

# Exploring Price Accuracy on Uniswap V3 in Times of Distress

Lioba Heimbach  
ETH Zürich  
Switzerland  
hlioba@ethz.ch

Eric Schertenleib  
ETH Zürich  
Switzerland  
ericsch@ethz.ch

Roger Wattenhofer  
ETH Zürich  
Switzerland  
wattenhofer@ethz.ch

## ABSTRACT

Financial markets have evolved over centuries, and exchanges have converged to rely on the order book mechanism for market making. Latency on the blockchain, however, has prevented *decentralized exchanges (DEXes)* from utilizing the order book mechanism and instead gave rise to the development of market designs that are better suited to a blockchain. Although the first widely popularized DEX, Uniswap V2, stood out through its astonishing simplicity, a recent design overhaul introduced with Uniswap V3 has introduced increasing levels of complexity aiming to increase capital efficiency.

In this work, we empirically study the ability of Uniswap V3 to handle unexpected price shocks. Our analysis finds that the prices on Uniswap V3 were inaccurate during the recent abrupt price drops of two stablecoins: UST and USDT. We identify the lack of agility required of Uniswap V3 liquidity providers as the root cause of these worrying price inaccuracies. Additionally, we outline that there are too few incentives for liquidity providers to enter liquidity pools, given the elevated volatility in such market conditions.

## KEYWORDS

blockchain, decentralized finance, decentralized exchange, constant product market maker, liquidity providers

### ACM Reference Format:

Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. 2022. Exploring Price Accuracy on Uniswap V3 in Times of Distress. In *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security (DeFi '22)*, November 11, 2022, Los Angeles, CA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3560832.3563435>

## 1 INTRODUCTION

While the first stock markets date back at least to the 17th century and have thus a long track record, their novel *decentralized finance (DeFi)* counterpart: *decentralized exchanges (DEXes)* are still in their teething phase. Most financial markets rely on an order book mechanism, where constant buy and sell orders determine the price in an auction-like fashion. The price lies between the highest *bid* (buy order) and lowest *ask* (sell order) and can adapt rapidly.

DEXes are built by smart contracts hosted on the blockchain, mainly Ethereum [26], and the significant latency between subsequent blocks has generally prohibited the adoption of the order

book mechanism on the blockchain. Instead DEXes, generally function as *constant function market makers (CFMMs)*. CFMMs rely on a predetermined price curve for on-chain market making. Individual liquidity providers aggregate liquidity in what is referred to as a liquidity pool for each tradeable cryptocurrency pair. Liquidity providers are remunerated for their service through transaction fees but bear the risk of high volatility.

The first widely popular CFMM was Uniswap V2 [9]. Uniswap V2 functions as a *constant product market maker (CPMM)*, where liquidity providers supply liquidity on the entire price range. The CPMM ensures that the product of the pool's two reserved cryptocurrencies remains constant. Thus, executed trades directly cause price shifts in the market. Large price differences between markets give rise to arbitrage opportunities which arbitrage traders are expected to exploit until the arbitrage opportunity vanishes.

Curve [1] later proposed a CFMM specialized in stablecoin (cryptocurrencies designed to have stable prices) liquidity pools. Stablecoins are often pegged to the US\$ and should therefore trade at an equal price. The CFMM proposed by Curve adapts the CPMM to reduce price changes for a given trade size as long as stablecoins are trading near their equal price. In a recent design overhaul, Uniswap implemented an updated protocol V3 [10], where liquidity providers have great freedom in selecting the price range for which they wish to facilitate trading. This update was marketed, in part, to be specially tailored for stablecoin pairs, as the more concentrated liquidity distribution at the intended price should increase capital efficiency.

For stablecoin pairs on Uniswap V3, the liquidity is highly concentrated around the equal price point of the two assets in the pool, i.e., both assets trade at 1 US\$. Thus, as soon as the price moves outside an extremely narrow price range and the liquidity is not adjusted accordingly, there is little to no liquidity left to support trading. The concentrated liquidity CFMM introduced by Uniswap V3 was tested in early May 2022. First, UST, a stablecoin with a market capitalization of 18.7B US\$ [6], lost its peg of 1 US\$ on 7 May and plummeted to a couple of cents within days. USDT, the largest stablecoin with a market capitalization of 83.2B US\$, also lost its peg of 1 US\$ and temporarily traded at 94 cents on 11 May. While USDT largely recovered from its price drop, UST did not.

In this work, we empirically analyze the ability of the concentrated liquidity CFMM utilized by Uniswap V3 to handle such unexpected price shocks and find that Uniswap V3 markets were not able to adjust to the dramatic price drops of UST and USDT in early May 2022. We further conclude that the lack of sophistication and agility of Uniswap V3 liquidity providers is largely responsible for the market's price inaccuracies. Additionally, we find that once the supposedly stable price of a pair turns volatile, there is a lack of incentives for liquidity providers to create positions that would facilitate trading at the current price, given the significant risk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

DeFi '22, November 11, 2022, Los Angeles, CA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9882-4/22/11...\$15.00

<https://doi.org/10.1145/3560832.3563435>

## 2 RELATED WORK

Limit order books dominate market making in traditional markets. Market making is largely in the hands of professional market makers, i.e., specialized firms. The professionals are tasked to ensure the accuracy of the price, and their ability to do so determines their own financial success. The high latency and transaction cost on the blockchain explain why CFMMs, thanks to their simplicity, have generally been used for market making. Yet, the recent introduction of concentrated liquidity CFMMs has brought market making on the blockchain closer to traditional market making as liquidity providers must select a price range. We study the price accuracy of concentrated liquidity CFMMs when prices swing unexpectedly.

Angeris et al. [12] were the first to formally study CPMMs and show that these markets must closely follow the price of the reference market. In subsequent work, Angeris et al. [11] extend the analysis to more general CFMMs. In contrast to their work, we empirically investigate the ability of Uniswap V3, an updated CFMM protocol, to track the price of the reference market.

Chitra et al. [16] compare on-chain CFMMs to conventional limit order books during periods of liveness loss. They find that both order books and concentrated liquidity CFMMs perform worse than unbounded CFMMs under loss of liveness. As opposed to studying concentrated liquidity CFMMs under liveness loss, we analyze their ability to track the price of the primary market during unexpected stablecoin price shocks.

Berg et al. [14] study the presence of market inefficiencies on Uniswap V2 by searching for and identifying cyclic arbitrage opportunities; indicative of price inaccuracies in the market. In our work, we measure price inaccuracies of Uniswap V3 by comparing the market's price to that of the primary market. Our analysis focuses on a time characterized by stablecoin price shocks.

Recently, several studies have emerged [19, 21, 23, 24] that focus on the novel CFMM design pioneered by Uniswap V3. These studies generally focus on performance and strategies for Uniswap V3 liquidity providers. On the other hand, we present the first empirical study on the ability of markets utilizing the concentrated liquidity model to track the price of the primary market.

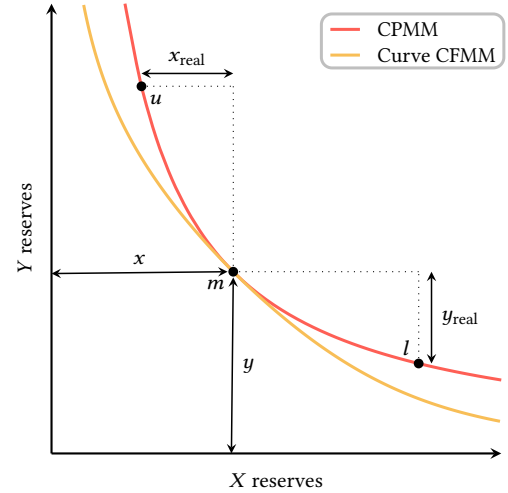
Multiple studies of cryptocurrency liquidity shocks in traditional finance have been published [17, 20, 22, 25]. Conversely, our analysis focuses on liquidity shocks on concentrated liquidity CPMMs and the role their unique market design plays therein.

## 3 CONSTANT FUNCTION MARKET MAKER

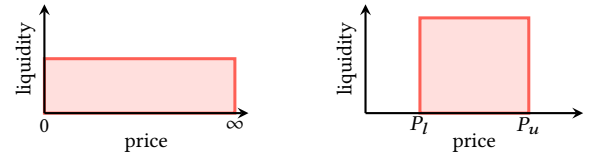
In CFMMs, trades execute automatically, and a predefined invariant controls the cryptocurrency prices. CFMMs aggregate the liquidity from many individual liquidity providers in *liquidity pools*, that are, smart contracts on the blockchain. Liquidity pools are created for each tradeable group of cryptocurrencies. Some DEXes restrict the number of cryptocurrencies in a liquidity pool to two, while others allow for liquidity pools with varying numbers of tokens.

### 3.1 Constant Product Market Maker

The most widely adopted subclass of the CFMM is the CPMM. Uniswap V3, its predecessors Uniswap V1 and V2 [9] as well as SushiSwap [5] functions as a CPMM. A CPMM ensures the product between the pool's reserves remains constant during trading. The



**Figure 1:** We draw the CPMM (red line) and Curve CFMM price curve (purple line). The CPMM price curve ensures the product invariant, and the Curve CFMM guarantees the Curve invariant. The (virtual) reserves of X- and Y-token at point  $m$  are  $x$  and  $y$ , while the marginal is given by the gradient of the slope. Note that Curve's price curve is less convex and, thus, has lower price slippage near the equal price. In turn, the price curve is steeper away from  $m$ .



(a) Uniswap V2 liquidity position (b) Uniswap V3 liquidity position



(c) collection of Uniswap V3 liquidity positions

**Figure 2:** Visualization of liquidity allocation in (concentrated liquidity) CPMMs. While a liquidity position's liquidity is distributed uniformly on Uniswap V2 (cf. Figure 2a), each Uniswap V3 liquidity position specifies a price interval  $[P_l, P_u]$  to distribute liquidity across (cf. Figure 2b). The liquidity distribution of a collection of positions across the price range on Uniswap V3 is non-uniform (cf. Figure 2c).

state of a liquidity pool with two cryptocurrencies, thus, moves along the red price curve shown in Figure 1.

Consider a liquidity pool  $X - Y$  for X-tokens and Y-tokens that holds  $x$  X-tokens and  $y$  Y-tokens. The pool's *marginal price* is then given by  $P = y/x$  and the pool's *liquidity* is  $L = \sqrt{x \cdot y}$  [9]. Note that in the original CPMM design, liquidity is unbounded, i.e., it supports trading on the pool's entire price range  $(0, \infty)$  (cf. Figure 2a).

Adams et al. [10] pioneered an overhauled CPMM design with Uniswap V3: a concentrated liquidity CPMM. A liquidity provider

concentrates their liquidity in their chosen price range  $[P_l, P_u]$ , see Figure 2b for a sample liquidity position. Importantly, each liquidity position only facilitates trading in the specified price range  $[P_l, P_u]$  and the liquidity provider only collects fees if trading occurs within this range. Therefore, liquidity providers intend to select their range around a price they view as a realistic future market price. The aggregation of the set of individual liquidity positions of the pool then gives the liquidity distribution of the pool across the entire price range (cf. Figure 2c).

To describe the pool's behavior on a price interval  $[P_l, P_u]$  with constant liquidity, Uniswap V3 utilizes the concept of *virtual reserves*. Virtual reserves simulate trading on the price interval as if the liquidity distribution on the entire price range  $(0, \infty)$  is constant (cf. Figure 2a) and matches that of the interval  $[P_l, P_u]$ . The virtual reserves behave according to the constant product price curve, as shown in Figure 1. Thus, the protocol ensures that the product of the virtual reserves  $x$  and  $y$  stays constant, i.e.,  $x \cdot y = L^2$ . Here,  $L$  is the liquidity reserved on the price interval  $[P_l, P_u]$  and the pool's marginal price is given by  $P = y/x$  [10]. The following relationship then holds between virtual reserves, real reserves, and liquidity:

$$x_{\text{real}} = x - \frac{L}{\sqrt{P_u}} = \frac{L}{\sqrt{P}} - \frac{L}{\sqrt{P_u}} \quad y_{\text{real}} = y - L\sqrt{P_l} = L\sqrt{P} - L\sqrt{P_l}.$$

We emphasize that the liquidity on the price interval  $[P_l, P_u]$  must only support trading within its price boundaries. The real reserves for the  $X$  token are fully depleted at the upper price limit  $P_u$  (cf. Figure 2b). The same holds for  $Y$ -tokens at the lower price boundary.

Note that most CPMs charge a percentage fee  $f$  on the trade input. While all Uniswap V2 pools have the same percentage fee (0.3%), each Uniswap V3 pool has a predefined percentage fee from the following set: {0.01%, 0.05%, 0.3%, 1%}.

### 3.2 Curve Constant Function Market Maker

Curve is a DEX that uses a unique CFMM that is supposed to be better suited for stablecoin pairs as the price curve is *flatter* near the peg but steeper further away compared to Uniswap V2 [18] (cf. Figure 1). However, as with Uniswap V2, liquidity is distributed across the entire price range, so we expect a qualitatively similar behavior in the event of price turbulences.

## 4 DATA COLLECTION

We analyze price data between various stablecoin pairs between 7 May and 12 May 2022. In particular, our analysis focuses on two stablecoins: UST and USDT. The price of the UST began to falter on 7 May and fell to 30 cents by 11 May. UST never fully recovered from the price drop and currently trades at a couple of cents. On 12 May, the price of USDT, the biggest stablecoin in terms of market capitalization, also dramatically dropped and landed 6% below its 1 US\$ peg. The price, however, recovered, and USDT has since been trading only slightly below its peg.

To analyze the impact of these unexpected price drops on the concentrated liquidity model introduced by Uniswap V3, we compare the price curve of stablecoin pairs, including UST or USDT, to that of Uniswap V2, Curve, and Binance. We analyze all stablecoin pairs between UST, USDT, and USDC with significant liquidity. USDC's price was stable in mid-May and is, therefore, used as a price reference for the other two stablecoins.

On Uniswap V3 we consider the following pools: USDT-USDC ( $f = \{0.01\%, 0.05\%\}$ ), UST-USDC ( $f = \{0.01\%, 0.05\%\}$ ), and UST-USDT ( $f = 0.05\%$ ). These include all pools between the three analyzed cryptocurrencies that had more than 30'000 US\$ liquidity on 6 May 2022. We include a pair's price from Uniswap V2 if a pool with more than 30'000 US\$ liquidity exists. Curve price data is inferred from the *3pool* liquidity pool, the biggest stablecoin pool on Curve that trades USDC, USDT, and DAI, as well as the *ust* liquidity pool, which trades UST, USDC, USDT, and DAI. For Binance, we collect the price data for UST-USDT and USDC-USDT, from which we infer UST-USDC price data. Binance is by far the leading cryptocurrency market in terms of volume [7], and thus we will consider its price the most reliable for the cryptocurrencies that are traded thereon.

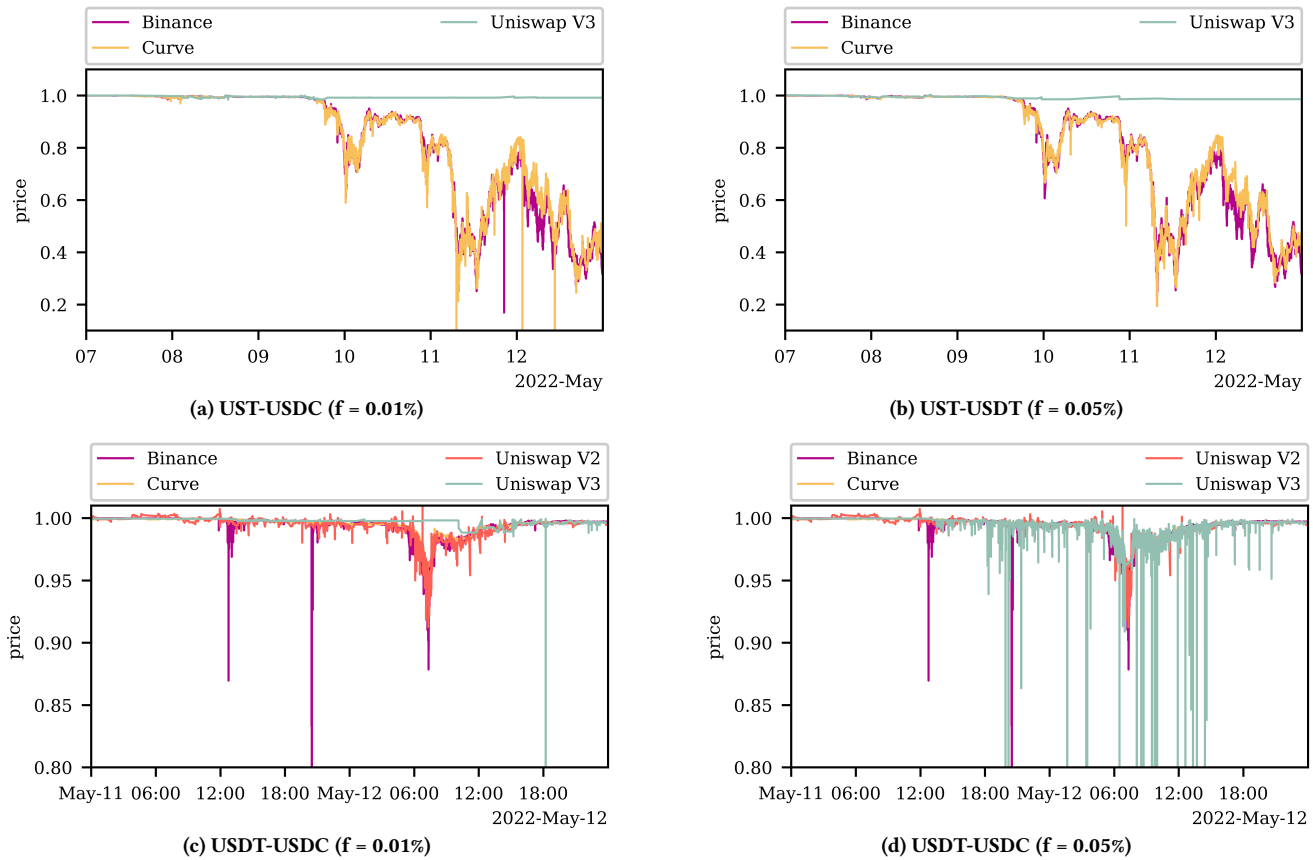
We launch an Erigon [2] client to collect data from the Ethereum blockchain between block 14'726'564 (first block of 7 May) and block 14'764'083 (last block of 12 May). In particular, we filter event logs for all events related to Uniswap V2 and V3 as well as Curve. Price data from Binance is obtained from their market data API [3] and downloading trade data for the relevant trading pairs.

## 5 ANALYSIS

We begin the analysis by studying the price of various stablecoin pools on Uniswap V3 and comparing them with the price on Binance, which we consider to be the most accurate, as well as Uniswap V2 and Curve. In Figure 3, we plot these price curves for four of the five Uniswap V3 pools we analyze. Note that we only plot the price curve for the UST-USDC pool with the smaller fee tier as the behavior of the other UST-USDC pool is very similar.

We immediately notice that the price between the pool's two assets is less likely to move for the analyzed Uniswap V3 pairs than the price on Uniswap V2, Curve, or Binance. For instance, we plot the price of UST in terms of USDC in Figure 3a and USDT in Figure 3b. We observe that UST's price on Uniswap V3 never falls significantly below its peg of 1 US\$. On Binance, however, the price of UST dropped from its peg to around 0.30 cents in the same time frame. The price on Curve (cf. Figure 3b) tracks the price on Binance closely, as opposed to Uniswap V3's price. It appears that there was no significant liquidity available for the Uniswap V3 stablecoin pools that held UST as one of their assets to continue to facilitate trading. Considering the very significant and unexpected price drop of UST, it would not be uncommon for financial markets to eventually stop trading assets given such events. Binance, for instance, stopped trading UST on 13 May due to its dramatic price drop [4]. However, looking at Uniswap V3, we find that it could not even follow the initial smaller price fluctuations on 9 May.

The picture becomes even more startling when we study the price curves of the two Uniswap V3 pools trading USDT and USDC (cf. Figures 3c and 3d). USDT experienced an abrupt but less significant price drop than UST on 11 and 12 May, from which it has since largely recovered. When analyzing Figures 3c and 3d, we observe that while Uniswap V3 struggles significantly to follow these small but unexpected price fluctuations, Curve and Uniswap V2 can easily keep up. In the USDT-USDC ( $f = 0.01\%$ ) pool, the price of Uniswap V3 appears to react only after 36 hours. By then, the price of USDT had already recovered significantly. However, the price of USDT in the Uniswap V3 pool still experienced a sudden and short price



**Figure 3:** We plot the price between various stablecoin pairs on Uniswap V3, Uniswap V2 (where applicable), Curve, and Binance. Figures 3a and 3b show the price of UST in terms of USDC and UST, respectively, in the Uniswap V3 pools along with the price curves from Curve and Binance. In Figures 3c and 3d we plot the price of USDT in terms of USDC for the Uniswap V3 pool with fees 0.01% and 0.05% respectively. The corresponding price curves for Uniswap V2, Curve, and Binance are included.

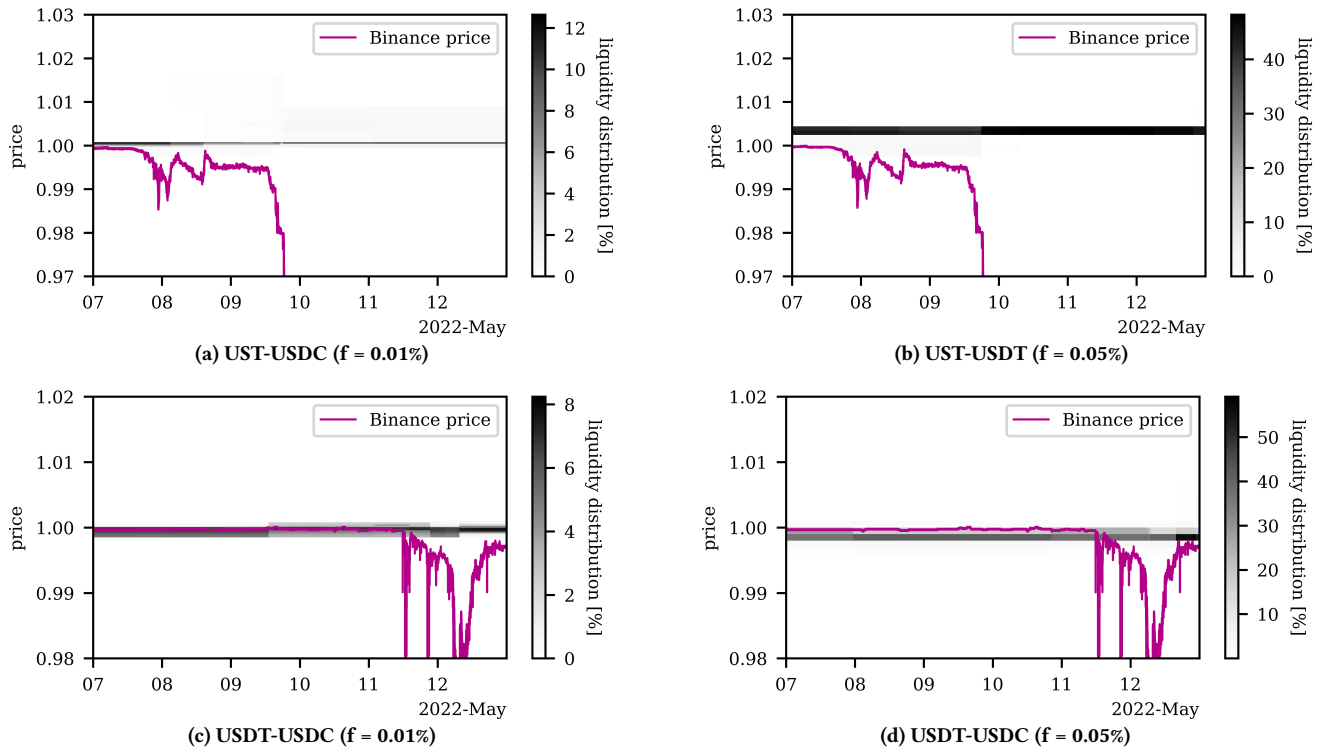
drop on 12 May at 18:00. Two similar price drops occurred on Binance a day earlier and in all three cases could be caused by low liquidity levels. The USDT-USDC ( $f = 0.05\%$ ) pool reacted much more quickly to the USDT price drop. However, the liquidity pool appears to overreact, i.e., experience sudden price drops repeatedly. The USDT price drops significantly below the price observed in the other three markets a couple of dozen times.

To better understand the failure of Uniswap V3 to react to UST and USDT price fluctuations, as well as sudden price drops, we visualize the liquidity distribution over time in the four Uniswap V3 pools (cf. Figure 4). We notice that, as expected, the vast majority of the liquidity is distributed around the equal price of the pool's assets in our four pools. In the two liquidity pools that hold UST as one of their assets (cf. Figures 4a and 4b), the relative liquidity distribution changes only very slightly over time. On 12 May, when the price of UST had not been near its peg for a couple of days, the pool's liquidity was still concentrated around the peg. In fact, by the end of 12 May, more than 99% of liquidity was still on the price interval  $[0.99, 1.01]$  in both of the UST liquidity pools shown. Thus, it appears that liquidity providers are not adjusting their positions quickly enough, which explains why the pool's price did not track the primary market's price.

In Figures 4c and 4d, we visualize the liquidity distribution in the two USDT-USDC Uniswap V3 pools and again notice that there are little to no changes in the liquidity distribution despite the price drop of USDT. Liquidity was not concentrated around the price but around USDT's intended peg. Thus, when the price moves outside the small interval around the equal price of the two stablecoins, there is little liquidity to support trading – illustrating why the frequent sharp price drops in the Uniswap V3 pools were possible.

While the previous analysis outlined that the liquidity distribution across the price interval only saw very few adjustments in the face of dramatic price drops, even after a few days, we analyze whether and how quickly liquidity was removed from the Uniswap V3 stablecoin pools that held UST as one of their assets in Figure 5. In Figures 5a, 5b and 5c we visualize the liquidity levels in three Uniswap V3 pools: UST-USDC ( $f = 0.01\%$ ), UST-USDC ( $f = 0.05\%$ ) and UST-USDT ( $f = 0.05\%$ ). Note that we plot both the value of the liquidity when it was deposited (in green) and the current value of the liquidity (in red) in US\$ for each timestamp. The current value of the liquidity accounts for UST's price drop. In all three pools, we observe that only a couple of significant liquidity providers remove their position from the liquidity pools before UST plummets below 0.95 cents on 9 May (cf. Figure 3). This is especially apparent in





**Figure 4:** We visualize the liquidity distribution on the price interval over time and include the price curve of the primary market (Binance) as a reference. The depth of liquidity in the price interval is visualized by the intensity of the shaded area; darker price intervals hold more liquidity than lighter price intervals at the same time. Liquidity is represented relatively at a point in time. Observe that the Binance price exits the price range for which trading is possible on Uniswap V3.

the UST-USDC ( $f=0.05\%$ ) pool, which we show in Figure 5b. Three large positions with an average value of 1M US\$, accounting for three-fourths of the pool’s liquidity, are removed from the pool by midday on 9 May. From that point on, we see the difference between the value deposited and current begin to divide significantly, i.e., the liquidity providers in the pool start to experience significant losses. We further note that the liquidity providers were no longer earning fees, as the pools had stopped trading (cf. Figure 3). Observe that even by the end of 12 May, by that time UST had not been trading near its peg for several days, more than half of the pool’s liquidity positions were still active (cf. Figure 5d). We conclude that the small number of liquidity providers in the three pools did not adapt to the changes in the price of the pool’s assets quickly enough to ensure that trading could continue or even to remove their assets from the pool in time to avoid significant financial losses.

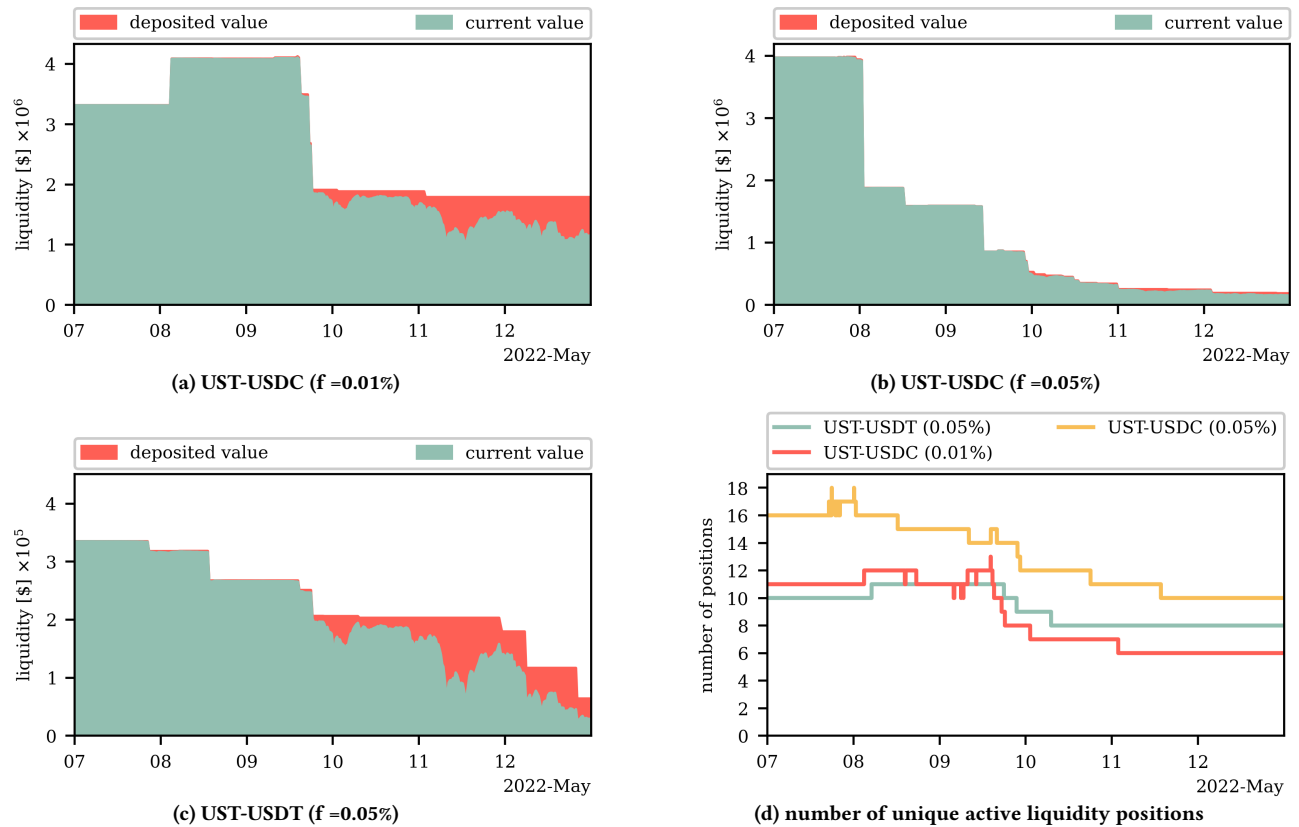
## 6 DISCUSSION

Uniswap V3 does not appear equipped to handle volatile cryptocurrency markets. Its Achilles heel partially lies in its reliance on liquidity providers to react to market changes. While market making has long been handled by a few professional market makers in traditional finance, we observed that Uniswap V3 liquidity providers as a whole currently lack the sophistication and agility required for complex market making. The liquidity providers’ apparent inexperience calls into question the suitability of the concentrated

liquidity CPMM for DeFi. Not only did we observe that Uniswap V3 could not track the price of the primary market, but also saw that both its predecessor and Curve were fit for the task. While Uniswap V2 does not offer the same capital efficiency that Uniswap V3 for stablecoin pairs that continuously trade at an equal price, Curve’s price curve can also reach high levels of capital efficiency around the equal price.

Providing liquidity bears the risk of seeing the value of the liquidity position decrease compared to the initial assets deposited. This phenomenon is known as *impermanent loss* and is particularly pronounced for high price volatility [21]. Given the speed of such an abrupt price drop and the general risk-off market sentiment that prevails in such a crisis, liquidity providers may be reluctant to open new liquidity positions.

Furthermore, note that pairs deemed to have a stable price, namely stablecoin pairs, generally have all liquidity in pools with a small percentage fee. Thus, once the price of a supposedly stable cryptocurrency pair becomes volatile, there are too few incentives for liquidity providers to create new liquidity positions. As a pool’s fee cannot be adjusted, the meager fees liquidity providers expect to receive in the pools do not compensate them for the risk they face given the elevated volatility. Therefore, before supplying liquidity, they might choose to create a new liquidity pool. In fact, new pools that hold UST with high fees were launched after the price drop, such as the UST-USDC ( $f = 1\%$ ) pool. However, the pool was only created on 12 May, days after UST had lost its peg and



**Figure 5:** We visualize the liquidity levels over time in three Uniswap V3 pools: UST-USDC ( $f=0.01\%$ ) (cf. Figure 5a), UST-USDC ( $f=0.05\%$ ) (cf. Figure 5b) and UST-USDT ( $f=0.05\%$ ) (cf. Figure 5c). We plot both the original and current US\$ value of the liquidity that remains in the pool. Note that a pool’s liquidity does not match the total value locked in the pool, as a consequence of the collected fees not acting as liquidity. In Figure 5d we further plot the number of active liquidity positions in the three pools over time. Notice that the number of unique liquidity positions is small and does not even half during UST price decline.

never acquired the same levels of liquidity as the existing pools with lower fees. Thus, liquidity was further diluted across several pools, which reduces capital efficiency.

Given this unfavorable risk-reward ratio for liquidity providers to deposit liquidity around the new price during times of market distress, DEXes could compensate liquidity providers for taking on this risk. In fact, liquidity providing on Uniswap V3 has been shown to exhibit similarities with financial derivatives [21]. Thus, just as option prices rise with volatility [15], protocols could adopt a fee structure that would increase with volatility. Such behavior also occurs in traditional financial markets. During times of market upheaval the bid-ask spread increases as market makers, who generate their profit from this spread, seek compensation for the risks arising due to the rapid price changes [8].

Finally, we also want to emphasize that the relative importance we should place on capital efficiency, in comparison to price accuracy, is unclear. The security implications of price inaccuracies and low levels of liquidity outside a small price interval around the equal price could allow for an inexpensive attack geometric mean time-weighted average price (TWAP) on Uniswap V3 [13]. During such an attack, the attacker profits from manipulating Uniswap V3’s spot price for the duration of a block. Not only do we already

observe repeated price drops in the USDT-USDC ( $f=0.05\%$ ) pool, but we also found that in the UST pool, very little liquidity was outside the price interval  $[0.99, 1.01]$  around the equal price at the end of 12 May. To give an example, on May 12 a single transaction exchanged 999’242 USDT for 995’221 USDC in the USDT-USDC ( $f=0.01\%$ ) pool. Thus, the trade executed at the price of 0.9959 (USDC/USDT). In the process, the trade moved the pool’s price from 0.9993 to 0.4999. At the same time, the price on Binance was 0.9966, and we see that the trade almost executed at market price but created a huge price inaccuracy as it entered regions on the price interval with little liquidity. Thus, it was relatively inexpensive for an attacker to move the price.

## 7 CONCLUSION

Our empirical analysis finds that the concentrated liquidity CPMM pioneered by Uniswap V3 is currently not ready to handle unexpected price drops. Not only did we observe significant price inaccuracies and financial loss on the liquidity provider side, but we also see the potential for TWAP attacks on Uniswap V3 oracles. For the success of Uniswap V3, it is thus imperative that liquidity providers will become more sophisticated and agile.

## REFERENCES

- [1] 2022. Curve. <https://curve.fi/>.
- [2] 2022. Erigon. <https://github.com/ledgerwatch/erigon>.
- [3] 2022. Historical Market Data. <https://www.binance.com/en/landing/data>.
- [4] 2022. LUNA & UST Trading Suspended. <https://www.binance.com/en/support/announcement/f68451879a1841a6a0f44025735d9236>.
- [5] 2022. Sushiswap. <https://sushi.com/>.
- [6] 2022. TerraUSD. <https://coinmarketcap.com/de/currencies/terrausd/>.
- [7] 2022. Top Cryptocurrency Spot Exchanges. <https://coinmarketcap.com/rankings/exchanges/>.
- [8] 2022. Understanding bid ask spreads. <https://www.rbcgam.com/en/ca/learn-plan/types-of-investments/understanding-bid-ask-spreads/detail>.
- [9] Hayden Adams, Noah Zinsmeister, and Dan Robinson. 2020. Uniswap v2 Core.
- [10] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. 2021. Uniswap v3 core.
- [11] Guillermo Angeris and Tarun Chitra. 2020. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*.
- [12] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2019. An analysis of Uniswap markets.
- [13] Bentley. 2021. Manipulating Uniswap V3 TWAP Oracles.
- [14] Jan Arvid Berg, Robin Fritsch, Lioba Heimbach, and Roger Wattenhofer. 2022. An Empirical Study of Market Inefficiencies in Uniswap and SushiSwap. In *The 2nd Workshop on Decentralized Finance (DeFi), Grenada*.
- [15] Fischer Black and Myron Scholes. 1973. The Pricing of Options and Corporate Liabilities. *The Pricing of Options and Corporate Liabilities* 81, 3 (1973), 637–654.
- [16] Tarun Chitra, Guillermo Angeris, and Alex Evans. 2021. How Liveness Separates CFMMs and Order Books.
- [17] Shaen Corbet, Yang Greg Hou, Yang Hu, Charles Larkin, Brian Lucey, and Les Oxley. 2022. Cryptocurrency liquidity and volatility interrelationships during the COVID-19 pandemic. *Finance Research Letters* 45 (2022), 102137.
- [18] Michael Egorov. 2019. StableSwap-efficient mechanism for Stablecoin liquidity.
- [19] Robin Fritsch. 2021. Concentrated Liquidity in Automated Market Makers. In *Proceedings of the 2021 ACM CCS Workshop on Decentralized Finance and Security (DeFi@CCS), Virtual Event, Republic of Korea*.
- [20] John M Griffin and Amin Shams. 2020. Is Bitcoin really untethered? *The Journal of Finance* 75, 4 (2020), 1913–1964.
- [21] Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. 2022. Risks and Returns of Uniswap V3 Liquidity Providers. In *4th ACM Conference on Advances in Financial Technologies (AFT), Cambridge, Massachusetts, USA*.
- [22] Agata Kliber. 2018. Price, liquidity and information spillover within the cryptocurrency market. The case of bitfinex. *Safe Bank* 73, 4 (2018), 62–79.
- [23] Stefan Loesch, Nate Hindman, Mark B Richardson, and Nicholas Welch. 2021. Impermanent Loss in Uniswap v3.
- [24] Michael Neuder, Rithvik Rao, Daniel J Moroz, and David C Parkes. 2021. Strategic Liquidity Provision in Uniswap v3.
- [25] Tao Tang and Yanchen Wang. 2022. Liquidity Shocks, Price Volatilities, and Risk-managed Strategy: Evidence from Bitcoin and Beyond. *Journal of Multinational Financial Management* (2022), 100729.
- [26] Gavin Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger.