

Behavior of Liquidity Providers in Decentralized Exchanges

Lioba Heimbach*, Ye Wang* and Roger Wattenhofer
{hlioba,wangye,wattenhofer}@ethz.ch
ETH Zurich

*The first two authors have equal contribution.

Abstract—Decentralized exchanges (DEXes) have introduced an innovative trading mechanism, where it is not necessary to match buy-orders and sell-orders to execute a trade. DEXes execute each trade individually, and the exchange rate is automatically determined by the ratio of assets reserved in the market. Therefore, apart from trading, financial players can also act as liquidity providers. Liquidity providers earn transaction fees of trades executed in DEXes. Although liquidity providers are essential for the functionality of DEXes, their behavior remains largely unstudied.

In this paper, we aim to understand how liquidity providers react to market information. Further, we analyze the returns and risks associated with providing liquidity in DEXes. We measure the operations of liquidity providers on Uniswap and analyze how they determine their investment strategy based on market changes. We also reveal the returns and risks of investments in different trading pair categories, i.e., stable pairs, normal pairs, and exotic pairs. Further, we investigate the movement of liquidity between trading pools. To the best of our knowledge, this is the first work that systematically studies the behavior of liquidity providers in DEXes.

I. INTRODUCTION

Traditionally, trading is executed on centralized exchanges (CEXes), using the limit order book mechanism. With this mechanism, each seller is matched with a buyer for a trade. In the crypto-space, when traders want to exchange Ether (ETH) with Bitcoin (BTC) through a CEX, they have to transfer their ETH to the account of the centralized operator, submit their sell orders of ETH, wait for matching corresponding buy orders of ETH, and then withdraw the incoming BTC from the market operators after the execution of their orders. This process is both a cumbersome and risky process as the CEX temporarily holds the funds.

Recently, an alternative in the form of decentralized exchanges (DEXes) surfaced. Currently, DEXes are popular in cryptocurrency markets. Eventually, fiat currencies, stocks, or other commodities might trade on DEXes as well. So understanding DEXes may be crucial beyond just cryptocurrency applications.

DEXes are smart contracts running on a blockchain. In contrast to the limit order book mechanism, traders do not need to be matched to a trading partner with opposite intentions. In a DEX, trades are completed immediately when the orders are recorded on the blockchain.

Uniswap is the most popular DEX. As an example, consider trades between BTC and ETH. The Uniswap platform offers a

smart contract (liquidity pool) with locked-in funds for these two cryptocurrencies. When a trader wants to exchange ETH for BTC, the trader only needs to send their ETH to the smart contract. The smart contract will then immediately send the appropriate amount of BTC back to the trader, while the ETH sent by the trader will be locked in the smart contract. The exchange rate is primarily determined by the ratio of BTC to ETH stored in the smart contract. The BTC and ETH in the smart contract are providing the liquidity for the trades between BTC and ETH on Uniswap.

In a DEX pool, liquidity is provided by liquidity providers. These liquidity providers lock their cryptocurrencies into the corresponding liquidity pools. DEXes generally charge a percentage transaction fee for each trade executed on the platform. Liquidity providers share transactions fees in proportion to their liquidity contributions. Therefore, with the emergence of DEXes, users have new investment opportunities in the cryptocurrency ecosystem: they can offer their assets in DEXes as liquidity and benefit from transaction fees.

Many questions emerge: How many different pools are liquidity providers invested in? What are the expected earnings from transaction fees? How do liquidity providers react to market forces? Do liquidity providers redistribute their assets? As DEXes are an emerging phenomenon that might cross over to other markets beyond cryptocurrencies, it is interesting to understand the motives of liquidity providers.

In this paper, we quantitatively measure the behavior of liquidity providers. We collect data from Uniswap, the most popular DEX in the cryptocurrency ecosystem. With this data, we study liquidity provider strategies, returns, and risks. First, we analyze the distribution of liquidity providers in Uniswap, examining the creation of liquidity pools, the distribution of liquidity, and the participation of liquidity providers in these pools. Although Uniswap allows users to create liquidity pools between any pair of tokens, more than 80% of liquidity pools include ETH, and six popular tokens dominate the market. Moreover, more than 60% liquidity is locked in the top 24 pools. Cryptocurrency holders act restrained concerning these new investments: approximately 70% of providers reserve their liquidity in a single pool. However, these seemingly conservative liquidity providers contribute more than 50% of the liquidity in the most popular pools. This observation indicates that individual providers, as opposed to professional market makers, control the liquidity of DEXes. In examining

the addition and removal of liquidity day by day, we find that the market change of liquidity is relatively stable, i.e., the correlation between the number of injections and withdrawals is 0.922.

To understand the behavior of liquidity providers, we study the risks and returns they face across three different pool categories introduced by Uniswap: stable, normal, and exotic. We find that while stable and normal pairs may provide attractive investment opportunities for liquidity providers depending on the individual risk tolerance and return expectations, the exotic pools investigated do not. Those demonstrated extremely negative returns accompanied by high risk.

Liquidity providers do not frequently move their assets across different pools, seemingly indifferent to price changes or other market indicators. However, we observe that many traders redistribute their liquidity when given the additional benefits of liquidity mining on top of transaction fees. We, thus, study what external factors are influencing the behavior of liquidity providers.

This paper makes the following contributions. First, we conduct a systematic investigation on the liquidity providers in Uniswap to outline the participants of users in such emerging trading activities. Second, we classify three categories of liquidity pools and analyze the investment returns and risks of liquidity providers in those pools. We present the variability of investment strategies for different types of liquidity pools. Finally, we demonstrate the movement of liquidity around the entire DEX and show the significant influence of external factors on liquidity providers, such as liquidity mining activities.

II. BACKGROUND AND RELATED WORK

A. Decentralized exchanges

Decentralized exchanges (DEXes) run as smart contracts on the blockchain and generally allow automatic trading of cryptocurrencies by an algorithm using liquidity provided in liquidity pools. Most DEXes adopt a similar automated market maker (AMM) mechanism as Uniswap, the largest DEX. Uniswap allows for the creation of liquidity pools between any pair of crypto tokens. Individual liquidity providers deposit both tokens at equal value in the pool, and these funds allow traders to swap the respective tokens. In the process, traders pay a relative transaction fee. The fee is then disrupted pro-rata amongst the pool's liquidity providers.

The exchange rate offered to a trader is determined by the smart contract. Uniswap uses the constant product market maker, i.e., the algorithm ensures that the product of the two pool reserves remains constant. Consider the liquidity pool between token A and token B , $A \rightleftharpoons B$, with respective reserves a_t and b_t at time t . A trader wishing to exchange δ_a tokens A at time t will receive δ_b tokens B , where

$$\delta_b = \frac{b_t(1-f)\delta_a}{a_t + (1-f)\delta_a},$$

and f is the transaction fee. The fee is charged on the input amount and is currently 0.3% in the case of Uniswap. This fee is distributed to liquidity providers.

B. Liquidity in DEXes

This paper aims to study liquidity providers in DEXes quantitatively. In this section, we review previous work in two categories: quantitative studies of market behavior in DEXes, and theoretical studies on liquidity providers in DEXes.

Since all transaction information is broadcast through the blockchain network and all transactions in DEXes are public to all market participants, researchers can study the market behavior of traders in cryptocurrency ecosystems. Chen et al. [1] provide a basic view of the ERC20 tokens on Ethereum. They visualize how cryptocurrencies are created, held, and transferred by traders. Other works explore more detailed behavior of traders in the markets. Daian et al. [2] measures front-running trades on DEXes. Because of the transparency and latency of DEXes transactions, traders can observe profitable transactions before their execution and front-run these orders with higher fees to front-run the target victim. Such behaviors result in a high miner-extractable value, which brings systemic consensus-layer vulnerabilities. Torres et al. [3] and Qin et al. [4] quantify the revenue of traders who conduct a combination of front-running and back-running transactions, i.e., sandwich attacks [5], demonstrating that normal transactions are vulnerable to arbitrageurs. Wang et al. [6] study another arbitrage behavior, i.e., cyclic arbitrages (triangular arbitrages) in Uniswap, which profits traders with price differences across different trading pools. They claim that implementing transactions with private smart contracts is more resilient to front-running attacks than directly calling public functions of DEXes smart contracts. Although previous studies have provided an in-depth understanding of traders in the cryptocurrency ecosystems, they mainly focus on the behavior of normal traders. The novel trading option of providing liquidity on DEXes has not been well studied. Therefore, this paper aims to fill the existing research gap and inspire more work in this direction.

A separate line of work has analyzed the returns of liquidity providers theoretically with microstructure models. Evans [7] studies the returns of liquidity providers when there is no transaction fee charged and claims that it is better to invest in constant-mix portfolios than providing liquidity. Moreover, Evans et al. [8] develop a framework for determining the optimal transaction fees for AMM DEXes and show that providing liquidity to DEXes is preferable to all alternative trading strategies as fees approach zero. Aoyagi [9] conducts a game-theoretical analysis between liquidity providers and traders to estimate the returns of liquidity providers in DEXes. These studies have ideal assumptions on information delivery and price fluctuations of assets, which do not coincide with the real scenarios. Moreover, they have not considered how liquidity providers choose between different trading pools. To the best of our knowledge, none of the previous studies have empirically analyzed the trading behavior of liquidity providers in real DEXes. To better understand the new trading option in DEXes, we measure liquidity provider behavior in Uniswap and study how they react to market indicators, which

will inspire better theoretical models for analyzing liquidity providers in DEXes.

III. DATA DESCRIPTION

We analyze Uniswap data to measure the liquidity provider behavior on the platform and to understand the influence of market factors on their trading strategies. Currently, there are three Uniswap versions, V1, V2, and V3. Uniswap V2 was launched on 4 May 2020 and succeeded by Uniswap V3 on 5 May 2021 and was the predominant market throughout the past year. Thus, we study Uniswap V2 as an example in this paper as it provides us with the biggest data set. Furthermore, other DEXes, e.g., Sushiswap, have similar mechanisms as Uniswap V2.

To collect data, we develop and launch a modified version of go-ethereum client, which exports all transactions executed on Uniswap V2. Further details concerning the mechanism of DEXes and related data extraction of the go-ethereum client can be found in Appendix A. We collect all transactions recorded on Ethereum from block 10000835 (where Uniswap V2 has been deployed, on 4 May 2020) to block 11709847 (on 23 January 2021)¹.

To further evaluate the behavior of liquidity providers and its dependence on pool characteristics on Uniswap, we require the price of each cryptocurrency in a common currency – US\$ in our case. As many cryptocurrencies do not have a market price, we first compute the price of each cryptocurrency studied in ETH. This data is obtained from the common pool with ETH of the cryptocurrency in question. Such a pool exists for most cryptocurrencies and we assume the prices to be accurate, as they would offer arbitrage opportunities otherwise. The ETH price then allows us to calculate the value of each token in US\$ historical data for the ETH price in US\$ from coinbase².

IV. LIQUIDITY DYNAMICS IN DEXES

In this section, we provide an overview of liquidity pools on Uniswap. Unlike traditional centralized exchanges where only the cryptocurrencies permitted by operators can be traded on the platforms, DEXes allow any Ethereum users to create trading pools between any pair of tokens. As shown in Figure 1, the number of liquidity pools increases quickly with the emergence of Uniswap and continues to rise at a rate of over 100 per day until November 2020. After that, the daily growth of liquidity pools has gradually slowed down and keeps at a rate of 50 pools per day by the end of January 2021. In total, we find 29,235 available liquidity pools, which involve 25,231 cryptocurrencies, while 22,828 tokens only have a single trading pool. The most popular token is ETH, 24,011 tokens share a liquidity pool with ETH. Except for ETH, the most popular cryptocurrencies are USDT (1321), USDC (627), DAI (542), UNI (270), and WBTC (148). These tokens are either stable coins whose value is anchored at \$1 or tokens with high value. Until 23 January 2021, Uniswap

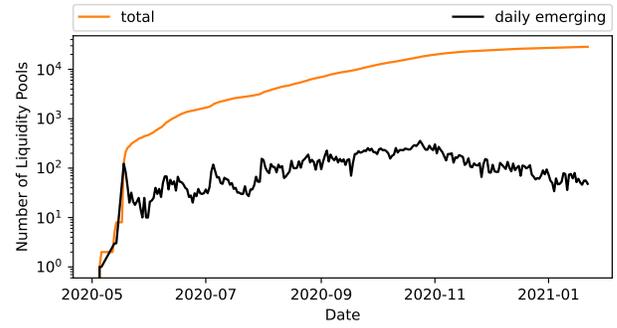


Figure 1: Number of liquidity pools on Uniswap from 4 May 2020 to 23 January 2021. The orange line represents the number of all liquidity pools on Uniswap, while the black line shows the number of daily emerging pools.

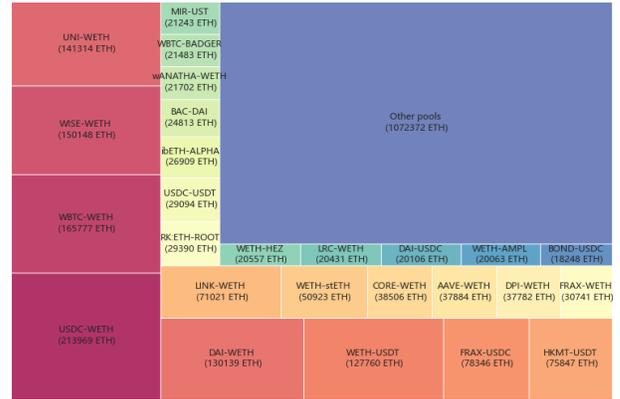


Figure 2: Distribution of liquidity in Uniswap on 23 January 2021.

has reserved more than \$2.5 billion and more than 60% of liquidity is located in the top 24 popular pools (Figure 2).

Apart from the distribution of liquidity pools, we also measure the participation of liquidity providers in Uniswap. Although some traders may use several addresses to provide liquidity on Uniswap, more than 82% of Ethereum users only control a single account [10]. Therefore, in this paper, we consider each address as a single liquidity provider. In total, we find 183,823 addresses add liquidity to liquidity pools on Uniswap, while 107,352 of them keep reserving their tokens in Uniswap by 23 January 2021 (Figure 3). Initially, most liquidity providers reserve their money only in one liquidity pool, and the average number of pools they participate in is less than 1.4 until September 2020. Later on, during October and November, liquidity providers are interested in more pools, as we see the average number of pools providers are involved in increases to 1.8.

As shown in Figure 4, the distribution of the number of liquidity pools that each liquidity provider participates in follows the power law. The largest providers participate in 120 pools, and more than 96.5% of the accounts reserve their money in no more than five pools.

We notice that most liquidity providers only participate in a single pool (Figure 3). This pattern also holds for the largest Uniswap liquidity pools. We see that more than 70%

¹Block data is available at <https://www.ethereum.org/>

²<https://www.coinbase.com/>

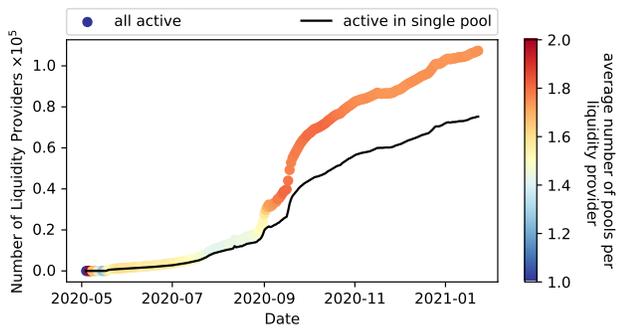


Figure 3: The number of liquidity providers over time. The black line demonstrates the number of providers that only participate in one pool. The colorful dotted line is the number of active liquidity providers in Uniswap, and the color indicates the average number of pools that liquidity providers participate.

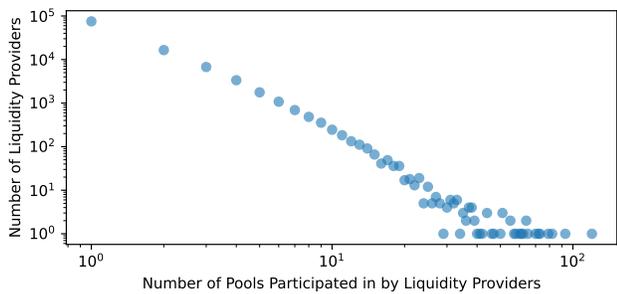
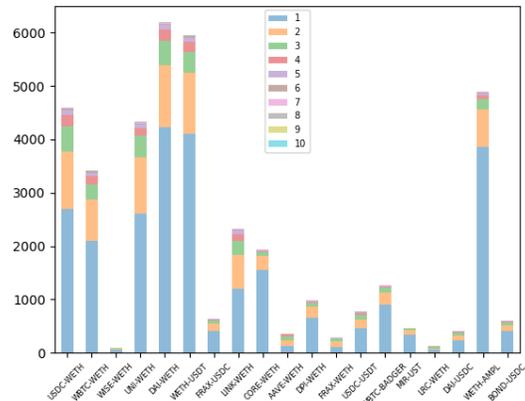


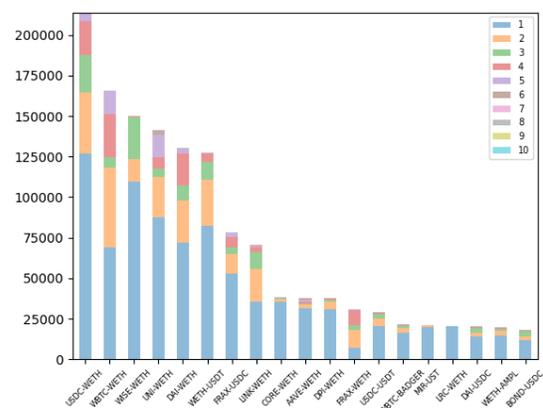
Figure 4: Distribution of the number of liquidity pools that each liquidity provider participates in. About 10,000 only participate in a single pool (top left), and some providers participate in many pools (bottom right).

of providers reserve their cryptocurrency assets only in one of them. Further, only a few providers spread their assets across more than ten pools (Figure 5a). Liquidity providers that reserve their cryptocurrency assets only in one pool have taken a serious role in the DEXes ecosystem, as they provide more than half liquidity in these popular pools (Figure 5b). This fact suggests that DEXes are not controlled by oligopoly and professional market makers but that ordinary users contribute the most to the operation of the decentralized market mechanism.

During the nine analyzed months, liquidity providers have gradually added and removed liquidity from different pools, while they reserve their tokens 1,011,524 times on Uniswap (Figure 6a) and withdraw them for 527,429 times (Figure 6b). We find that the day-by-day correlation between the number of liquidity injections and liquidity withdrawals is very high (0.992), indicating that the market size is growing steadily. Moreover, in terms of days, the market behaviors of liquidity providers are relatively consistent. Almost half of the liquidity operations from August to September take place in the most popular pools. However, these pools become less active from October onwards. On the one hand, we may consider the liquidity in these popular pools to become stable after five months of the development. On the other hand, this fact also



(a) Number of liquidity providers per pool. The DAI-WETH pool has more than 6,000 individual liquidity providers. More than 4,000 of them only contributed to this one (1) pool. The other colors represent providers which are providing liquidity to more than one (2,3,...) pool.



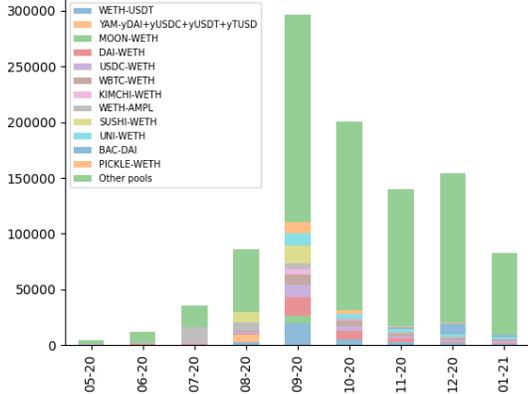
(b) Total liquidity per pool. The biggest pool is USDC-WETH with more than 200,000 ETH liquidity. About 125,000 ETH is funded by liquidity providers which only contributed to this one (1) pool. The other colors represent providers which are providing liquidity to more than one (2,3,...) pool.

Figure 5: Liquidity providers and liquidity per pool.

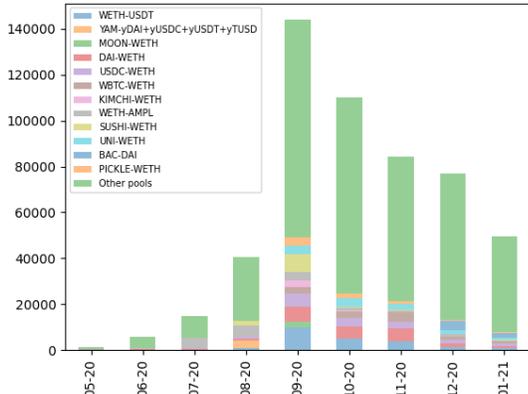
indicates a rapid change of interests in the liquidity pools of investors in the cryptocurrency ecosystems.

We have specified 12 pools with the most liquidity injections and withdrawals in Figure 6a and Figure 6b. As most liquidity operations occur in these pools, the patterns of liquidity changes in these pools can represent the trading behavior of liquidity providers in DEXes well. For instance, we observe apparent differences between the three types of liquidity pools: normal, stable, and exotic.

In the normal liquidity pools, both cryptocurrencies traded are established currencies, such as $USDC \rightleftharpoons WETH$, $WBTC \rightleftharpoons WETH$, and $DAI \rightleftharpoons WETH$. These cryptocurrencies are recognized in the cryptocurrency ecosystems. To support trades between these cryptocurrencies, the liquidity in the pools is relatively high (Figure 2). Although the price of normal tokens may fluctuate, the price trend is relatively stable and in line with the development of the cryptocurrency ecosystem.



(a) Number of mint events per month executed in different liquidity pools in Uniswap. In September 2020 there were more than 300,000 mint events in total.



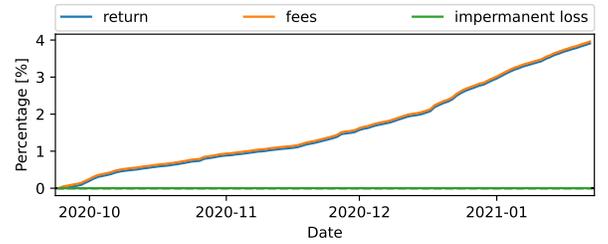
(b) Number of burn events per month executed in different liquidity pools in Uniswap. In September 2020 there were more than 140,000 burn events.

Figure 6: Mint and burn events.

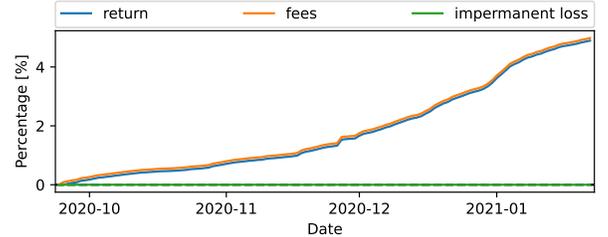
In stable pools, both tokens traded in the pool are stable coins, such as $USDC \rightleftharpoons USDT$ and $DAI \rightleftharpoons USDC$. The price fluctuations of stable coins are negligible. Thus, the market environment has little impact on stable pool liquidity providers. Since the price of stable coins generally remains constant, liquidity providers earn profits by charging transaction fees.

The remaining pools are referred to as exotic pools. In such pools, the price of one trading token is extremely volatile. The price of these tokens changed by more than a hundred times during our measuring period. As we show in Figure 2, although many liquidity operations have been applied on these pools, they are not part of the high liquidity value by 23 January 2021, as little liquidity remains in the pool – for example, YAM, MOON, and KIMCHI. For these exotic cryptocurrencies, liquidity providers take more risks because of their dramatic price fluctuations. When the price of the exotic tokens changes drastically, the other coin they have reserved in the pool may be emptied instantly by other traders.

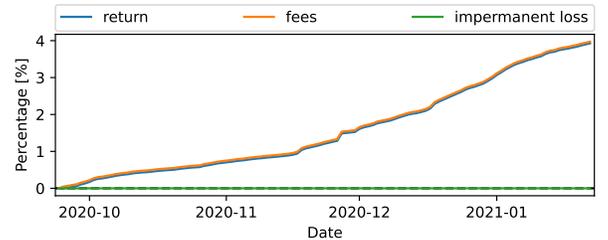
Given the significant differences of these three kinds of



(a) $USDC \rightleftharpoons USDT$



(b) $DAI \rightleftharpoons USDT$



(c) $DAI \rightleftharpoons USDC$

Figure 7: Evolution of cumulative returns, fees and impermanent loss over four months for stable pairs. Stable pairs do not suffer from impermanent loss.

pools, we infer that the distribution of liquidity providers, their trading strategies, and the investment return across these three categories differ. Therefore, we measure the activities of liquidity providers in these pools separately and analyze how they react to market changes.

V. RETURNS AND RISKS OF PROVIDING LIQUIDITY

Evans et al. [7] suggest that the returns of providing liquidity are lower than investing in the constant-mix portfolios if the transaction fee is zero. However, in reality, liquidity providers can benefit from their contribution to the liquidity pools through strictly positive transaction fees, which results in different returns and risks compared to the previous theoretical analysis. To better understand liquidity providers' motivation in DEXes, we compare the return received by liquidity providers to holding the respective assets according to the given token ratio during the initial liquidity injection.

The return of a liquidity provider between time t_1 and t_2 in percent is given as

$$\text{return}_{t_1 \rightarrow t_2} = 100 \cdot \frac{\text{invest}_{t_2} - \text{hold}_{t_2}}{\text{hold}_{t_2}},$$

where invest_t , which is the current value in US\$ of the liquidity placed in the pool at time t and hold_t is the value in US\$ of the constant-mix portfolio at time t .

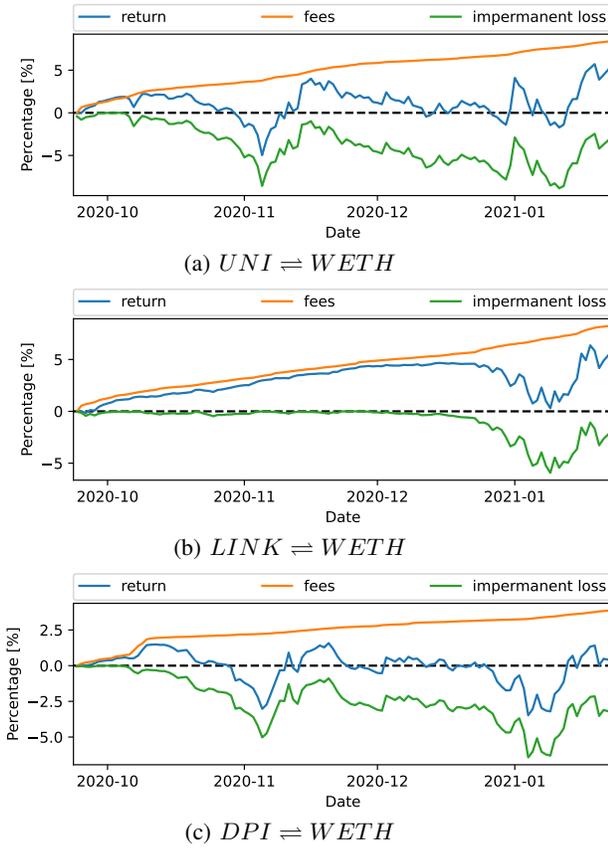


Figure 8: Evolution of cumulative returns, fees and impermanent loss over four months for normal pairs. A liquidity provider funding LINK-WETH in November, and pulling out in January would have suffered a loss, since the accumulated fees do not make up for the impermanent loss.

This return is positively influenced by the fees collected from trades performed in the pool and negatively impacted by the *impermanent loss*, otherwise referred to as divergence loss. Consider a liquidity pool $A \rightleftharpoons B$ between token A and B , where the amount of A in $A \rightleftharpoons B$ at time t is denoted as a_t and the amount of B in $A \rightleftharpoons B$ at time t is denoted as b_t . The fees collected between time t_1 and t_2 as a percentage of the liquidity are given as,

$$\text{fees}_{t_1 \rightarrow t_2} = 100 \cdot \left(1 - \frac{\sqrt{k_{t_1}}}{\sqrt{k_{t_2}}} \right),$$

where $k_t = a_t \cdot b_t$ [11]. The impermanent loss, on the other hand, describes the risk for liquidity providers of seeing the value of their reserved tokens decrease in comparison to holding the assets. This occurs with any price change in the pool. More precisely, the impermanent loss between t_1 and t_2 is given as,

$$\text{impermanent loss}_{t_1 \rightarrow t_2} = 100 \cdot \left(\frac{2 \cdot \sqrt{\frac{p_{t_2}}{p_{t_1}}}}{1 + \frac{p_{t_2}}{p_{t_1}}} - 1 \right),$$

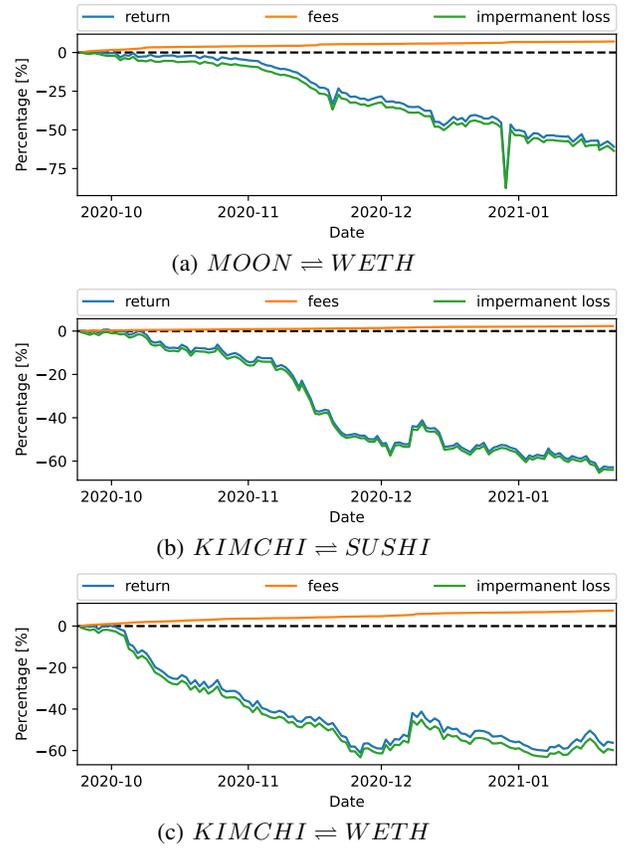
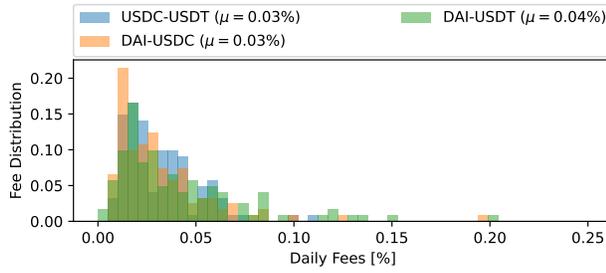


Figure 9: Evolution of cumulative returns, fees and impermanent loss over four months for exotic pairs. The fees should be much higher considering the risk of impermanent loss.

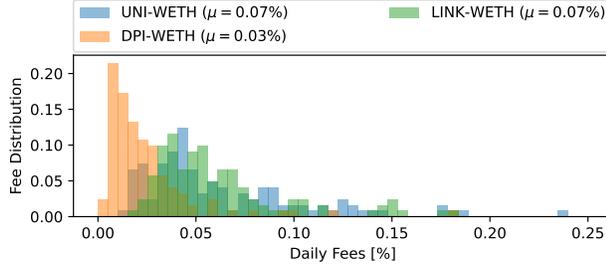
where $p_t = \frac{b_t}{a_t}$, is the ratio between tokens in the pool at time t [12]. Due to the impermanent loss, reserving of tokens runs the risk of under-performing a buy and hold strategy of a constant-mix portfolio.

The returns and the impermanent loss may vary greatly across the different liquidity pool categories, namely normal pools, stable pools, and exotic pools. Thus, to provide a comprehensive understanding of returns and risks of providing liquidity in Uniswap, we analyze them separately. We look at nine pools in detail - three of each kind (stable pairs: USDC \rightleftharpoons USDT, DAI \rightleftharpoons USDT, and DAI \rightleftharpoons USDC; normal pairs: UNI \rightleftharpoons WETH, LINK \rightleftharpoons WETH, DPI \rightleftharpoons WETH; exotic pairs: MOON \rightleftharpoons WETH, KIMCHI \rightleftharpoons SUSHI, KIMCHI \rightleftharpoons WETH)³. We choose pools through a combination of size and variety. When analyzing the data, we observe the same four-month period between the end of September 2020 and the end of January 2021, i.e., the time during which all nine sample pools were active. Further, when looking at daily returns, we consider the average daily return, computed from the average reserves ratio in the pool, as opposed to the closing daily return to mitigate the effects of short-term in-balances

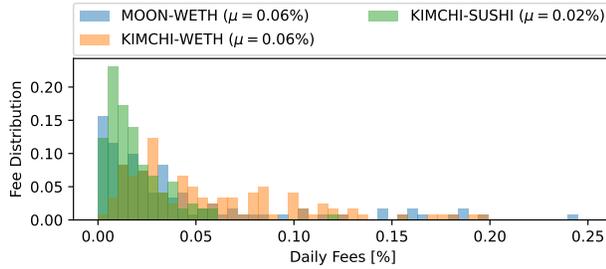
³In picking pools, we excluded liquidity mining pools (i.e., WETH \rightleftharpoons USDT, WETH \rightleftharpoons USDC, WETH \rightleftharpoons DAI, and WETH \rightleftharpoons WBTC), as the influence of the liquidity mining program clearly presents itself in the data.



(a) stable pairs



(b) normal pairs

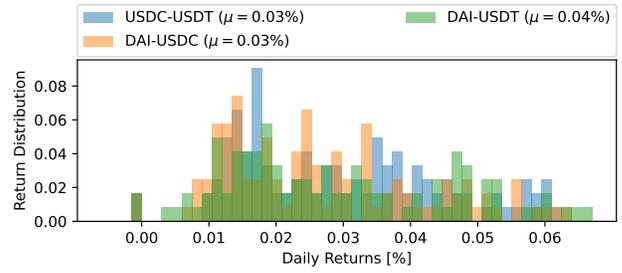


(c) exotic pairs

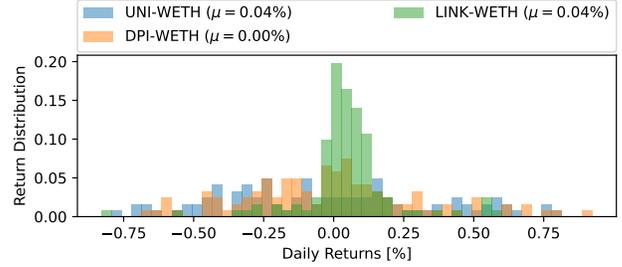
Figure 10: Daily percentage fees of sample pairs in three categories over a four month period. The histograms are normalized, i.e., the sum of the bar heights of each data set are equal to one. μ are the average daily fees in percent.

in pool reserves. Such imbalances occasionally occur after trades that are large in comparison to the pool reserves. Due to the influence of the impermanent loss on the returns, the inaccurate price ratio causes starkly negative returns. However, these temporary imbalances recover quickly as other trades see an arbitrage opportunity due to the inaccurate price ratio. We further note that percentage returns we extract, correspond to the realized returns of any liquidity provider in the respective pool during the analyzed period.

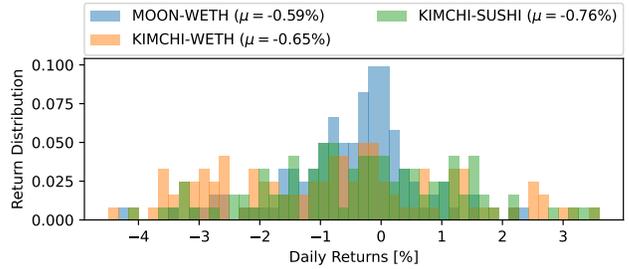
In Figures 7, 8 and 9, we show the evolution of the return, fees and impermanent loss of stable pools, normal pools, and exotic pools over time, respectively. Each of the three stable coins (USDT, USDC, and DAI) tracks the US\$. Thus, the ratio between each pair is close to one at all times, and impermanent loss plays a subordinate role in determining the returns and is negligible. Rather, the fees received through trades in the pool dominate the return. This dependency is apparent from the almost overlaying curves for return and fees shown across all three pairs shown in Figure 7. Moreover, due to negligible price fluctuations and continuous collection of fees, liquidity providers in stable pools can expect positive returns



(a) stable pairs



(b) normal pairs



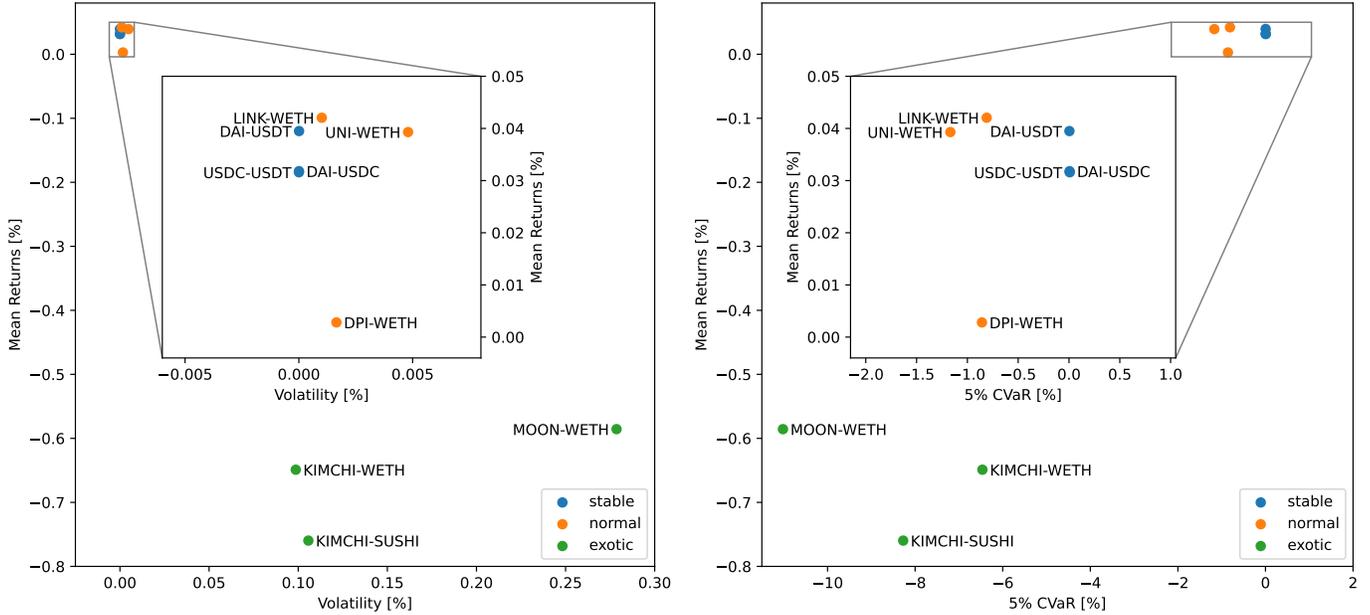
(c) exotic pairs

Figure 11: Daily percentage returns of sample pairs in three categories over a four month period. The histograms are normalized, i.e., the sum of the bar heights of each data set are equal to one. μ are the average daily returns in percent.

independent of the time of the liquidity injection. Therefore, providing liquidity in stable pools appears almost risk-free with consistent and stable revenue. Despite the comparatively high fees, users appear to exchange stable coins over Uniswap.

The influence of the impermanent loss on the profits of liquidity providers becomes more apparent for the two other types of pools – normal and exotic. For normal pairs (Figure 8), the cumulative return fluctuates below and above zero, influenced both by the ever-changing impermanent loss and the steadily increasing fees collected. Finally, we observe even starker domination of the impermanent loss for the return rate in Figure 9. Due to the high price volatility characteristic for exotic pairs, we observe impermanent losses of around 70% over a four-month period, which the collected fees cannot compensate – leading to deeply red returns.

Figure 10 illustrates the fee distribution observed in the nine sample pools. While we can detect differences in the distributions of the fees collected in the various pools, there is no apparent pattern between the different types of pools. The fees collected by liquidity providers appear to depend more



(a) Average daily returns and their volatility for nine sample pools. (b) Average daily returns and their 5% CVaR for nine sample pools.

Figure 12: Risk and return analysis of nine sample pools. Providing liquidity for exotic tokens is beyond just risky.

on the specific pair than on whether the pair is stable, normal, or exotic.

Turning to the return rate observed in our sample pools, Figure 11, we clearly observe patterns between the three different types of pools. Due to the previously observed negligible influence of impermanent loss on the profits expected in stable pairs, the daily returns in stable pools are rarely, if ever, negative – Figure 11a. Thus, all have positive daily average returns of around 0.03%. For the normal pairs, Figure 11b, we observe significantly higher volatility in the returns, accompanied by both higher (0.04% for $LINK \rightleftharpoons WETH$) and lower (0.00% for $DPI \rightleftharpoons WETH$) daily returns than previously seen for the stable pairs. Finally, the exotic pairs seen in Figure 11c are characterized by even larger return volatility and very negative daily returns (-0.76% for $KIMCHI \rightleftharpoons SUSHI$).

To summarize the risks and returns associated with providing liquidity, we turn to Figure 12. We use historical data to compare the daily mean returns to both the volatility and conditional value at risk (CVaR) of the returns.

Figure 12a plots the daily mean returns observed in each sample pool against the volatility. Generally, more volatile returns present greater risks to investors. The three exotic pairs both have the highest volatility and the lowest average return associated with them – making them an unattractive investment. Stable and normal pairs exhibit similar behavior, except for $DPI \rightleftharpoons WETH$, which presents significantly smaller returns than the other five pools.

We further analyze the risks and returns of the nine pools, and we plot the daily mean returns observed in each sample pool against the daily CVaR in Figure 12b. CVaR (sometimes referred to as expected shortfall) is a coherent risk measure that represents the average worst-case scenario commonly used

to quantify the risks of an investment. More precisely, CVaR at 5% level is the expected return on an investment in the worst 5% of cases [13]. CVAR has succeeded volatility and value at risk (VaR) in becoming the most frequently used risk measure [14]. While volatility is also affected by higher than average returns, VaR is blind to the size of the losses. Analyzing Figure 12b allows us to extract further patterns emerging between the three different kinds of pools that are in line with our previous observations. The three exotic pairs have the highest risk, measured through the CVaR, and the lowest average return associated with them. Exotic pairs make an inefficient investment to liquidity providers, as all three exotic pairs are dominated, i.e., higher return and lower risk, by each of the remaining pairs. The difference between stable and normal pairs becomes more apparent when comparing return and CVaR. While the $DPI \rightleftharpoons WETH$ pair is dominated by all stable pairs, the remaining two normal pairs are not. $LINK \rightleftharpoons WETH$ and $UNI \rightleftharpoons WETH$ exhibit higher returns than at least the majority of stable pairs, but with that also carry higher risks. All three stable pairs are very similar in terms of risk and return, with $DAI \rightleftharpoons USDT$ having a slight edge over the others in return. Stable and normal pairs may provide attractive opportunities for liquidity providers depending on their individual risk tolerance and return expectations, their difference becoming more apparent when analyzing tail behavior.

VI. MOVEMENT BETWEEN LIQUIDITY POOLS

As profits and risks associated with providing liquidity vary widely across different pools, we analyze how users redistribute their investments across different pools, i.e., how they move their assets from one pool to another. We investigate

	USDC \rightleftharpoons USDT	DAI \rightleftharpoons USDT	DAI \rightleftharpoons USDC
Volume	0.82	0.80	0.82
Price token 0	0.09	0.03	0.09
Price token 1	0.09	0.07	0.13

(a) stable pairs

	UNI \rightleftharpoons WETH	DPI \rightleftharpoons WETH	LINK \rightleftharpoons WETH
Volume	0.00	0.14	0.48
Price token 0	-0.31	-0.29	0.71
Price token 1	-0.36	-0.16	0.67

(b) normal pairs

	MOON \rightleftharpoons WETH	KIMCHI \rightleftharpoons WETH	KIMCHI \rightleftharpoons SUSHI
Volume	0.01	-0.20	-0.05
Price token 0	0.31	-0.50	0.08
Price token 1	-0.11	0.51	-0.05

(c) exotic pairs

Table I: Correlation between the liquidity in the pool and various pool characteristics. The correlations are recorded for daily data.

the movement of liquidity providers with the goal of better understanding their motivations.

In Table I we record the correlation between the pool liquidity and other pool characteristics, namely volume, and token prices in US\$. The correlations are calculated for daily data. For all stable pairs, Table Ia, volume correlates highly with pool liquidity. This strong correlation is expected, as liquidity provider’s returns in stable pools are directly related to the fees accumulated (Figure 7), which are in turn proportional to the volume. The high correlation between liquidity and trading volume indicates the dynamic balance of the liquidity providers in the market. When the trading volume increases, liquidity providers will earn more fees, then the high return attracts more liquidity into the pools. Symmetrically, when the trading volume decreases, liquidity will withdraw their funds from the pools, allowing the remaining liquidity providers to earn sufficient benefits again. Thus, the correlation between liquidity and trading volume appears natural. We see that liquidity tracks volume in stable pools in Figure 13a, where the daily volume and liquidity are visualized for $USDC \rightleftharpoons USDT$.

The data shown in Table Ib for normal pairs and Table Ic for exotic pairs reveals a less obvious picture. The liquidity in all but one of the normal and exotic pairs appears largely uncorrelated with the volume – $LINK \rightleftharpoons WETH$ being the exception. Figure 13b shows that while liquidity and volume are somewhat correlated for the $LINK \rightleftharpoons WETH$ pair, the link between them is not as apparent as for the stable pairs. In the remaining pools, an example is shown with $KIMCHI \rightleftharpoons SUSHI$ in Figure 13c, liquidity and volume are uncorrelated. The lack of correlation might partially be due to generally low volume and low liquidity but could also stem from the less predictable returns in normal and exotic pairs. Thus, liquidity providers might pay less attention to the volume when adding or removing liquidity. In general, the price of the tokens appears uncorrelated with the liquidity.

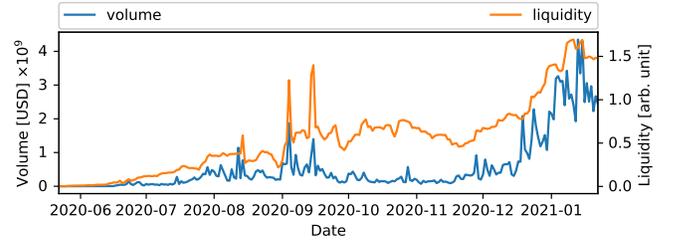
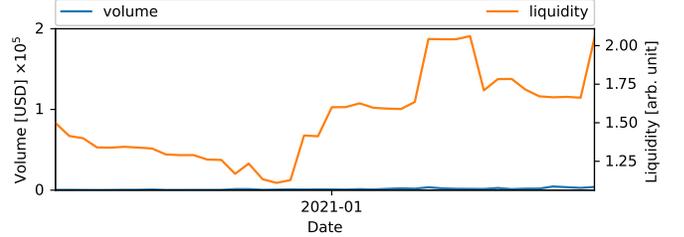
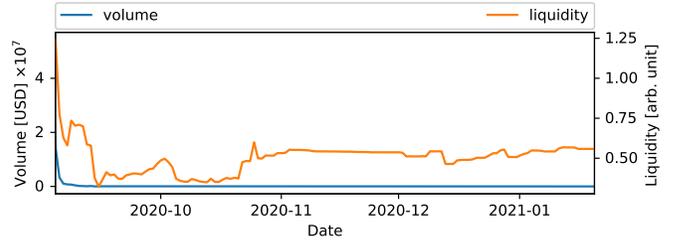
(a) $USDC \rightleftharpoons USDT$ (b) $LINK \rightleftharpoons WETH$ (c) $KIMCHI \rightleftharpoons SUSHI$

Figure 13: Volume and liquidity in three sample pools.

While we do not necessarily expect a strong correlation, the token price influences the liquidity providers’ returns via the impermanent loss. Thus, liquidity providers do not seem to react to price changes, indicating that either they hope for the ratio to recover to previous value or they are insensitive to the effects of the impermanent loss.

We further look at the movement of liquidity between pools. When observing liquidity movements, we consider the entire data set. We record a movement if the same address removes liquidity from one pool and adds liquidity to another pool within 6000 blocks – roughly a day. Further, we restrict our analysis to the pools with the most liquidity movements (mint and burn events) in the following and only consider pools with more than 5000 liquidity movements.

In Figure 14 we plot a colormap of the movement between the 72 pools most active pools. We order the pools smallest to largest by their average size, i.e., liquidity in the pool, since their creation. For better visibility, all values are capped at 500. We immediately draw three conclusions from Figure 14. First, the movements of liquidity in Uniswap are rare. Additionally, the matrix appears symmetric. In other words, if providers move liquidity from one pool to another, there also seem to be providers that move liquidity in the opposite direction at similar numbers. Lastly, the number of movements increases, especially for the six largest pools. Otherwise, movements do not appear to correlate with the pool size. The ordered

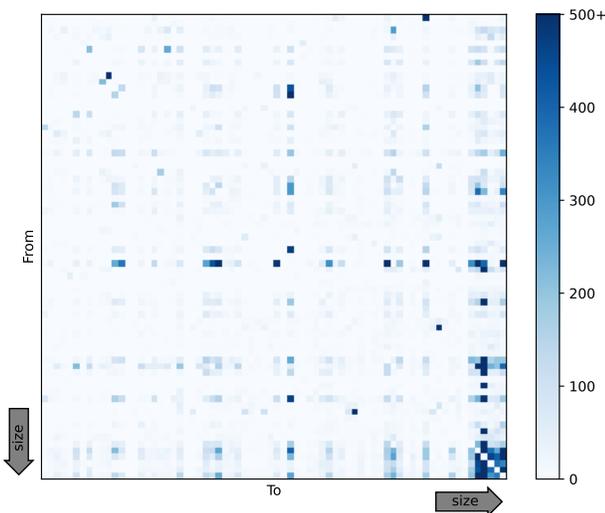


Figure 14: The number of directional movements between the 72 most active pools visualized as an adjacency matrix. The pools are ordered by their average size, i.e., liquidity in the pool, on 23 January 2021. The order list can be found in Appendix B. We cap the number of movements at 500 for better visibility.

pool pair with the most movements is $DAI \rightleftharpoons WETH \rightarrow WETH \rightleftharpoons SURF$ with a total of 13389 movements. We count less than 3500 movements for all remaining ordered pairs.

When looking at the data of some ordered pool pairs with many movements in more detail, we observe the emergence of two patterns in Figure 15. While movements between $DAI \rightleftharpoons WETH \rightarrow WETH \rightleftharpoons SURF$ all occur within a rather short period and appear to be driven by an individual event, the movements we observe between $USDT \rightleftharpoons WETH \rightarrow USDC \rightleftharpoons WETH$ and $USDC \rightleftharpoons WETH \rightarrow DAI \rightleftharpoons WETH$ happen over a longer period. However, even though we observe different patterns, movements appear to relate mostly to liquidity mining. The single spike in between $DAI \rightleftharpoons WETH \rightarrow WETH \rightleftharpoons SURF$ (Figure 15a) coincides with the time at which liquidity mining for $SURF$ started in the $WETH \rightleftharpoons SURF$ pair⁴. Further, the movement between the $USDT \rightleftharpoons WETH \rightarrow USDC \rightleftharpoons WETH$ and $USDC \rightleftharpoons WETH \rightarrow DAI \rightleftharpoons WETH$ reaches its peaks around the UNI liquidity mining period, 18 September 2020 to 17 November 2020⁵. In conclusion, the movement of liquidity does not appear too common among liquidity providers unless driven by external motivations, such as liquidity mining.

VII. CONCLUSION

This paper addresses three fundamental problems in understanding liquidity providers in DEXes: Who provides liquidity in DEXes? What are the returns and risks of providing liquidity? How do liquidity providers react to market changes?

⁴https://www.reddit.com/r/CryptoMoonShots/comments/jjivadf/surf_finance/

⁵<https://uniswap.org/blog/uni/>

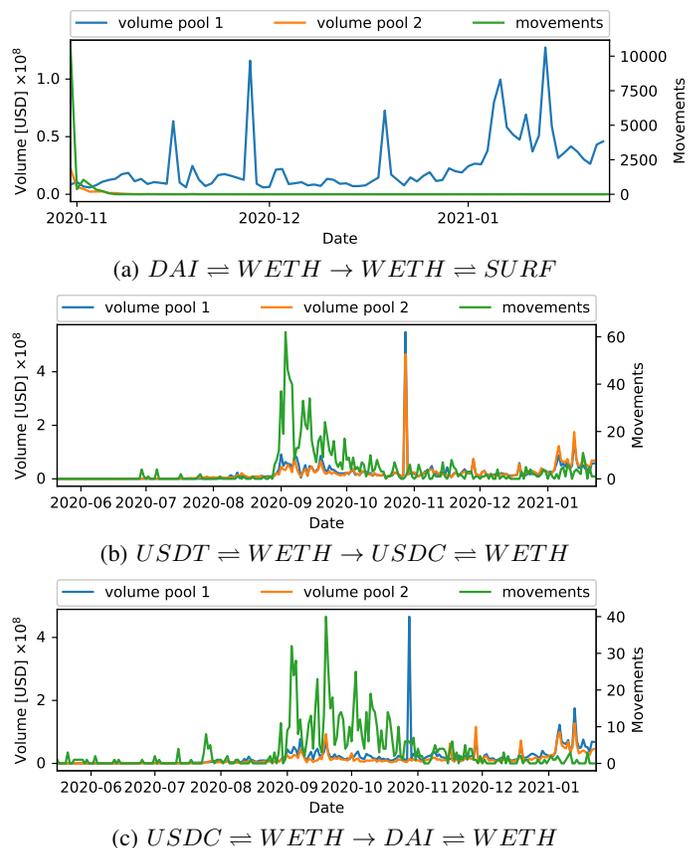


Figure 15: Directional liquidity movements and volume for a set of ordered pool pairs.

Our analysis suggests that users of cryptocurrency ecosystems have gradually become interested in providing liquidity in DEXes. However, users still act with caution – preferring to participate in few liquidity pools. We find that the returns and losses of providing liquidity in different types of pools vary a lot. Stable pools enable a seemingly risk-free and profitable investment opportunity compared to constant-mix portfolios. Providing liquidity in exotic pools, on the other hand, appears to perform much worse than the corresponding constant-mix portfolios. Reacting to the unique return opportunities and risks, liquidity providers perform different trading strategies across pool categories: they respond to trading volume changes in stable pools and pay less attention to market changes of normal and exotic pools. Besides market indicators, liquidity providers are also motivated by external market factors, i.e., liquidity mining activities, to redistribute their liquidity investments.

In this paper, we extend the research scope of liquidity provider behavior in DEXes from theoretical analysis to empirical studies. By studying the behaviors of liquidity providers on the Uniswap market, our work provides a comprehensive insight into the new trading options in the cryptocurrency ecosystems for users. More, it inspires future work aimed to better understand DEXes market mechanisms.

REFERENCES

- [1] W. Chen, T. Zhang, Z. Chen, Z. Zheng, and Y. Lu, "Traveling the token world: A graph analysis of ethereum ERC20 token ecosystem," in *Proceedings of The Web Conference 2020*, 2020, pp. 1411–1421.
- [2] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 910–927.
- [3] C. F. Torres, R. Camino, and R. State, "Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain," *arXiv preprint arXiv:2102.03347*, 2021.
- [4] K. Qin, L. Zhou, and A. Gervais, "Quantifying blockchain extractable value: How dark is the forest?" *arXiv preprint arXiv:2101.05511*, 2021.
- [5] L. Zhou, K. Qin, C. F. Torres, D. V. Le, and A. Gervais, "High-frequency trading on decentralized on-chain exchanges," *arXiv preprint arXiv:2009.14021*, 2020.
- [6] Y. Wang, Y. Chen, S. Deng, and R. Wattenhofer, "Cyclic arbitrage in decentralized exchange markets," *Available at SSRN 3834535*, 2021.
- [7] A. Evans, "Liquidity provider returns in geometric mean markets," *arXiv preprint arXiv:2006.08806*, 2020.
- [8] A. Evans, G. Angeris, and T. Chitra, "Optimal fees for geometric mean market makers," *arXiv preprint arXiv:2104.00446*, 2021.
- [9] J. Aoyagi and Y. Ito, "Liquidity implication of constant product market makers," *Available at SSRN 3808755*, 2021.
- [10] F. Victor, "Address clustering heuristics for ethereum," in *International Conference on Financial Cryptography and Data Security*. Springer, 2020, pp. 617–633.
- [11] H. Adams, N. Zinsmeister, and D. Robinson, "Uniswap v2 core," 2020.
- [12] Pintail, "Uniswap: A good deal for liquidity providers?" 2021. [Online; accessed 27-May-2021]. [Online]. Available: <https://pintail.medium.com/uniswap-a-good-deal-for-liquidity-providers-104c0b6816f2>
- [13] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [14] "Fundamental review of the trading book: A revised market risk framework," Bank for International Settlements, 2013.

APPENDIX A

DEX MECHANISM AND DATA EXTRACTION

In this section, we introduce the DEXes mechanism on the Ethereum blockchain and the details of the data extraction.

A. Ethereum Blockchain and ERC20 Standard

Most DEXes run on Ethereum: a public blockchain platform, which supports Turing complete smart contracts. Compared to earlier blockchain systems, such as Bitcoin, Ethereum provides a decentralized virtual machine, the Ethereum Virtual Machine (EVM), to execute smart contract code. A smart contract is a set of programs written in high-level languages, e.g., Solidity. These programs will be compiled into executable byte-code. After the creation of the smart contract, the executable byte-code will be stored in an independent database of the blockchain. Any Ethereum users can invoke functions defined in a smart contract.

Ethereum supports three kinds of transactions: a *simple transaction*, where the recipient is another address (account) to transfer the native currency, ETH; a *smart contract creation transaction*, without recipient is the *null* to create a new smart contract; a *smart contract execution transaction*, where the recipient is a smart contract address to execute a specific function of that contract.

When a transaction is included in a block by a miner, the operation corresponding to the message takes effect. The miner who creates the block modifies the state of corresponding accounts based on the messages. Each step of the miner's

operation consumes a certain amount of gas, and the amount of gas consumed in each block is capped. Users need to specify a gas price for the operation execution when sending transactions. The fee paid by the initiator of a transaction to miners is determined by the amount of gas consumed and the gas price (gas fee = gas price \times gas consumption). The miner will include a receipt of the executed transactions in Ethereum blocks as well, including the information on whether the transaction has been executed successfully, the gas fee, the identity of the transaction and the block, and other information generated during the execution.

Based on the support of the smart contract, users can create cryptocurrency other than ETH on Ethereum. These smart contracts have to follow some standards; the most widely used standard is the ERC20 standard, which requires an `approve` function and a `transferFrom` function. When an address (account) $addr_a$ calls the function `approve(addr_b, v)`, then the address $addr_b$ can transfer at most v tokens in total from $addr_a$ to other accounts. After this approval, $addr_b$ can transfer v' tokens from $addr_a$ to another account $addr_c$ by calling the function `transferFrom(addr_a, addr_c, v')`, where $\sum v' \leq v$.

B. DEXes

DEXes are smart contracts on Ethereum. Users send messages to a DEX address to invoke functions for performing market operations. DEXes support these operations: create the trading (liquidity) pool between a pair of tokens, adding/removing liquidity, and exchanging tokens. We take Uniswap as an example to present these operations.

1) *Creating Liquidity Pools*: In Uniswap, exchanges between two tokens are conducted through a liquidity pool, i.e., a smart contract that keeps the pair of tokens. There are two participants involved in the market: the liquidity provider and the trader. Providers reserve their tokens in the liquidity pool, while traders exchange their tokens with the liquidity pool. Because providers contribute to the market liquidity, they benefit from the transaction fees incurred with transactions in DEXes.

Assume we have two tokens A and B , and we want to create a liquidity pool between A and B on Uniswap. We first send a smart contract execution transaction with the ERC20 smart contract address of A and B to the Uniswap to claim the creation of the liquidity pool $A \rightleftharpoons B$. The smart contract will then check whether the pool between A and B exists according to the addresses of two tokens. If not, Uniswap smart contract will create a new liquidity pool $A \rightleftharpoons B$, i.e., a new smart contract reserving these two tokens.

2) *Adding/Removing Liquidity*: After creating the liquidity pool between A and B , liquidity providers can add a token pair to the liquidity pool. Liquidity providers need to approve the Uniswap address to transfer their A token and B token from their address to the liquidity pool address. When a Uniswap contract receives a liquidity providers call to add liquidity, it will invoke the `transferFrom` function in ERC20 contracts

to transfer tokens from the provider’s address to the liquidity pool address.

If there are no tokens reserved in the liquidity pool, users can supply any amount of A and B to the liquidity pool, and the pool will return *liquidity* tokens as proof of the deposit. If the amounts of A and B provided by the providers are a and b , respectively, then the provider will get $\lambda = \sqrt{a \times b}$ *liquidity* tokens. Meanwhile, the total supply of *liquidity* tokens of $A \rightleftharpoons B$ pool is $\Lambda = \sqrt{a \times b}$.

With a of token A and b of token B already in the liquidity pool, a provider can reserve δ_a of its asset A and δ_b of its asset B in the liquidity pool simultaneously, where $\frac{\delta_a}{\delta_b} = \frac{a}{b}$. Then they will earn $\delta_\lambda = \frac{\delta_a}{a}$ *liquidity* tokens for the $A \rightleftharpoons B$ and the total supply of *liquidity* tokens becomes $\Lambda = \lambda + \delta_\lambda$.

Providers can also remove their tokens from the liquidity pool. The amount of tokens providers can redeem is related to the amount of *liquidity* tokens they own. Assume a provider has δ_λ *liquidity* tokens of the liquidity pool $A \rightleftharpoons B$ and the total supply of *liquidity* tokens is λ . The provider can withdraw δ_a of A and δ_b of B from the $A \rightleftharpoons B$ pool with $\delta'_\lambda \leq \delta_\lambda$ *liquidity* token, where $\frac{\delta_a}{a} = \frac{\delta_b}{b} = \frac{\delta'_\lambda}{\lambda}$. The δ'_λ of *liquidity* tokens will be burned (destroyed) after they redeem the money and the total supply of *liquidity* tokens becomes $\Lambda = \lambda - \delta'_\lambda$.

3) *Exchanging Assets*: In Uniswap, tokens are not exchanged between two traders but between a trader and the liquidity pool. The exchange of assets is realized in two steps. First, the traders sends their tokens to the liquidity pool. Second, the liquidity pool computes the exchange rate and returns the targeted token to the traders.

Assume traders wants to exchange δ_a of A for B token and the liquidity of A and B are a and b . They first need to let the Uniswap address get the approval for transferring his A token to other accounts. After receiving the exchange order from the traders, the Uniswap contract transfers δ_a of A to the liquidity pool address and returns δ_b of B back to the traders.

The following equation always holds during the exchange: $a \cdot b = (a + \delta_a \cdot r_1) \cdot (b - \frac{\delta_b}{r_2})$, where r_1 and r_2 denote the transaction fee ratio in asset A and B respectively. In Uniswap, $r_1 = 0.997$ and $r_2 = 1$, which indicates that the transaction fee is equal to $3\% \cdot \delta_a$. The remaining liquidity in the pool equals to $(a + \delta_a, b - \frac{r_1 \cdot r_2 \cdot b \cdot \delta_a}{a + r_1 \cdot \delta_a})$ and the amount of *liquidity* tokens does not change.

C. Data Extraction

The addresses of liquidity pool smart contacts are stored in the *UniswapV2Factory* contract. Our modified go-ethereum client queries this contract to get the contract addresses and token pairs of 29,235 available trading pools until 23 January 2021. Then, we find 21,830,282 transactions interacting with these liquidity pools as Uniswap trades and filter those containing *Mint* and *Burn* events, representing adding liquidity, and removing liquidity, respectively.

In each *Mint* event, *liquidity* tokens of the liquidity pool will be generated by the liquidity pool smart contract and then transferred from the 0 address to the address of liquidity providers. Similarly, in each *Burn* event, liquidity providers

transfer their *liquidity* tokens to the 0 address. Moreover, some traders may exchange *liquidity* tokens in other transactions, while liquidity pools will generate *Transfer* events to record such *liquidity* token movements. With the information of *Mint*, *Burn*, and *Transfer* events, we compute the balance of *liquidity* tokens of liquidity providers in each pool.

APPENDIX B ORDERED LIST OF PAIRS

Pair	Liquidity [ETH]
YAM-yDAI+yUSDC+yUSDT+yTUSD	1.428872
WETH-ZZZ	1.485706
USDC-LUA	1.751875
XSP-WETH	14.462280
WETH-YFIG	17.151782
MOON-WETH	18.254465
EDC-WETH	23.734045
WETH-SASHIMI	23.880723
ORB-WETH	40.352911
MAGGOT-ROT	42.561145
WETH-ROT	44.211202
KIMCHI-WETH	60.737109
KIMCHI-SUSHI	62.282785
NAMI-WETH	64.820192
DAI-ZAI	71.726004
SAKE-WETH	116.067252
YAM-yyDAI+yUSDC+yUSDT+yTUSD	125.538101
sUSD-\$BASED	129.872536
DRC-WETH	258.590228
BNSD-WETH	262.181825
STA-WETH	369.155391
FARM-USDC	371.331104
CIBS-WETH	413.760214
DAI-ONC	674.910787
DOKI-WETH	738.239800
LEND-WETH	751.000403
BAND-WETH	761.538398
sUSD-WETH	774.386045
YFL-WETH	852.991518
WETH-SHROOM	1045.601330
CREAM-WETH	1149.861260
ESD-USDC	1176.263457
BASE-WETH	1200.543511
SAV3-WETH	1380.415120
USDC-DSD	1548.107930
BID-WETH	1671.099336
UMA-WETH	1799.774446
SAND-WETH	1863.114363
SUSHI-WETH	2287.350792
WETH-SURF	2587.357197
YFIM-WETH	2668.451593
XIO-WETH	2783.801616
KP3R-WETH	2825.997391
REN-WETH	3455.682119
COMP-WETH	3554.598169
MTA-WETH	3995.430011
WETH-CRV	4134.174818
MPH-WETH	4168.733445
DAI-BAS	4631.640805
STAKE-WETH	5209.986423
ORN-WETH	6125.185442
POLS-WETH	6741.259787
XOR-WETH	7473.624516
YFI-WETH	8580.269632

PICKLE-WETH	9027.224902
SNX-WETH	11902.224024
LON-WETH	15406.317841
MKR-WETH	16361.398631
BOND-USDC	18248.992687
WETH-AMPL	20063.122512
WBTC-BADGER	21483.810705
BAC-DAI	24813.984020
ibETH-ALPHA	26909.981701
USDC-USDT	29094.132246
DPI-WETH	37782.513858
CORE-WETH	38506.273571
LINK-WETH	71021.684491
WETH-USDT	127760.291154
DAI-WETH	130139.207443
UNI-WETH	141314.373502
WBTC-WETH	165777.023010
USDC-WETH	213969.777843

Table II: Order list of pairs for Figure 14. The pairs are ordered by their liquidity on 23 January 2021.