta X)$ corresponds to amount of the other swapped out $(-\Delta Y)$. The black solid lines show relative price of the two tokens $(P \equiv -\frac{\partial Y}{\partial X})$.

Panel A

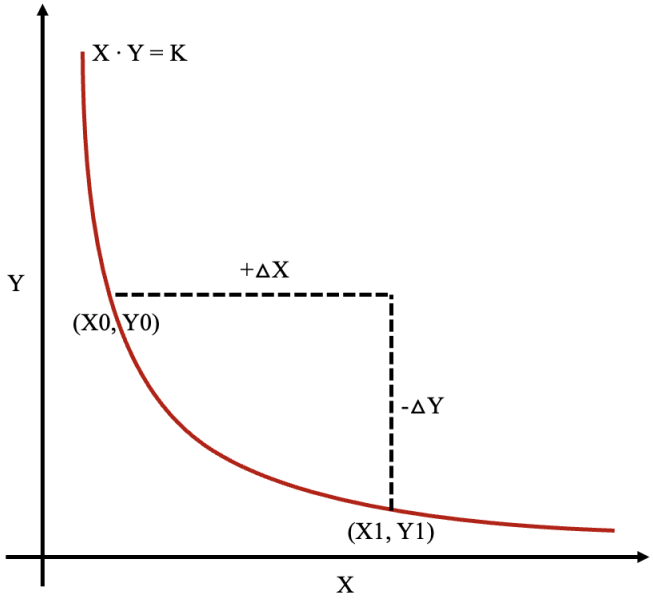

Panel B

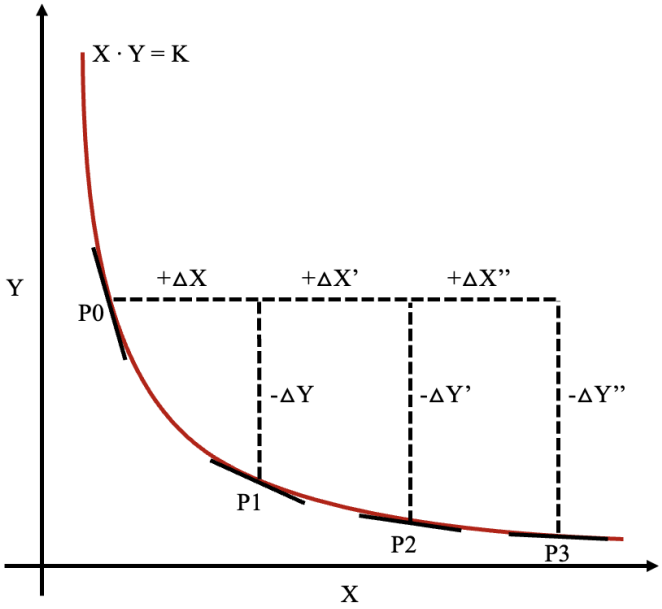

Figure 3: **Impermanent Loss on CPAMM**. X-axis is change in relative price of the two tokens in a CPAMM pool. Y-axis is impermanent loss, i.e. the opportunity cost of providing liquidity relative to holding on to the two tokens, assuming no fee revenue. The red line plots the deterministic relationship between impermanent loss and price change, given by Equation 6.

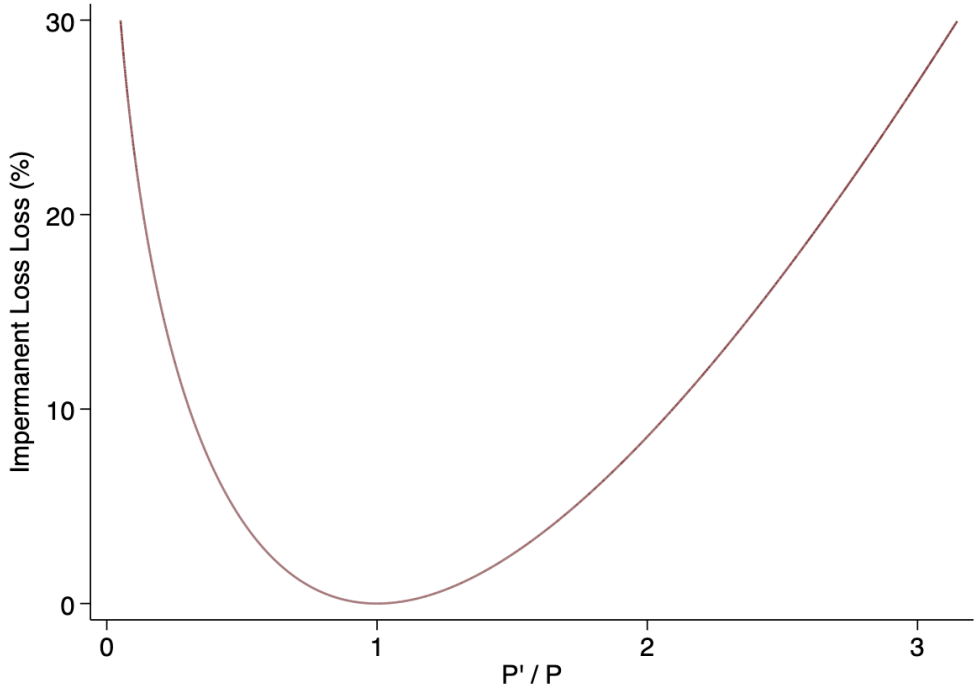

Figure 4: **The Display of Fee Revenue as APY**. The top panel shows how fee revenue is displayed on Uniswap V2's website. The bottom panel shows how fee revenue is displayed on SushiSwap's website. These screenshots were taken in October 2022.

Uniswap V2

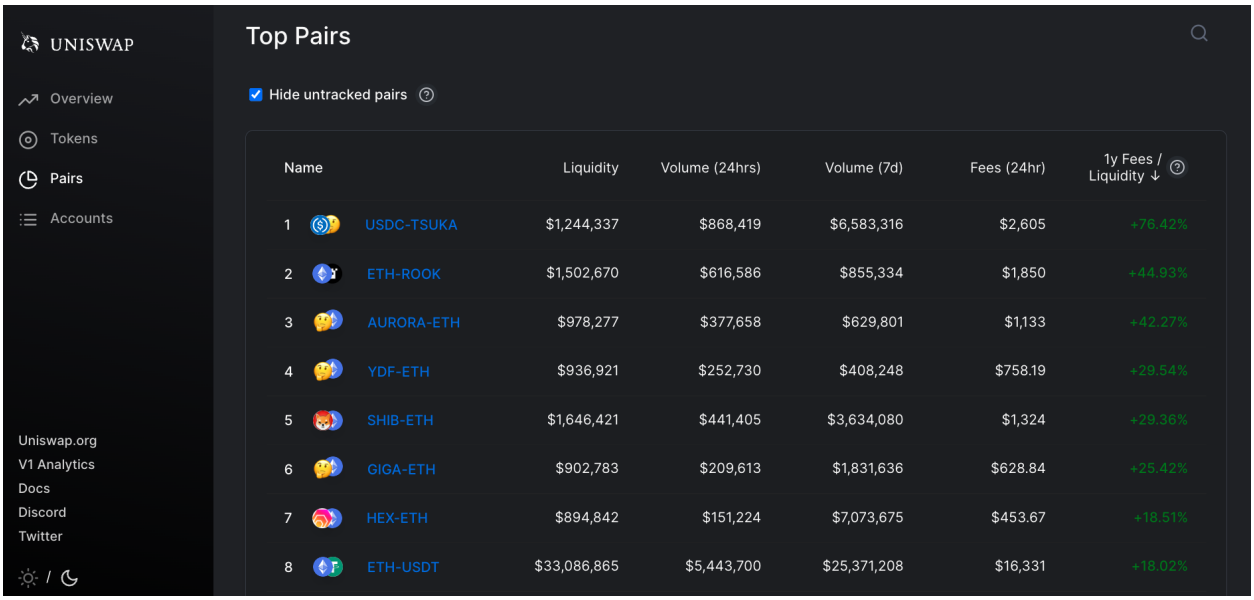

SushiSwap

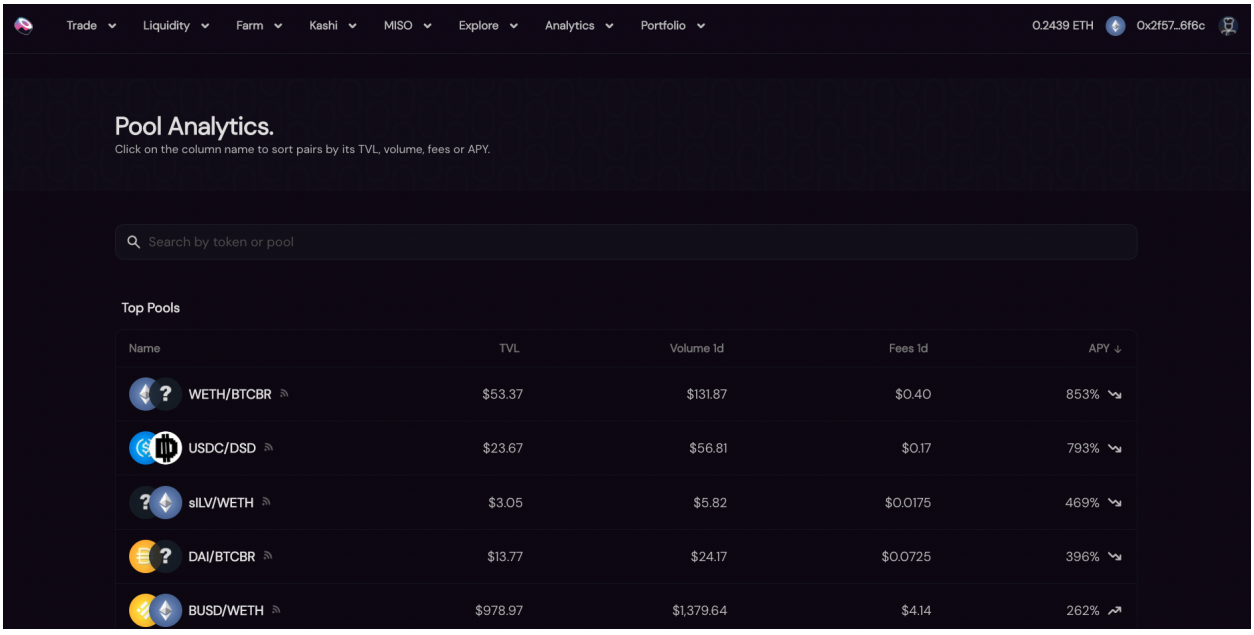

Figure 5: **Liquidity Misallocation over Time**. Panel A examines the evolution of liquidity return anomaly and shows 10-week rolling average of LP net return difference between sorted portfolios in Section 4.1. Panel B examines the evolution of liquidity flow suboptimality and shows 10-week rolling average of estimated flow sensitivity to fee revenue ($\beta_{1,t}$) and impermanent loss ($\beta_{2,t}$) for from the Fama-MacBeth (11).

Panel A: Return Predictability over Time

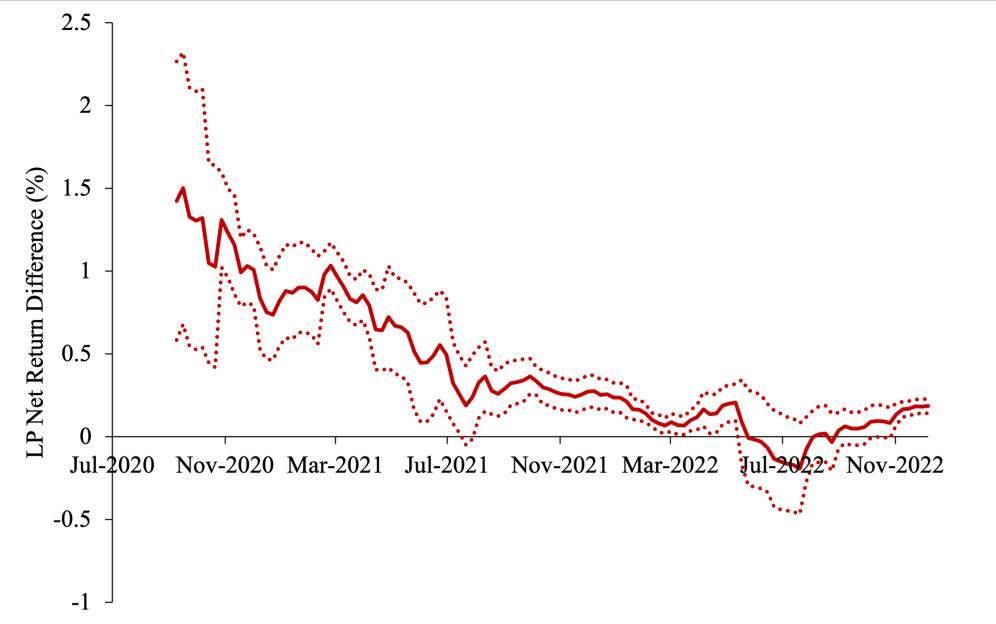

Panel B: Flow Sensitivity over Time

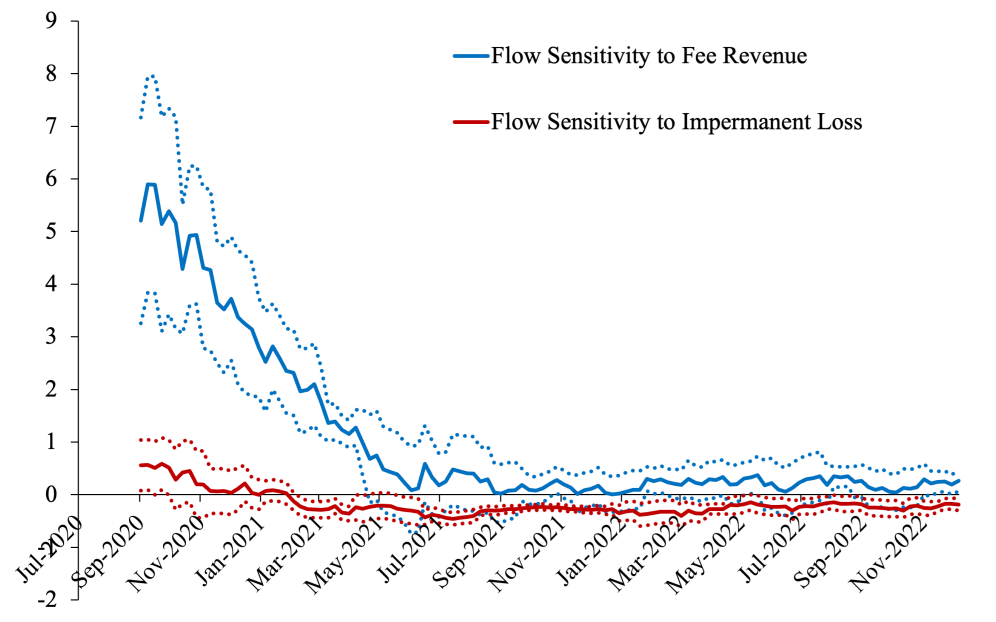

Figure 6: **Options-Implied Liquidity Compensation**. The lines show options-implied vs realized liquidity compensation for the BTC-USD pools and the ETH-USD pools. Options-implied liquidity compensation are derived based on Equation 8 and Monte Carlo simulation. Realized liquidity compensation refers to fee revenue, given by Equation 5.

BTC-USD

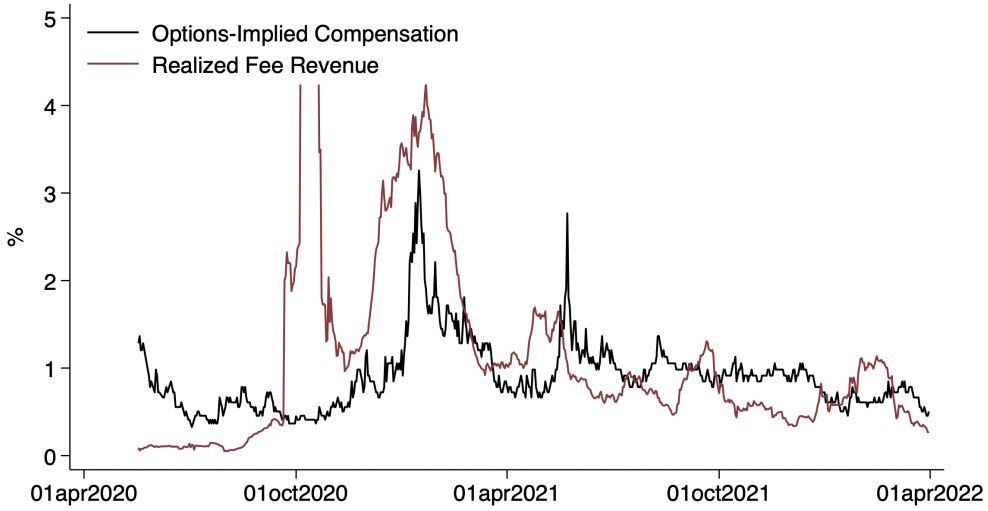

ETH-USD

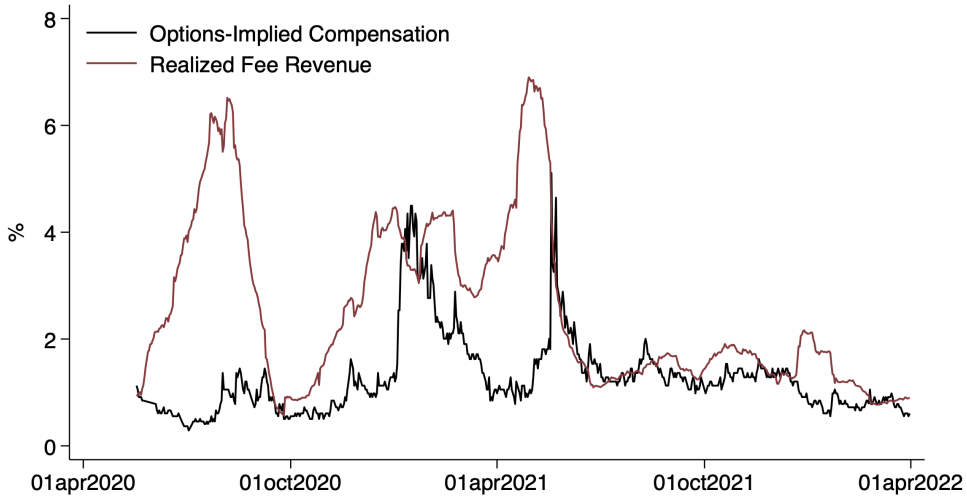

Table 1: **Liquidity Returns Sorted by Past Net Returns**. At the beginning of each week from May 2020 to December 2021, I sort my sample of Uniswap V2 pools into five quintiles based on their net returns over the past week, form a size-weighted portfolio for pools in each quintile, hold the portfolios for one week and dissolve them. The tables show the time series mean (and its standard error in parentheses) for each portfolio. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Panel A

Future Outcome (%, week t to t+1)	Fee Revenue	Impermanent Loss	Net Return Average	Net Return 3-Factor Alpha	Net Return 6-Factor Alpha
	(1)	(2)	(3)	(4)	(5)
Quintile of Past Net Return (week t-1 to t)					
Low	0.54	0.89	-0.35***	-0.33***	-0.33***
	(0.03)	(0.04)	(0.07)	(0.07)	(0.08)
2	0.21	0.28	-0.08***	-0.05*	-0.05
	(0.02)	(0.02)	(0.03)	(0.03)	(0.04)
3	0.17	0.23	-0.06*	-0.06*	-0.06
	(0.02)	(0.02)	(0.04)	(0.04)	(0.05)
4	0.28	0.27	0.01	-0.00	0.03
	(0.02)	(0.02)	(0.04)	(0.05)	(0.05)
High	0.61	0.32	0.29***	0.29***	0.29***
	(0.04)	(0.03)	(0.06)	(0.06)	(0.07)
High – Low	0.07	-0.58***	0.64***	0.62***	0.62***
	(0.06)	(0.07)	(0.08)	(0.09)	(0.11)

Panel B

Future Outcome (%, week t to t+1)	Net Return Volatility	Underlying Return		Transaction Cost
		Average	Volatility	
	(1)	(2)	(3)	(4)
Quintile of Past Net Return (week t-1 to t)				
Low	0.37	1.01	0.11	0.32
	(0.05)	(0.01)	(0.02)	(0.03)
2	0.21	1.01	0.10	0.24
	(0.03)	(0.01)	(0.01)	(0.02)
3	0.24	1.03	0.10	0.23
	(0.03)	(0.01)	(0.01)	(0.02)
4	0.25	1.02	0.09	0.29
	(0.04)	(0.01)	(0.01)	(0.02)
High	0.38	1.02	0.09	0.31
	(0.05)	(0.01)	(0.02)	(0.02)
High – Low	0.01	0.01	-0.02	-0.01
	(0.11)	(0.02)	(0.03)	(0.04)

Table 2: **Pool-Level Liquidity Flows**. The table shows determinants of net flows across liquidity pools using pooled regression 9 (with pool fixed effects and time fixed effects) and the Fama-MacBeth regression 11. The sample includes weekly observations of 875 Uniswap V2 pools with more than $100K in TVL from May 2020 to December 2021. Standard errors (two-way clustered by pool and by time for the pooled specification and robust for the Fama-MacBeth specification) 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: -0.05, SD: 17.00					
	Pooled			Fama-MacBeth		
	(1)	(2)	(3)	(4)	(5)	(6)
Past Net Return (%, t-1 to t)	0.28			0.82		
	(0.36)			(1.14)		
Past Fee Revenue (%, t-1 to t)		3.81***			2.25*	
		(1.12)			(1.24)	
Past Impermanent Loss (%, t-1 to t)		-0.35			0.85	
		(0.32)			(1.71)	
Past Volume (standardized, t-1 to t)			2.31***			2.29**
			(0.71)			(1.02)
Past Volatility (standardized, t-1 to t)			-0.82			-0.35
			(0.54)			(1.91)
Controls	Lagged Flow, Log TVL, Airdrop Indicator					
Observations	46108	46108	46191	46108	46108	46191
R2	0.39	0.39	0.39	0.05	0.06	0.07

Table 3: **Liquidity Returns Sorted by Past Fee Revenues Only**. At the beginning of each week from May 2020 to December 2021, I sort my sample of Uniswap V2 pools into five quintiles based on their fee revenues over the past week, form a size-weighted portfolio for pools in each quintile, hold the portfolios for one week and dissolve them. The table shows the time series mean (and its standard error in parentheses) for each portfolio. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Future Outcome (%, week t to t+1)	Fee Revenue	Impermanent Loss	Net Return Average	Net Return 3-Factor Alpha	Net Return 6-Factor Alpha
	(1)	(2)	(3)	(4)	(5)
Quintile of Past Fee Revenue (week t-1 to t)					
Low	0.13	0.23	-0.10	-0.05	-0.05
			(0.08)	(0.05)	(0.08)
2	0.28	0.41	-0.13	-0.06	-0.10
			(0.09)	(0.12)	(0.14)
3	0.51	0.63	-0.12	-0.22	-0.21
			(0.11)	(0.14)	(0.14)
4	0.79	0.53	0.25**	0.22	0.24
			(0.12)	(0.15)	(0.16)
High	1.30	1.12	0.18	0.22	0.21
			(0.17)	0	(0.21)
High – Low	1.17***	0.89***	0.28*	0.27	0.27
	(0.23)	(0.18)	(0.16)	(0.20)	(0.18)

Table 4: **Liquidity Returns Sorted by Past Fee Revenues and Past Volatilities**. At the beginning of each week from May 2020 to December 2021, I first sort my sample of 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 log price volatilities over the past week, form a size-weighted portfolio for pools in each sub-quintile, hold the portfolios for one week and dissolve them. The table shows the three-factor alpha (and its standard error in parentheses) of weekly returns in percentage points for each portfolio. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

	Future Net Return, Three-Factor Alpha (week t to t+1, %)				
Quintile (week t-1 to t)	Low Fee	2	3	4	High Fee
Low Volatility	0.04	0.03	0.05	0.20***	0.51***
	(0.06)	(0.04)	(0.05)	(0.05)	(0.08)
2	-0.09*	-0.07	-0.01	0.14***	0.05
	(0.05)	(0.09)	(0.04)	(0.05)	(0.21)
3	-0.11***	-0.14***	-0.10	0.06	0.07
	(0.03)	(0.06)	(0.07)	(0.07)	(0.14)
4	-0.26***	-0.29***	-0.24***	-0.18*	0.15
	(0.08)	(0.09)	(0.08)	(0.10)	(0.15)
High Volatility	-0.35***	-0.33***	-0.36***	-0.32**	-0.48**
	(0.10)	(0.13)	(0.11)	(0.14)	(0.20)
Low Vol – High Vol	0.39***	0.36***	0.41***	0.52***	0.99***
	(0.14)	(0.16)	(0.15)	(0.15)	(0.29)

Table 5: **Counterfactual Liquidity Returns Sorted by Past Net Returns**. This table repeats the sorted portfolio exercise in Table 1 using counterfactual liquidity derived from Equation (12) and Equation (13). *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Future Outcome (%, week t to t+1)	Actual			Counterfactual		
	Average	3F Alpha	6F Alpha	Average	3F Alpha	6F Alpha
	(1)	(2)	(3)	(4)	(5)	(6)
Quintile of Past Net Return (week t-1 to t)						
Low	-0.35***	-0.33***	-0.33***	-0.01	0.01	0.00
	(0.07)	(0.07)	(0.08)	(0.05)	(0.07)	(0.05)
2	-0.08***	-0.05*	-0.05	-0.04**	-0.04	-0.02
	(0.03)	(0.03)	(0.04)	(0.02)	(0.03)	(0.04)
3	-0.06*	-0.06*	-0.06	0.01	0.03	0.02
	(0.04)	(0.04)	(0.05)	(0.03)	(0.04)	(0.03)
4	0.01	-0.00	0.03	0.05**	0.05	0.05
	(0.04)	(0.05)	(0.05)	(0.02)	(0.03)	(0.03)
High	0.29***	0.29***	0.29***	0.32***	0.29***	0.35***
	(0.06)	(0.06)	(0.07)	(0.06)	(0.07)	(0.12)
High – Low	0.64***	0.62***	0.62***	0.33***	0.28***	0.35***
	(0.08)	(0.09)	(0.11)	(0.09)	(0.09)	(0.15)

Table 6: **Address-Level Liquidity Flows**. The table shows determinants of address-level liquidity flows through regression (14). The sample includes weekly observations of 10,000 sampled LP addresses that provide liquidity to at least 10 pools from May 2020 to December 2022. Standard errors are two-way clustered by address and by time and 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: -98.45, SD: 1083.63	
	(1)	(2)
Past Fee Revenue (%, t-1 to t)	22.98***	23.81***
	(7.12)	(7.05)
× 1(Prior Withdrawal at High Gain)		8.49*
		(4.93)
× 1(Prior Withdrawal at High Loss)		-5.67
		(3.75)
Past Impermanent Loss (%, t-1 to t)	-4.95	-4.05
	(5.38)	(4.99)
× 1(Prior Withdrawal at High Gain)		-0.93
		(5.15)
× 1(Prior Withdrawal at High Loss)		-19.43**
		(8.79)
Controls	Lagged Pool Flow, Log TVL, Airdrop Indicator	
Pool FE	Y	Y
LP Address FE × Week FE	Y	Y
Standard Errors	Clustered by Pool and by Week	
Observations	94,896,236	94,896,236
R2	0.33	0.33

Appendix A Additional Figures and Tables

Figure A1: **Illustration of impermanent Loss**. Panel A (B) figures give a numerical example of impermanent loss when BTC price goes up (down).

X = 1 and Y = 10,000

K = X × Y = 10,000

P = Y / X = 10,000

P = 10,000 → P' = 40,000

X' × Y' = K = 10,000

Y' / X' = P' = 40,000

X' = 0.5 and Y' = 20,000

HODL Value = 1 × 40,000 + 10,000 = 50,000

LP Value = 0.5 × 40,000 + 20,000 = 40,000

Impermanent Loss: 50,000 – 40,000 = 10,000 (or 50% of Original Value)

X = 1 and Y = 10,000

K = X × Y = 10,000

P = Y / X = 10,000

P = 10,000 → P' = 2,500

X' × Y' = K = 10,000

Y' / X' = P' = 2,500

X' = 2 and Y' = 5,000

HODL Value = 1 × 2,500 + 10,000 = 12,500

LP Value = 2 × 2,500 + 5,000 = 10,000

Impermanent Loss: 12,500 – 10,000 = 2,500 (or 25% of Original Value)

Figure A2: **Characteristics of Liquidity Positions**. Panel A shows distribution of size of liquidity positions (in thousand U.S. dollar) at year-end 2020, 2021, and 2022. Panel B shows cumulative distribution of duration of liquidity positions, defined as the number of days between when a liquidity position is first created and when it receives an update.

Panel A: Size of Liquidity Position

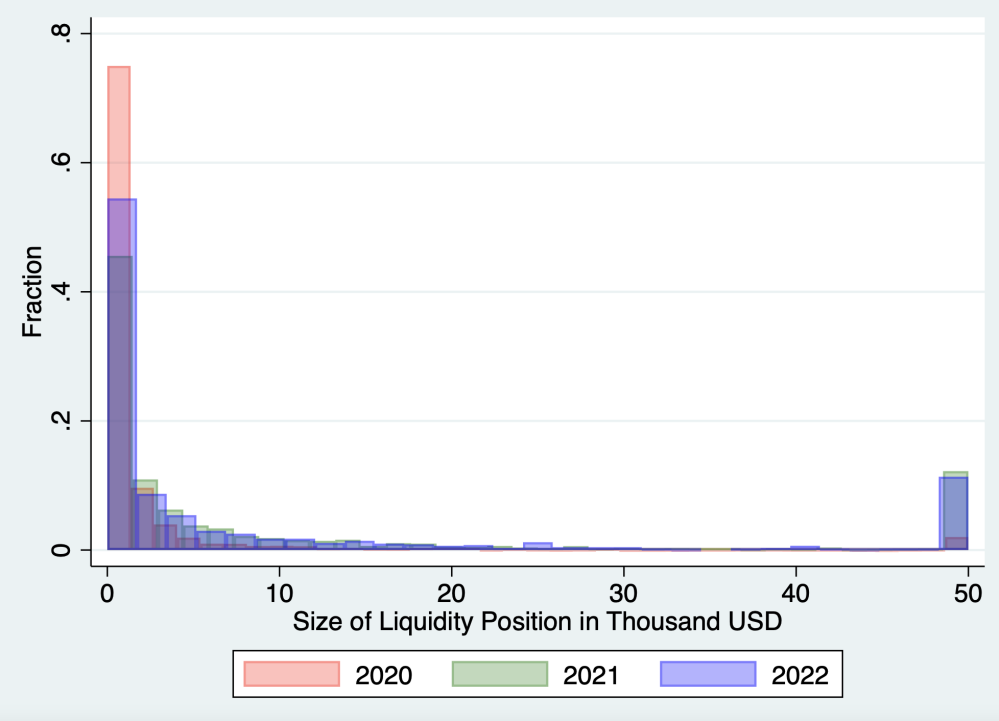

Panel B: Duration of Liquidity Position

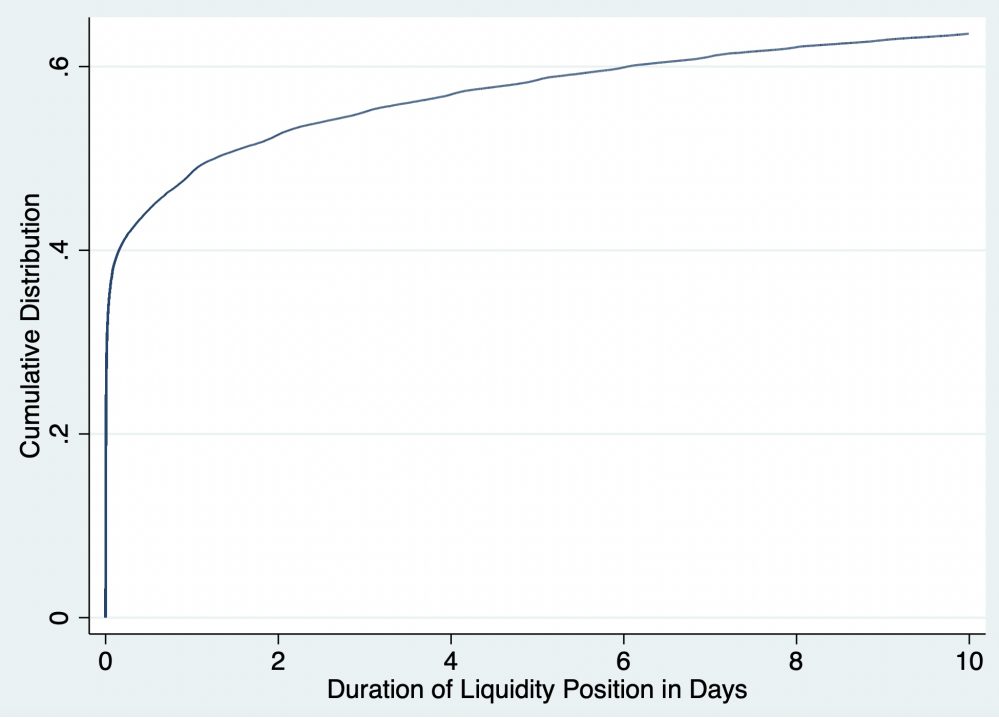

Figure A3: **Cross-Sectional Distribution of LP Net Returns**. The figure plots the cross-sectional distribution of weekly LP net returns across Uniswap V2 pools with at least $100K total value locked in the first week of September 2020, December 2020, and March 2021.

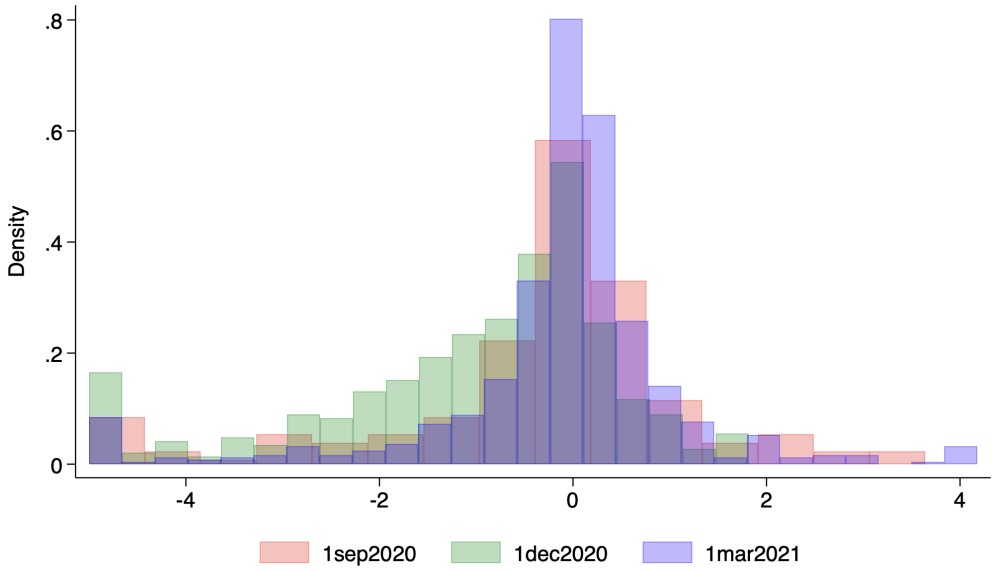

Figure A4: **Counterfactual Aggregate Liquidity**. The graph plots total value locked in Uniswap V2 liquidity pools, actual vs counterfactual, from May 2020 (when V2 was launched) to May 2021 (when V3 was launched). Counterfactual total value locked is calculated assuming that liquidity flows are as sensitive to impermanent losses as they are to fee revenues, using Equation (12) and Equation (13).

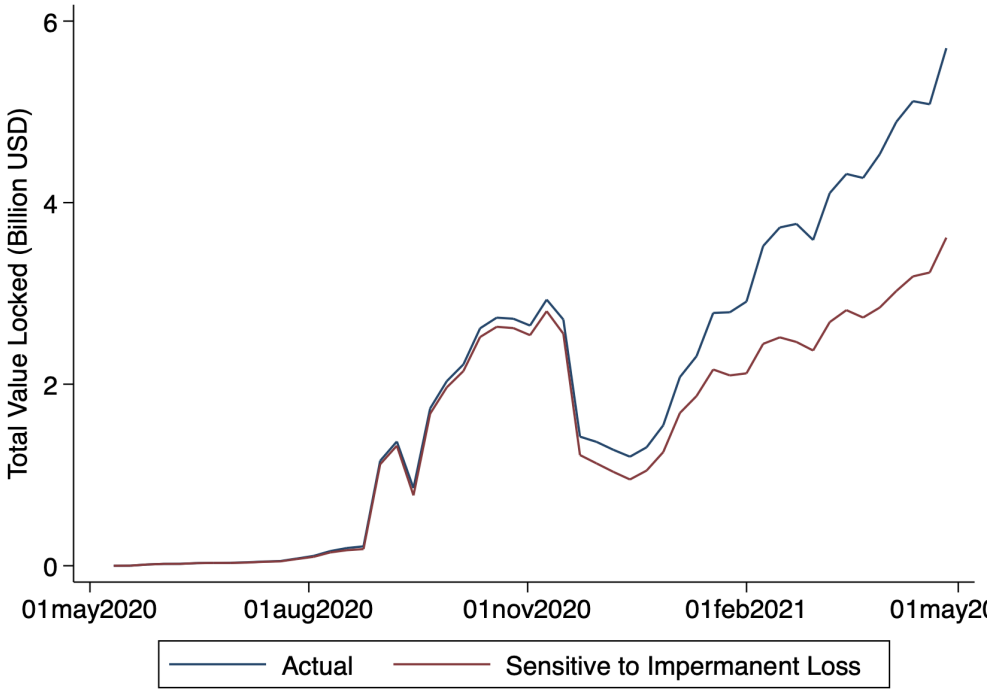

Table A1: **Summary Statistics**. The sample includes weekly observations of 875 Uniswap V2 pools from May 2020 to December 2022. All variables are winsorized at the 5th percentile in each cross section.

	N	Mean	SD	P25	P50	P75
Liquidity Flow (%)	47089	-0.42	15.28	-0.88	0.07	0.55
Fee Revenue (%)	47089	0.47	0.93	0.07	0.19	0.47
Adverse Selection Loss (%)	47089	1.13	4.53	0.03	0.18	0.69
Net Return (%)	47089	-0.66	4.24	-0.38	0.00	0.16
Trading Volume (%)	47089	139.10	202.49	23.37	62.02	157.69
Log Price Volatility	47089	8.68	8.74	3.59	6.30	10.61
Total Value Locked (mn $)	47089	6.04	37.76	0.50	1.08	2.49
Airdrop Indicator	47089	0.06	0.09	0.00	0.00	0.00

Table A2: **Determinants of Liquidity Flows, Robustness**. The tables show various alternative specifications of the pooled regression (9) (with pool fixed effects and time fixed effects) and the Fama-MacBeth regression (11). The sample includes daily or monthly observations of the top 20 or 100 Uniswap V2 pools with more than $100K in TVL from 2020 to 2022. Standard errors (two-way clustered by pool and by time for the pooled specification and robust for the Fama-MacBeth specification) are reported in parentheses. *, **, *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Panel A: Top 20 Pools

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: 10.51, SD: 118.76					
	Pooled			Fama-MacBeth		
	(1)	(2)	(3)	(4)	(5)	(6)
Past Net Return (%, t-1 to t)	-3.38			22.18		
	(3.31)			(14.73)		
Past Fee Revenue (%, t-1 to t)		15.48**			50.17**	
		(7.00)			(19.34)	
Past Impermanent Loss (%, t-1 to t)		3.67			-51.97	
		(3.41)			(80.90)	
Past Volume (standardized, t-1 to t)			7.00**			28.46**
			(3.33)			(11.00)
Past Volatility (standardized, t-1 to t)			4.33			-10.10
			(4.73)			(10.99)
Controls	Lagged Flow, Log TVL, Airdrop Indicator					
Observations	1639	1639	1639	1639	1639	1639
R2	0.41	0.42	0.42	0.29	0.36	0.36

Panel B: Top 100 Pools

Dependent Variable	Weekly Liquidity Flow (%, t to t+1) Mean: 1.69, SD: 45.73					
	Pooled			Fama-MacBeth		
	(1)	(2)	(3)	(4)	(5)	(6)
Past Net Return (%, t-1 to t)	0.95			2.52		
	(1.49)			(2.33)		
Past Fee Revenue (%, t-1 to t)		12.77**			6.59*	
		(5.84)			(3.85)	
Past Impermanent Loss (%, t-1 to t)		-0.94			3.91	
		(1.57)			(6.79)	
Past Volume (standardized, t-1 to t)			6.92**			4.54*
			(2.87)			(2.48)
Past Volatility (standardized, t-1 to t)			-1.70			0.18
			(1.11)			(4.62)
Controls	Lagged Flow, Log TVL, Airdrop Indicator					
Observations	7204	7204	7232	7204	7204	7232
R2	0.44	0.44	0.44	0.09	0.11	0.12

Panel C: Daily Flows

Dependent Variable	Daily Liquidity Flow (%, t to t+1) Mean: 0.08, SD: 8.26					
		Pooled			Fama-MacBeth	
	(1)	(2)	(3)	(4)	(5)	(6)
Past Net Return (%, t-7 to t)	0.01			0.07		
	(0.02)			(0.27)		
Past Fee Revenue (%, t-7 to t)		0.87**			0.65***	
		(0.38)			(0.24)	
Past Impermanent Loss (%, t-7 to t)		-0.02			-0.20	
		(0.03)			(0.38)	
Past Volume (standardized, t-7 to t)			0.79**			1.54
			(0.37)			(2.05)
Past Volatility (standardized, t-7 to t)			-0.13			1.16
			(0.17)			(1.26)
Controls	Lagged Flow, Log TVL, Airdrop Indicator					
Observations	323155	323155	323747	323155	323155	323747
R2	0.33	0.33	0.33	0.05	0.07	0.07

Panel D: Monthly Flows

Dependent Variable	Monthly Liquidity Flow (%, t to t+1) Mean: -0.57, SD: 64.53					
		Pooled			Fama-MacBeth	
	(1)	(2)	(3)	(4)	(5)	(6)
Past Net Return (%, t-1 to t)	0.37			4.76		
	(0.30)			(4.78)		
Past Fee Revenue (%, t-1 to t)		4.01***			11.23*	
		(1.25)			(6.56)	
Past Impermanent Loss (%, t-1 to t)		-0.42			-4.65	
		(0.27)			(4.54)	
Past Volume (standardized, t-1 to t)			9.73***			20.01*
			(3.37)			(11.43)
Past Volatility (standardized, t-1 to t)			-3.21*			-9.35
			(1.56)			(6.05)
Controls	Lagged Flow, Log TVL, Airdrop Indicator					
Observations	9873	9873	9905	9873	9873	9905
R2	0.48	0.48	0.48	0.17	0.20	0.21

Table A3: **Replication of Existing Literature**. Panel A replicates Table 2 of Lehar and Parlour (2021). Panel B replicates Table 2 of Capponi and Jia (2021). Standard errors are reported in parentheses. *, **, and *** denote p-values less than 0.10, 0.05, and 0.01, respectively.

Panel A: Table 2 of Lehar and Parlour (2021)

Dependent Variable	Daily Total Value Locked (million $)	
	Pooled	Fama-MacBeth
Past Volume (million $, t-30 to t)	0.16***	0.27***
	(0.01)	(0.01)
Past Volatility (%, t-30 to t)	-0.12***	-0.14***
	(0.02)	(0.00)
Airdrop Indicator	260.96***	37.60***
	(63.98)	(4.64)
Observations	326252	326252
R2	0.66	0.62

Panel B: Table 2 of Capponi and Jia (2021)

Dependent Variable	Weekly Liquidity Flow (%, t to t+1)	
	Pooled	Fama-MacBeth
Past Volume (%, t-1 to t)	0.02*	0.02*
	(0.01)	(0.01)
Past Volatility (%, t-1 to t)	-0.53*	-0.72*
	(0.27)	(0.40)
Observations	468	468
R2	0.18	0.28