

An Empirical Study of Market Inefficiencies in Uniswap and SushiSwap

Jan Arvid Berg, Robin Fritsch*, Lioba Heimbach*, and Roger Wattenhofer

ETH Zürich

{bergar, rfritsch, hlioba, wattenhofer}@ethz.ch

Abstract. Decentralized exchanges are revolutionizing finance. With their ever-growing increase in popularity, a natural question that begs to be asked is: how efficient are these new markets?

We find that nearly 30% of analyzed trades are executed at an unfavorable rate. Additionally, we observe that, especially during the DeFi summer in 2020, price inaccuracies across the market plagued DEXes. Uniswap and SushiSwap, however, quickly adapt to their increased volumes. We see an increase in market efficiency with time during the observation period. Nonetheless, the DEXes still struggle to track the reference market when cryptocurrency prices are highly volatile. During such periods of high volatility, we observe the market becoming less efficient – manifested by an increased prevalence in cyclic arbitrage opportunities.

Keywords: blockchain · automated market maker · market efficiency.

1 Introduction

Nakamoto introduced the first fully decentralized cryptocurrency, Bitcoin [13], in 2008. In the following years, several blockchains followed, notably Ethereum [5] which introduced smart contracts. In their initial phase, blockchains only had a few niche applications, and the excitement surrounding them was mainly fueled by the hopes of continuously rising cryptocurrency prices.

However, this changed with the introduction of *decentralized finance (DeFi)*. Suddenly, blockchains had a new purpose: offering financial services without the need of a middleman. *Decentralized exchanges (DEXes)*, which allow users to trade in a fully noncustodial manner, are a main pillar of DeFi. Instead of requiring traders to give up custody of their funds by depositing into a centralized exchange (CEX), traders can now directly swap tokens with a smart contract on the blockchain. The popularity of DEXes is undeniable. The trading volume on all DEXes, which include Uniswap, SushiSwap, Balancer, Bancor, and Curve, exceeded \$50 billion in January 2021 and in every month since [9].

Most of these DEXes use a novel market-making mechanism. Rather than matching orders using a limit order book like traditional exchanges, most DEXes

* Corresponding authors.

use an automated market maker mechanism that executes orders against a *liquidity pool* holding token reserves. The exchange rate in a liquidity pool is determined by a trading function and the amount of funds in the pool. Temporarily, these rates can be inaccurate and vary across different DEXes. Such inaccuracies in the price lead to trades being executed at unfavorable rates if traders are not alert. Further, they can create cyclic arbitrage opportunities. Cyclic arbitrage opportunities indicate erroneous rates and thereby stem from lacking market efficiency, which measures how well the prices reflect all relevant information [12]. As the market is becoming increasingly complex with an ever-growing number of DEXes and liquidity pools, studying the market’s efficiency is ever more important.

We investigate the existence and severity of market inefficiencies in two market leading DEXes, Uniswap and SushiSwap, between 12 September 2020 and 23 January 2021. With the optimal routing problem, we find trades that executed with an unfavorable rate – indicating both the presence of price inaccuracies in the market and their effects on traders. Further, we look for past cyclic arbitrage opportunities, which stem from price differences in the market, and use them as a tool to determine how efficient the market as a whole is.

2 Related Work

Numerous studies analyze market efficiency on CEXes [15, 12, 4, 18]. Conversely, we measure the market efficiency on DEXes, by studying the prevalence of price inaccuracies across the market. More precisely, we focus on DEXes using an *automated market maker (AMM)*.

AMMs have been around for quite some time, e.g., the logarithmic market scoring rule used in prediction markets [10]. However, only the recent emergence of DEXes has made AMMs a popular alternative to traditional central limit order book systems. Uniswap and SushiSwap, along with most other DEXes, implement a new type of AMM design called *constant function market maker (CFMM)* [2]. More specifically, Uniswap and SushiSwap use a form of CFMM called *constant product market maker (CPMM)*. Angeris et al. [3] formally analyze how closely CPMMs track reference markets and demonstrate the numerical stability of CPMMs under a wide range of market conditions. We, on the other hand, empirically study the efficiency of Uniswap and SushiSwap in tracking the reference market. By looking for optimizable trades and cyclic arbitrage opportunities, both caused by price inaccuracies in the market, we identify moments when some prices on the DEXes do not accurately track the reference markets.

Several works quantify and study *Blockchain Extractable Value (BEV)* in DeFi. BEV measures the extractable profit from the blockchain by including, excluding, and re-ordering the transactions in a block. A broad study on both the amount of extractable and extracted BEV in DeFi is provided by Qin et al. [14]. Daian et al. [6] study front-running in decentralized exchanges and show that these arbitrage opportunities can cause a priority gas auction, which drives up the gas fee for all blockchain users. Through studying back-running on DEXes,

Zhou et al. [17] show that this kind of BEV can cause back-run flooding, a denial of service practice on the blockchain. As opposed to focusing on the effects of arbitrage opportunities on DEXes, we empirically study one of their root causes – price inaccuracies in the market.

Danos et al. [8] theoretically study the optimal routing of trades through a network of CFMMs. We apply the special case of independent paths studied by Danos et al. [7] to analyze past trading data. In examining big swaps, we find that a significant fraction executes suboptimally on Uniswap and SushiSwap. Traders are, thus, suffering from price inaccuracies across the market.

Wang et al. [16] focus on cyclic arbitrage trades, a type of BEV, analyzing them theoretically and empirically. In contrast to this work, we study the availability of cyclic arbitrage opportunities in this paper and use it to identify price inaccuracies in the market. We further look for the causes of the inaccuracies and use them as a tool to study the market’s efficiency over time.

3 Constant Product Market Makers

CPMMs, which include Uniswap and SushiSwap, are smart contracts running on the Ethereum blockchain. On a CPMM, anyone can create a liquidity pool for an arbitrary token pair. Once created, liquidity providers deposit amounts of equal value of both tokens into the pool. Traders swap tokens with the pool’s liquidity and pay a small trading fee, which is distributed pro-rata among all liquidity providers. To execute such a transaction on the blockchain, traders first submit their transaction to the Ethereum *mempool* and wait for a miner to include them in the next block.

The CPMM smart contract determines the output amount the trader receives. The amount returned by the CPMM ensures that the product of the two reserve amounts stays constant. More precisely, consider a pool that holds reserves of a token A and token B . Assume the reserves are R_A token A and R_B token B , and that a trader wants to swap t_A token A . The trader then receives

$$t_B = R_B - \frac{R_A R_B}{R_A + (1-f)t_A} = \frac{R_B(1-f)t_A}{R_A + (1-f)t_A},$$

where f is the transaction fee charged on the input amount (0.3% in Uniswap) [1]. Thus, when swapping an amount of t_A token A , the price per token B is $\frac{R_A + (1-f)t_A}{R_B(1-f)}$. We note that the price per token is increasing in the input amount. The larger the trade gets, the more the trader has to pay per desired token. This effect is a consequence of a trade’s *price impact* – the impact of an individual trade on the market price. We also see in the formula above, that a trade’s price impact is lower in pools with larger reserves. Thus, it is both the ratio of its assets and the liquidity of the pool that determine the price.

4 Data Description

We analyze data from Uniswap and SushiSwap: Ethereum’s two largest DEXes by trading volume. Together they represent the majority of the market, as they account for more than half of Ethereum’s DEX trading volume [9]. Further, both use the exact same trading mechanism, making them ideal for analyzing price inaccuracies between markets. More precisely, we collect pool reserve and transaction data between 5 August 2020 and 23 January 2021 for all pools between any of the following five cryptocurrencies: Bitcoin (BTC), Ethereum (ETH), Tether (USDT), USD Coin (USDC), and Dai (DAI). These are some of the most traded tokens and are among those which share the highest number of pools with other cryptocurrencies [11]. Thus, they are well suited for our analysis. The high number of pairwise pools provides a large set of independent paths between cryptocurrency pairs to find optimizable routes and cyclic arbitrage opportunities. For a thorough data description, see Appendix A.

5 Identifying Market Inefficiencies

To find market inefficiencies and study how they evolved over time, we analyze both the optimizability of past transactions and the opportunity for cyclic arbitrage. Price differences, the root cause of suboptimal trades, and cyclic arbitrage opportunities reflect the market’s ability to reflect the relevant information.

5.1 Suboptimal Trade Routing

We start with an example of a trade that was executed at an unfavorable rate to illustrate the meaning of suboptimal trade routing (Figure 1). On 15 November 2020, the original trade exchanged ETH for 16.61 BTC. The trade routed through the ETH-BTC SushiSwap pool, as we show on the left in Figure 1. We visualize an optimization that executes 99.75% of the trade in the Uniswap pool and only 0.25% of the trade in the SushiSwap pool. By routing the transaction as shown, the trader would have received 18.07 BTC – an 8.79% improvement. We notice that the Uniswap trade offered a better price for this trade. Thus, the market was unsynced at this moment, and thereby prices in at least one pool did not reflect all relevant information – indicative of market inefficiency. As a result of the large trade size, it was optimal to route a small part of the trade through the SushiSwap pool to reduce the trade’s price impact.




Fig. 1: Example of original and optimizable swap. The Uniswap pools are represented by solid lines and SushiSwap pools are represented by dashed lines.

We analyze past transactions over \$30’000 from 12 September 2020 to 23 January 2021 and check whether they were routed optimally. We focus on large

trades since they have a more significant price impact, and the benefit from an optimal routing is greater for these trades compared to smaller ones. These trades are, thus, particularly well suited to study in order to detect inefficiencies in the market. We also note that we exclude swaps that are part of a transaction including multiple swaps. While some of these swaps could be optimizable, others might already be part of optimized routing. Then optimizing one of them without considering the others would overestimate the possibility for optimization.

To identify trades that received unfavorable prices, we apply a special case of the optimal trade routing problem studied by Danos et al. [7]. They provide the solution to the optimal routing problem for trades through a set of independent paths. The solution accounts for the CPMM’s transactions fees. We construct the set of independent paths as follows: we include both direct routes (Uniswap and SushiSwap) if they exist. The direct route is the pool containing the input and output tokens. For paths containing multiple pairs, we consider the more liquid Uniswap pool (Figure 6a). We search for transactions whose output can be increased by more than \$30 with optimized routing and consider them optimizable. The threshold accounts for additional gas fees related to potentially routing the trade through several routes. The potential for such optimization indicates either the existence of price inaccuracies or that the market is illiquid.

29’611 out of 108’667 analyzed transactions (Figure 5) were optimizable – a share of 27%. On average, it is possible to increase the output of an optimizable trade by 0.15% (Table 1) by using an optimized routing. While it might not appear to be a significant proportion, the 0.15% is an invisible tax placed on trades stemming from price inaccuracies present in the market. We further observe that the mean gain is considerably higher than the median gain, suggesting that the tax on the most affected traders is significantly higher. Looking at the top 5% trades, where optimized routing would improve by price most significantly (in percentage terms), we find that the average achievable gain was 0.71% (Table 1). Further, indicating that traders are suffering from these market inefficiencies.

	All Trades	Top 5% Trades
Mean Gain	0.15%	0.71%

Table 1: Gains achieved by optimized routings.

To check whether the market inefficiencies stem from price inaccuracies or the potential lack of liquidity in the market, we analyze how many paths were used by the optimized routings (Figure 2). We count a path if at least 0.1% of the trade routes through it. 18% of optimizable trades only use a single path with the new routing. Thus, this optimized routing does not include the original path at all. The unfavorable price received, thereby, solely stems from price inaccuracies. We note that the optimized routing for a small proportion of trades consists of at least three paths. There are at most five possible paths in the network (Appendix B, Figure 6a). These relatively complex routings indicate that the liquidity of the pools also causes unfavorable rates received by large trades.

When we repeat the analysis and include the less liquid SushiSwap pools (Appendix B, Figure 6b) in the set of independent paths, we find that the op-



Fig. 2: Proportions of number of paths used by optimized routings.

timization potential increases. We perform this adjacent analysis on a smaller set of 43'321 swaps, which include all trades originally executed in the following pools: USDC-ETH (Uniswap and SushiSwap) and DAI-ETH (SushiSwap). Here, we find that an even larger proportion of swaps, 33.75% compared to 27% previously, could have been optimized. Further, the average gains are higher in the less liquid pools (Table 2). We find that routing through lower liquidity pools leads to better optimizing possibilities. Thus, it seems to be price inaccuracies across the market driving the unfavorable rates as opposed to the lack of liquidity in parts of the market – indicative of the presence of market inefficiencies in (parts of) the market. The significant number of trades with unfavorable rates shows that traders are experiencing the consequences of these inefficiencies.

	More Liquid Pools	Less Liquid Pools
Mean Gain	0.15%	0.16%

Table 2: Gains achieved by optimized routings percent for the two sets of pools.

5.2 Cyclic Arbitrage Opportunities

Cyclic arbitrage opportunities result from temporary price inaccuracies. By taking advantage of non-equilibrium exchange rates within and between markets, arbitrage traders can profit by trading their funds in a cycle. However, the existence of such opportunities suggests market inefficiency. The exchange rate of at least one pool in such a cycle must not accurately reflect all relevant information. Studying the prevalence and duration of these cycles brings inside new market insights beyond the unfavorable prices experienced by traders.

We identify cyclic arbitrage opportunities retrospectively by searching for directed cycles of pools where a cyclic swap with α tokens A , the returns $\hat{\alpha}$ tokens A where $\hat{\alpha} > \alpha$. In a cycle $c = (e_1, \dots, e_n)$ each edge e_i present in the cycle represents a pool used in the transactions. For each arbitrage opportunity, we compute the maximum possible profit. The problem’s convexity [8] lets us find the optimal, unique solution that maximizes the profit $\hat{\alpha} - \alpha$.

Over six months, from 5 August 2020 to 23 January 2021, we analyze the occurrence of cyclic arbitrage opportunities between ETH, USDC, and USDT on Uniswap. These three cryptocurrencies have the largest pairwise Uniswap pools between them. We plot the daily number of blocks with arbitrage opportunities exceeding 30\$ and ETH’s daily price movement in Figure 3. Daily price movement is a volatility measure and given by is $100\% \cdot \frac{p_{\text{high}} - p_{\text{low}}}{p_{\text{low}}}$. Here, p_{high} is the day’s highest price and p_{low} is the day’s lowest price.

Most arbitrage opportunities are between August and September 2020 (Figure 3). In this period, Uniswap experienced an incredible increase in daily trading volume. Within six weeks, the volume increased steadily by around 700%, reaching \$500 million at the end of September. Similar daily trading volumes were not consistently reached again until early January 2021. Arbitrage opportunities




Fig. 3: Daily number of blocks with cyclic arbitrage opportunities and ETH price movement between 5 August 2020 and 23 January 2021.

also experience a sharp decrease starting at the end of September, and we only begin to see a consistent uptake at the end of December. In general, the number of arbitrage opportunities appears to decrease throughout the analyzed period – indicating that the market is becoming more efficient. We observe a second trend when comparing the number of daily arbitrage opportunities to ETH’s price movement. With 0.15, the correlation between the two is relatively low. However, after the initial explosion in volume in the summer/autumn of 2020, we find that on days with exceptionally high price movements, we also observe exceptionally many arbitrage opportunities.

To further investigate the influence of external cryptocurrency prices on the number of arbitrage opportunities, we search for cycles in five pools on Uniswap and SushiSwap between ETH, USDC, USDT, DAI, and BTC during the following two periods: (1) 11 September to 3 October 2020, and (2) 23 December 2020 to 23 January 2021. The first period analyzed is characterized by relatively stable cryptocurrency prices – the price of ETH moved by less than 8%. The second period, however, is characterized by highly volatile cryptocurrency prices – the price of ETH more than doubled. By then, the pools containing BTC had accumulated enough liquidity, and we included them in the analysis. The analyzed networks during both periods are shown in Figures 7a and 7b (Appendix B).

	Blocks with Cyclic Arb.	Mean Profit	Avg. Duration
11.09.20 - 03.10.20	84	0.24%	2.33 blocks
23.12.20 - 23.01.21	1'061	0.35%	1.43 blocks

Table 3: Statistics of cyclic arbitrage profits in the two analyzed periods.

We find that 84 out of 140,000 blocks have arbitrage cycles with a mean profit of 0.24% in the first period (Table 3). Stable cryptocurrency prices appear to not allow for many arbitrage opportunities, even in the face of the recent introduction of SushiSwap at the time. The markets’ exchange rates are synced. In the second period, we find 1061 blocks with cyclic arbitrage opportunities, and, on average, they present a profit of 0.351% (Table 3).




Fig. 4: Daily number of blocks with cyclic arbitrage opportunities and ETH price movement from 23 December 2020 to 23 January 2021.

As suspected, the number of cyclic arbitrage opportunities is significantly higher in the second period, where cryptocurrencies prices are highly volatile. To show the correlation, we plot the distribution of arbitrage opportunities and ETH’s price movement in Figure 4. Two days, 4 January and 11 January 2021, make up for over 50% of all identified opportunities. The price of ETH experienced temporary movements exceeding 20% on both days. Further, the correlation between the number of arbitrage opportunities and the ETH price movement is 0.72. We conclude that, as expected, higher price volatility is more favorable for arbitrage. The floating exchange rates on Uniswap and SushiSwap do not adjust to market price sufficiently fast. For every cyclic arbitrage opportunity, there is at least one pool whose price does not reflect all relevant information. We also observe that arbitrage bots, who profit from the BEV caused by market inefficiencies, become more efficient in the months between the two periods. The arbitrage opportunities are available for 2.33 blocks on average in the first period and only available for 1.43 blocks on average in the second period (Table 3). Since the minimum availability of an arbitrage opportunity we observe is 1, this improvement in the efficiency of arbitrage bots is significant.

6 Conclusion

Traders are actively suffering from market inefficiencies on DEXes: nearly 30% of analyzed trades were executed at an unfavorable price. We show that these erroneous rates stem largely from price inaccuracies across the market’s liquidity pools. The market struggled to accurately reflect all relevant price information, especially during the initial volume explosion on DEXes in the late summer of 2020 – evident from the increased number of cyclic arbitrage opportunities. However, the market quickly adapts, and the price inaccuracies largely disappear. Later, we only observe market inefficiencies on days when cryptocurrency prices are highly volatile. Parts of the market do not adjust to external price changes properly, thereby creating cyclic arbitrage opportunities. Arbitrage bots are also becoming more efficient and take advantage of these prevailing market inefficiencies within a block at most times.

References

1. Adams, H., Zinsmeister, N., Robinson, D.: Uniswap v2 core (2020)
2. Angeris, G., Chitra, T.: Improved price oracles: Constant function market makers. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies. p. 80–91. AFT '20, Association for Computing Machinery, New York, NY, USA (2020)
3. Angeris, G., Kao, H.T., Chiang, R., Noyes, C., Chitra, T.: An analysis of uniswap markets. arXiv preprint arXiv:1911.03380 (2019)
4. Bernard, V., Thomas, J., Wahlen, J.: Accounting-based stock price anomalies: Separating market inefficiencies from risk. *Contemporary Accounting Research* **14**(2), 89–136 (1997)
5. Buterin, V.: Ethereum whitepaper (2013), <https://ethereum.org/en/whitepaper/>
6. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 910–927. IEEE (2020)
7. Danos, V., Khalloufi, H.E.: Optimal execution in uniswap like amms: Independent paths case. Working Draft (04 2021)
8. Danos, V., Khalloufi, H.E., Prat, J.: Global order routing on exchange networks. In: Bernhard, M., Bracciali, A., Gudgeon, L., Haines, T., Klages-Mundt, A., Matsuo, S., Perez, D., Sala, M., Werner, S. (eds.) *Financial Cryptography and Data Security. FC 2021 International Workshops*. pp. 207–226. Springer Berlin Heidelberg, Berlin, Heidelberg (2021)
9. DeFiprime: Dex tracker - decentralized exchanges trading volume, <https://defiprime.com/dex-volume>
10. Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* **5**(1), 107–119 (1 2003). <https://doi.org/10.1023/A:1022058209073>
11. Heimbach, L., Wang, Y., Wattenhofer, R.: Behavior of Liquidity Providers in Decentralized Exchanges. In: *Crypto Valley Conference on Blockchain Technology (CVCBT 2021)*, Rotkreuz, Switzerland (October 2021)
12. Malkiel, B.G.: Is the stock market efficient? *Science* **243**(4896), 1313–1318 (1989)
13. Nakamoto, S.: Bitcoin whitepaper (2008), <https://bitcoin.org/bitcoin.pdf>
14. Qin, K., Zhou, L., Gervais, A.: Quantifying blockchain extractable value: How dark is the forest? arXiv preprint arXiv:2101.05511 (2021)
15. Stiglitz, J.E.: The inefficiency of the stock market equilibrium. *The Review of Economic Studies* **49**(2), 241–261 (1982)
16. Wang, Y., Chen, Y., Deng, S., Wattenhofer, R.: Cyclic arbitrage in decentralized exchange markets (2021)
17. Zhou, L., Qin, K., Gervais, A.: A2mm: Mitigating frontrunning, transaction re-ordering and consensus instability in decentralized exchanges. arXiv preprint arXiv:2106.07371 (2021)
18. Zunino, L., Tabak, B.M., Figliola, A., Pérez, D.G., Garavaglia, M., Rosso, O.A.: A multifractal approach for stock market inefficiency. *Physica A: Statistical Mechanics and its Applications* **387**(26), 6558–6566 (2008)

A Data Description

In this section we provide a detailed data description.

A.1 Reserve Data

By launching a go-ethereum client, we collect all pool reserves recorded on Ethereum from block 10000835 (4 May 2020, day of Uniswap V2 deployment) to block 11709847 (23 January 2021) for Uniswap. For SushiSwap, the client exports all pool reserves from block 10750000 (28 August 2020, day of SushiSwap deployment) to block 11709847 (23 January 2021).

A.2 Transaction Data

We collect the transaction data required for the optimal trading analysis from The Graph’s Uniswap V2 subgraph¹ and The Graph’s SushiSwap subgraph² with GraphQL. The Graph is a decentralized protocol for indexing and querying data from blockchains like Ethereum.

For each transaction between block 10000835 and block 11709847, we collect the amounts of tokens swapped, as well as the value of the swap in \$. While the subgraphs we use also provide pool reserves, these are updated infrequently and therefore could not be used. For this reason, we collect pool reserves with the go-ethereum client.

We pre-process the transaction data to reflect the reserves at the beginning of each block. These reserves coincide with those at the end of the previous block and are the only information most traders use when submitting their transaction (assuming their transaction is included in the next block). Depending on other transactions executed ahead of the swap in the same block, the pool reserves might shift slightly before trade execution and thus change the exchange rate. We correct this effect by pre-processing the data accordingly such that the transactions reflect the trader’s view – based on which we optimize in Section 5.1.

A.3 Pool Network

Figure 5 visualizes the number of independent swaps exceeding \$30’000 in both Uniswap and SushiSwap between 12 September 2020 and 23 January 2021. In Uniswap we have a pool between almost every pair of tokens, while for SushiSwap the liquid pools during the observed time period all included ETH. Nonetheless, even on Uniswap most large trades are executed in pools that include ETH.

¹ <https://thegraph.com/explorer/subgraph/uniswap/uniswap-v2>

² <https://thegraph.com/explorer/subgraph/benesjan/sushi-swap>




Fig. 5: The number of independent trades exceeding \$30'000 in each existing pool between any pair of the following cryptocurrencies: BTC, ETH, USDC, USDT and DAI, in Uniswap (Figure 5a) and SushiSwap (Figure 5b). Data was collected between 12 September 2020 and 23 January 2021.

B Networks

We show the networks used to find optimized routing in Figure 6 and the networks used to find cyclic arbitrage opportunities in Figure 7.




Fig. 6: Paths from USDC to ETH for both cases. Uniswap pools are represented by solid lines and SushiSwap pools are represented by dashed lines.




Fig. 7: Network of pools considered during the two analyzed time periods. Cryptocurrencies are nodes and edges represent pools between the respective cryptocurrencies. Uniswap pools are represented by solid lines and SushiSwap pools are represented by dashed lines.