Statistical vulnerabilities in Uniswap V3 Price Oracle

Pavel Anokhin

***Abstract*—Decentralized applications in blockchain make extensive use of the price oracles for a variety of purposes. All price oracle providers possess some unique security risks. This paper studies the statistical properties of one of the very popular price oracle providers – Uniswap v3 price oracle, and the statistical vulnerabilities which might arise from the usage of this price oracle. The study is performed using the Uniswap v3 historical trading data from the blockchain and the historical trading data from the Binance exchange for the ETH / USDT cryptocurrencies trading pair. The security consequences and possible statistical vulnerabilities which arise from the usage of Uniswap v3 oracle price based on the results of the study are discussed and the recommendations on procedures and measures to mitigate these vulnerabilities are given. The study can be of a practical use to blockchain software developers and smart contract auditors to better secure applications using the Uniswap v3 price oracles.**

***Keywords*—statistical vulnerability, smart contract, blockchain security, price oracle, Uniswap.**

I. Introduction

Decentralized finance (DeFi) is a term used for financial services using blockchain technology which do not depend on intermediaries such as banks or brokers. Many DeFi services and technologies are novel and vastly different from the traditional finances [1].

One of such novel DeFi technologies is automated market maker (AMM) – a mechanism used to calculate buy and sell price for any asset based on some formula, which allows to provide ongoing market pricing for market participants [2]. Currently, one of the most popular and one of the largest AMMs (in terms of total value locked and trading volume) used in practice is Uniswap, which is an AMM based on the constant product pricing formula [3].

One of the advantages of the Uniswap is price discovery – it doesn't need any external oracles to price the assets. The price only depends on the amount of 2 assets in the Uniswap pool, requiring the product of the amounts of 2 assets in the pool to be the same (or greater) after each trade. Uniswap price will be close to market price, because arbitrageurs will extract near risk-free profit if the price is different from the other exchanges thus pushing Uniswap price close to the market price.

Due to this useful price discovery property, Uniswap price is often used as a blockchain-native price oracle for DeFi applications that need to know some asset price (for example, lending services need asset prices to calculate account health based on debt and collateral in different assets). However, using instant (spot) Uniswap price is known to be very vulnerable to price manipulation and flash loan attacks which was proven by many hacks. Different examples of such attacks are given in [4].

In order to solve this problem, Uniswap v2 has introduced an accumulator (sum of time-weighted prices), which can be used to calculate arithmetic time-weighted average price (TWAP) of any asset for arbitrary period, although the usage in practice was not very convenient as the applications had to save accumulator value to be able to calculate TWAP for the period it needed.

Uniswap v3 has introduced a better price oracle with built-in recording of the accumulator value so that mean price can be calculated for arbitrary (although small) time period. However, Uniswap v3 price oracle records prices in *ticks*, where the price is calculated using the equation (1):

$$price = 1.0001^{tick} \tag{1}$$

The oracle returns average *tick*, rather than *price*, which means that Uniswap v3 price oracle is the *geometric* mean whereas Uniswap v2 price oracle is the *arithmetic* mean. This has certain security consequences, which were extensively researched [5]. However, to the best of the author's knowledge, all existing research of Uniswap v3 price oracle focuses on its resistance to price manipulation, but no research exists which analyzes statistical properties and possible statistical vulnerabilities of the Uniswap v3 price oracle.

The goal of this study is to analyze statistical properties of the Uniswap v3 price oracle compared to market price and determine possible statistical vulnerabilities based on real trading data from blockchain.

II. Data Source

In order to analyze Uniswap price oracle, it has to be compared to a market price. We use real trading price data in 1 second intervals from the leading centralized crypto-exchange Binance [6], which provides historical data at its website [7]. For the spot market analyzed in this study, Binance's trading volume exceeds $300 million per day, which is sufficient to consider the Binance trading price to be close to real market price.

For the Uniswap v3 price data, a program was coded in C# to download the data directly from Ethereum main-net blockchain. For each block, the program downloaded the following data: block timestamp, current pool price, TWAP oracle price with time periods of 1, 5 and 20 minutes. The main part of the program's listing is given below:

Manuscript received July 30, 2023.

Pavel Anokhin is independent smart contract security expert (e-mail: pavel@anokhin.name)

```
public      double      GetAvgPrice(ContractHandler
contract, uint timeSec, ulong block = 0)
{
  BlockParameter blockParam =
    block == 0 ? null : new BlockParameter(block);
  var observeFunction =
    new UniswapV3Pool.ObserveFunction();
  List<uint> timeList = new List<uint>();
  timeList.Add(timeSec);
  timeList.Add(0);
  observeFunction.SecondsAgos = timeList;
  var res =
    contract.QueryDeserializingToObjectAsync
    <UniswapV3Pool.ObserveFunction,
    UniswapV3Pool.ObserveOutputDTO>
    (observeFunction, blockParam).Result;
  long deltaTick = res.TickCumulatives[1] -
    res.TickCumulatives[0];
  double avgTick = 1.0 * deltaTick / timeSec;
  double price = Math.Pow(1.0001, avgTick);
  return price * Math.Pow(10, 12);
}
…
for (ulong block = fromBlock; block <= toBlock;
  block++)
{
  long blockTimestamp = (long)mWeb3.Eth.Blocks.
    GetBlockWithTransactionsHashesByNumber.
    SendRequestAsync(new
    BlockParameter(block)).Result.Timestamp.Value;
  var slot0 = contract.
    QueryDeserializingToObjectAsync<
    UniswapV3Pool.Slot0Function,
    UniswapV3Pool.Slot0OutputDTO>(null,
    new BlockParameter(block)).Result;
  BigInteger bigPrice = slot0.SqrtPriceX96;
  bigPrice = (bigPrice * bigPrice
    * BigInteger.Pow(10, 18+12)) >> 192;
  double price = MyMath.ToDouble(bigPrice, 18);
  double price1 = GetAvgPrice
    (contract, 60 * 1, block);
  double price5 = GetAvgPrice
    (contract, 60 * 5, block);
  double price20 = GetAvgPrice
    (contract, 60 * 20, block);
  string s = string.Format(
    "{0};{1};{2};{3};{4};{5}",
    block, blockTimestamp,
    price, price1, price5, price20);
  writer.WriteLine(s);
}
…
```

The trading data was downloaded for the market of swaps between cryptocurrencies Ethereum (ETH) and Tether (USDT) for the period from June 21, 2023 to July 20, 2023. In Ethereum main-net this period corresponded to block numbers from 17524329 to 17737810. For Uniswap v3, the data from the following pool addresses was downloaded:

1. ETH / USDT (0.3% fee) pool address: 0x4e68ccd3e89f51c3074ca5072bbac773960dfa36
2. ETH / USDT (0.05% fee) pool address: 0x11b815efb8f581194ae79006d24e0d814b7697f6

Additionally, to verify certain properties during extreme market conditions, 1-minute market data for swaps between LUNA and USDT for the period from May 8, 2022 to May 13, 2022 was downloaded from Binance (LUNA cryptocurrency price has crashed from $60 to less than $0.01).

Since Ethereum blocks are currently mined every 12 seconds and Binance price data is available in 1 second intervals, in order to match the prices between Binance and Uniswap, the open Binance price for the timestamp of each block's timestamp was taken, meaning for each block (for every 12 seconds) the following data is available:

1. Timestamp.
2. Binance price (opening price of 1-second interval starting at the block's timestamp).
3. Uniswap v3 current price, 1-minute average, 5-minute average and 20-minute average prices for 0.3% fee pool and for 0.05% fee pool.

III. Comparison of Uniswap v3 price oracle and Binance price

A. *Factors influencing Uniswap v3 oracle price*

In this study we will compare Uniswap v3 oracle price with arithmetic mean of the Binance price for the same time periods as oracle. We use arithmetic mean for comparison because it is possible to execute real trades at the arithmetic mean price (trading equal amounts in Binance each 12 seconds), but it is not possible to execute real trades at the geometric mean price. As such, real security vulnerabilities may arise from deviation of oracle price from real trading opportunities (arithmetic mean prices), but not from the geometric mean price.

The following factors influence deviation of Uniswap v3 oracle price from the arithmetic mean of the market (Binance) price:

1. Usage of *tick* instead of *price* in Uniswap v3 oracle, which loses price precision.
2. Difference between the geometric mean and the arithmetic mean.
3. Uniswap trading fee.

It is possible to analyze each factor's impact separately from the other factors and then analyze all factors combined, which is done in the following sections.

B. *Price precision loss due to usage of ticks instead of price in Uniswap v3 oracle*

Uniswap v3 price oracle records *tick* and returns an average *tick* over a time period, rather than price. The price is then calculated using the equation (1). This means that the price is truncated to the closest power of 1.0001 when calculating average. For example, if the current price is 2980, the average price returned for 1 second will be 2979.77, because price at tick 80000 is 2979.77 and price at tick 80001 is 2980.06, thus the current price of 2980 is truncated to a smaller tick.

The impact of the usage of *ticks* instead of price in Uniswap v3 oracle is reduction of the oracle price by at most 0.01% from actual trading price. Average reduction is 0.005% due to random uniform distribution of price between tick boundaries: from 0% to 0.01%.

C. *Difference between the geometric mean and the arithmetic mean*

Uniswap v3 price oracle uses the geometric mean, while real trading is only possible using the arithmetic mean price. The geometric mean equals the arithmetic only when the price is constant throughout the whole time period. In all the other cases the geometric mean is strictly less than the arithmetic mean. In normal trading circumstances the difference is usually small enough to ignore it, but in extreme cases the difference becomes large enough to be able to exploit it, meaning it becomes a statistical vulnerability.

The worst case for the largest deviation of the geometric mean from the arithmetic mean is a sharp drop (or rise) of the price. For comparison, 2 cases of the price charts are shown in Figure 1: case 1 is a sharp drop in price and case 2 is a gradual price decrease over 100 blocks, in both cases price decreases by 10% and is constant before and after the drop.

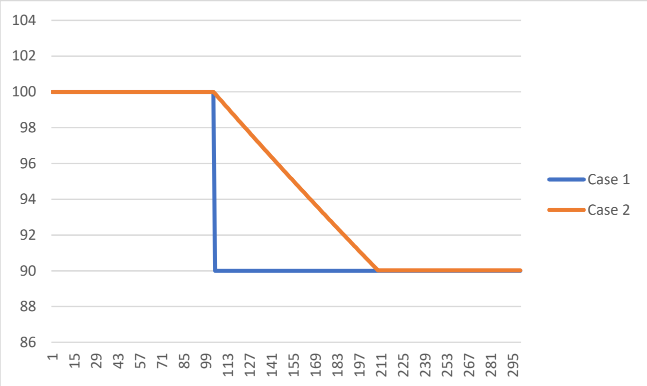

Figure 1. Price chart for the worst case of the arithmetic and the geometric mean deviation

Corresponding deviation of the geometric mean from the arithmetic mean over a moving window of 100 blocks for both cases is shown in Figure 2.

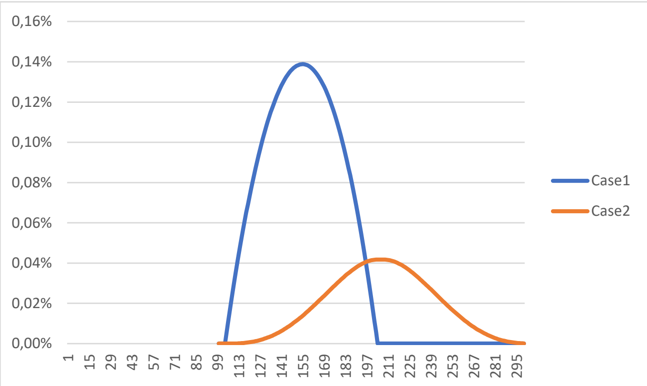

Figure 2. Deviation of the geometric mean from the arithmetic mean during instant price drop (Case 1) and gradual price decrease (Case 2)

As can be seen from the chart, a sharp drop by 10% makes the geometric mean to deviate by at most 0.14% from the arithmetic mean. And a more gradual decline in price reduces maximum deviation to 0.04%.

The actual deviation of the geometric mean from the arithmetic mean calculated from Uniswap ETH / USDT 0.3% fee current pool prices over the studied 1-month period using the moving window of 100 blocks (20 minutes) is shown in Figure 3. The max deviation over 1 month period was 0.0066% and average deviation was 0.00003%. The numbers were even lower for the 0.05% fee pool.

Since the worst-case scenario for the deviation is during a sudden large price drop, we have also tested deviation of the geometric mean from the arithmetic mean during the LUNA cryptocurrency crash starting on the 9th of May, 2022. Unfortunately, on-chain data was very unreliable during that period of time due to network congestion, so we have used Binance trading data instead.

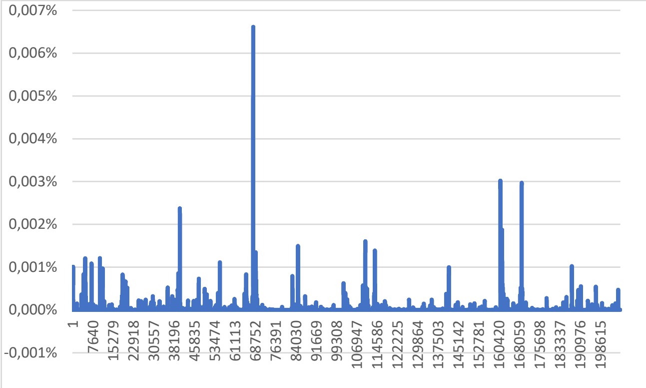

Figure 3. Deviation of the geometric mean from the arithmetic mean of the Uniswap ETH / USDT 0.3% fee pool prices using moving window of 100 blocks

The price chart during the LUNA crash is shown in Figure 4 and the deviation of the geometric mean from the arithmetic mean for moving windows of 5, 25, 100 blocks (1, 5 and 20 minutes) is shown in Figure 5.

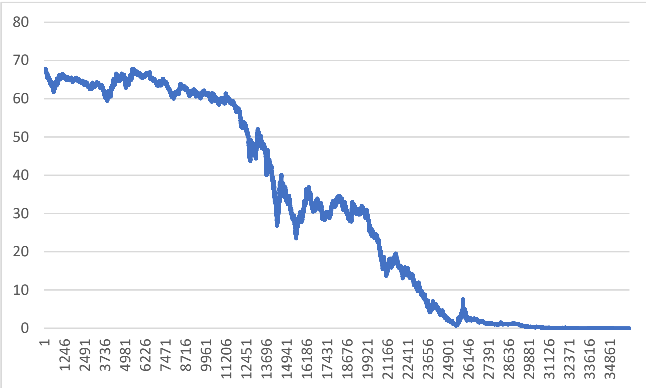

Figure 4. Price chart of LUNA crash for the May 8 – May 13, 2022 period.

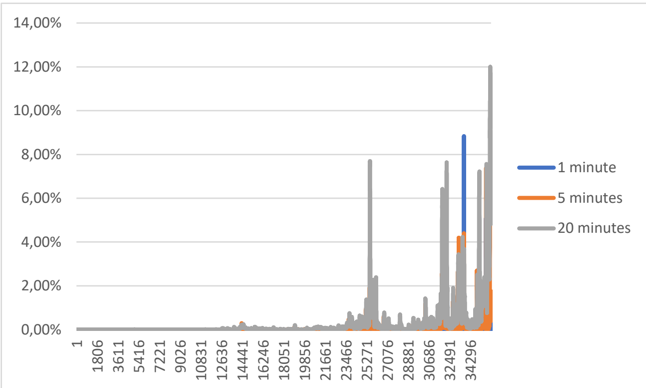

Figure 5. Deviation of the geometric mean from the arithmetic mean of the Binance LUNA / USDT trading price during LUNA crash

Maximum deviation observed during LUNA crash was 12% and it was above 1% for extended periods of time.

The impact of the Uniswap v3 oracle being geometric mean vs natural arithmetic mean is almost non-existent (below 0.01%) during normal market conditions. However, during extreme market conditions such as quick crash of the cryptocurrency, the impact increases to 1%-10%.

Table I. Statistical properties of Binance and Uniswap prices

Property	Binance Price	Uniswap 0.05% fee pool		Uniswap 0.3% fee pool	
		Price	Deviation from Binance	Price	Deviation from Binance
Volatility (annualized)	38%	29%	73%	17%	261%
Standard Deviation	0.023%	0.018%	0.045%	0.01%	0.16%
Autocorrelation (12 seconds)	0.038	0.056	0.88	0.138	0.99
Autocorrelation (1 minute)	0.005	0.037	0.61	0.066	0.95
Autocorrelation (5 minutes)	0.017	0.002	0.19	0.021	0.82
Autocorrelation (20 minutes)	0.011	0.006	0.027	0.005	0.54

D. *Influence of Uniswap trading fee*

In the ideal world the Uniswap pool price should be equal to market price (so Uniswap price equal to Binance price). However, due to fees paid by traders for swapping via Uniswap, it is not economically profitable to push the price to equivalence. For example, if 1 ETH price is 1003 USDT in the Binance exchange and 1000 USDT in the Uniswap pool with 0.3% fee, then the real Uniswap price is actually 1000 USDT + fee of 3 USDT = 1003 USDT total, meaning the arbitrage is absent even though the prices differ by 0.3%. There is also a gas cost associated with Uniswap transaction. Even though the gas cost is fixed (regardless of swapped amount), it can be very significant in percentage terms for swaps of small amounts, further restricting arbitrageurs from bringing the price closer to the market.

Due to fees, we can expect that Uniswap and Binance prices should be very close, but can deviate from each other by about the fees (0.3% or 0.05% depending on the pool fee). This is visible in the Figure 6, which shows the chart of the ETH/USDT pair price in Binance and Uniswap (0.3% and 0.05% fee pools) during the first 1000 blocks of the period studied.

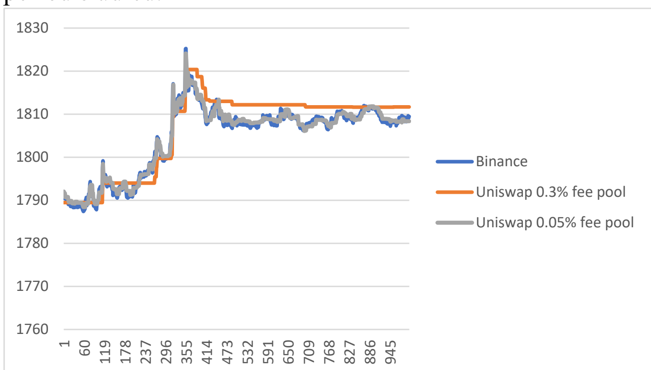

Figure 6. Comparison of ETH/USDT prices between Binance and Uniswap (0.3% and 0.05% fee pools)

As can be seen from the charts in Figure 6, Uniswap 0.05% fee pool tracks the Binance price much better than Uniswap 0.3% fee pool. Additionally, Uniswap 0.3% fee pool can be above (or below) Binance price for extended period of time due to fees. Figure 7 shows deviation of Binance and Uniswap 0.3% fee pool prices over the first 1000 blocks of the period studied.

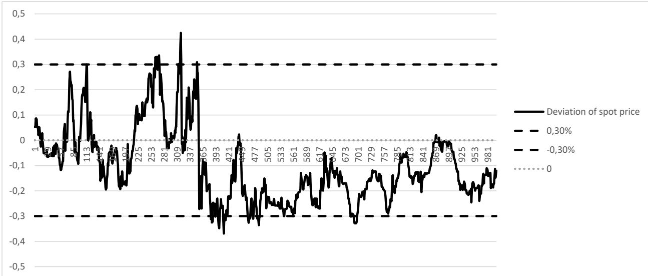

Figure 7. Deviation of Binance and Uniswap 0.3% fee pool prices over the first 1000 blocks of the period studied

As expected, the deviation is mostly within the boundaries of the 0.3% fee, but it can also be seen that the sign of deviation stays the same for extended periods of time, which is a sign of strong autocorrelation and predictability of the deviation. Calculated statistical properties of Binance and Uniswap (0.3% fee pool and 0.05% fee pool) prices, including statistical properties of deviation of Binance and Uniswap prices is show in Table I. Deviation between Uniswap and Binance prices (both 0.3% and 0.05% fee pools) exhibits very strong autocorrelation, which falls for longer periods of time (higher order autocorrelation). The autocorrelation coefficient of the 5th order (1 minute or 5 blocks) is high both for 0.05% fee pool (0.61) and especially for 0.3% fee pool (0.95), which means that current deviation is a very strong predictor of the future deviation (1 minute into the future). For the 0.3% fee pool autocorrelation stays very high even for 20 minutes (100 blocks) with a value of 0.54.

High autocorrelation coefficient of the Uniswap and Binance price deviation allows to use a very simple algorithm to predict future mean Uniswap price deviation from Binance:

- The expected sign of future deviation is the same as the sign of current deviation.
- The greater the absolute value of the current deviation, the greater the expected absolute value of future deviation.

To find the expected future value of deviation based on current value of deviation, we have performed a mathematical modeling, which has calculated average arithmetic mean of deviation over the future time interval *T* based on the interval of the current deviation. The current deviation was broken down into 60 intervals (from -0.3% to 0.3% with a 0.01% step) and for each interval of current deviation, the arithmetic mean of price deviation over the time interval from current time to current time + *T* was calculated, and average of all these arithmetic means was calculated. The arithmetic mean was used to isolate the fee factor from the other factors of Uniswap v3 oracle.

The results of the mathematical modeling for the 0.3% fee pool are shown in Figure 8 and for the 0.05% fee pool in Figure 9.

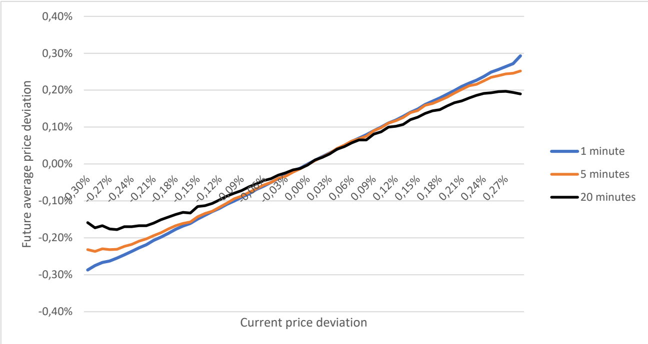

Figure 8. Dependance of future average Binance and Uniswap 0.3% fee pool price deviation on the current price deviation for 1, 5 and 20 minutes.

As can be seen from Figure 8, for the 0.3% fee Uniswap pool it is possible to predict the price deviation of up to 0.3% for 1 minute, up to 0.25% for 5 minutes and up to 0.15% for 20 minutes time interval.

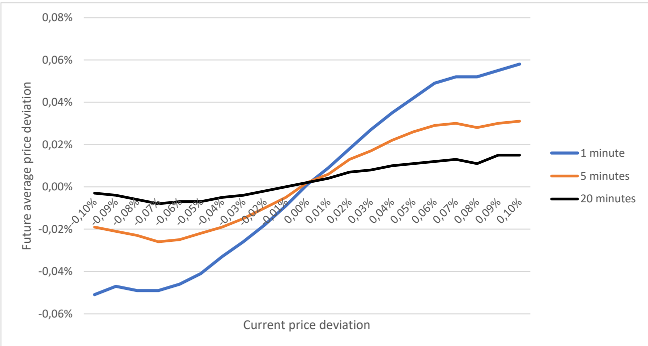

Figure 9. Dependance of future average Binance and Uniswap 0.05% fee pool price deviation on the current price deviation for 1, 5 and 20 minutes.

For the 0.05% fee Uniswap pool it is possible to predict the price deviation of up to 0.05% for 1 minute, up to 0.025% for 5 minutes and up to 0.01% for 20 minutes time interval.

E. *Influence of all factors combined*

Analysis of influence of all 3 factors to Uniswap v3 oracle price is summarized in the Table II.

As can be seen from the Table II, first 2 factors only reduce the oracle price by a small amount, while the 3rd factor allows to predict future average price deviation within the fee interval and also depends on the time interval used. Factor 2 can lead to a large deviation during extreme market conditions, but most of the time the main dominating factor is factor 3 (influence of fees to oracle price).

In order to estimate the influence of all factors combined, we have modified the mathematical modeling used to analyze factor 3: instead of using arithmetic mean of Uniswap pool prices over period from current time to current time + *T*, we have used the actual value returned by the Uniswap pool oracle at the time of current time + *T* for the *T* seconds ago. Using such modified mathematical model has allowed us to estimate the predictability of the Binance and Uniswap oracle price deviation based on current deviation of spot Binance and Uniswap prices.

Table II. Influence of different factors to Uniswap v3 oracle price

Factor	**Influence**
Usage of *ticks* (Loss of precision)	Always reduces oracle price from 0% to 0.01% due to truncation. Average price reduction by 0.005%.
Geometric mean vs arithmetic mean	Always reduces oracle price, because geometric mean is always less than or equal to arithmetic mean. Average reduction of oracle price by 0.00003%. Max reduction of 0.0066% during normal market conditions. Reduction of oracle price by 1%-10% during extreme market conditions, such as LUNA price crash.
Trading fee	Can increase or reduce oracle price by up to fee percentage. The deviation of oracle price from market price can be predicted for small enough time intervals.

The results of the modified modeling for the Uniswap 0.3% fee pool are shown in Figure 10. The results are almost identical to the results shown in Figure 8.

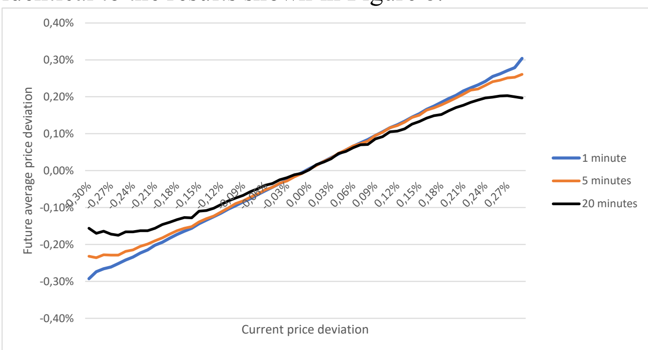

Figure 10. Dependance of average future deviation of Binance price and Uniswap 0.3% fee pool price oracle on the current price deviation for 1, 5 and 20 minutes.

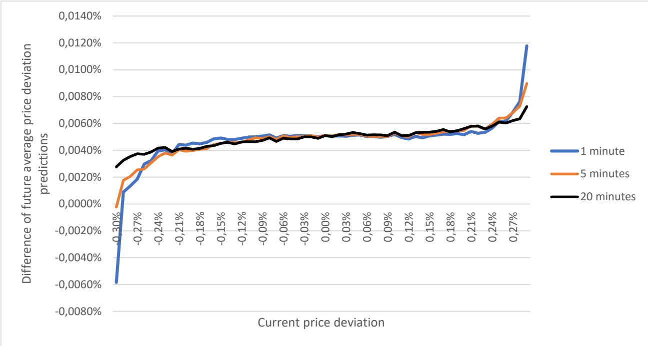

Figure 11. Difference of dependance of average future deviation of Binance and Uniswap 0.3% fee pool price due to factor 3 and due to all 3 factors.

To better see the difference, a chart of the difference between Figure 8 and Figure 10 (impact of factor 3 and impact of all 3 factors) is shown in Figure 11. As can be seen from Figure 11, the difference is mostly close to 0.005% due to factor 1. Only at the extreme ends of current deviation (-0.3% and 0.3%), factor 2 starts to play a role. However, the main influence to overall Uniswap v3 oracle price predictability plays the factor 3: influence of pool fee.

For Uniswap 0.05% fee pool the results are very similar and we omit them as they don't provide any additional value to the study.

IV. Discussion

The results of the analysis of statistical properties of Uniswap v3 price oracle and its deviation from market (Binance) prices show, that Uniswap v3 price oracle tracks Binance price very closely, but only up to the fees percentage of the corresponding Uniswap pool. According to this study, deviation between Uniswap and Binance price for the percentages less than Uniswap pool fees exhibits a very high autocorrelation, meaning that future deviation is highly predictable based on current deviation. The accuracy of prediction depends on the prediction time interval, but remains high even for 20 minutes (100 blocks) time interval, especially for the 0.3% fee pool.

These results have direct security consequences for the protocols using Uniswap v3 price oracles: the Uniswap v3 price oracle is only accurate up to the pool fees percentage; if the protocol requires a higher accuracy, then usage of Uniswap v3 price oracle can introduce statistical vulnerability, which can be abused to slowly drain all funds out of protocol using the statistical properties of the oracle price.

For example, if the DeFi application allows users to trade using future 5-minute Uniswap v3 oracle price from the 0.3% fee pool, but only charges users 0.1% fee, this application will have a statistical vulnerability: a smart user can calculate, that when current deviation between Uniswap and Binance price is 0.3%, then expected future deviation between Uniswap oracle price and average Binance price over the next 5 minutes will be 0.25%, meaning the user can sell 1 ETH in the DeFi application and at the same time buy 0.04 ETH every 12 seconds in Binance. After 5 minutes – the average price the user will buy 1 ETH for will be 0.25% lower than DeFi application sale price. After 0.1% fee, the user's profit over 5 minutes will be 0.15%. Doing many similar operations, the user will be able to quickly extract significant profit from the DeFi application at the expense of the other application users, which is a critical statistical vulnerability.

Another security consequence of the results of this study is behavior of the Uniswap v3 oracle price during the extreme market conditions: as can be seen from the LUNA crash, geometric mean can be less than arithmetic mean by 10% during black swan events, which can have very serious security consequences for any application using such oracle price. Depending on the usage, it can have low impact (for example, lending application calculating account health with a 50% safety net will lose only 10% of its safety net due to this property) all the way to loss of all protocol funds (for example, if the safety net is 5%, or if application allows to trade using such deflated price).

We recommend the following procedures to determine and prevent the statistical vulnerabilities in the applications using Uniswap v3 price oracle:

1. Determine the accuracy of the oracle price required for the application: consider what deviation from market (Binance) price can your application tolerate?
2. If the application requires accuracy less than Uniswap pool's fee, then it's highly likely the application is susceptible to statistical vulnerability during normal market operation.
3. If the application requires accuracy more than Uniswap pool's fee, but less than 10%, then it can be vulnerable during extreme market conditions due to difference of geometric and arithmetic mean.
4. If the application can tolerate the oracle price to be 10% less than market price, then it should not have statistical vulnerabilities due to Uniswap price oracle.

If the application requires accuracy less than Uniswap pool's fee, then the following measures can be taken:

1. Increase the application's fees to the level which will make the possible attack infeasible. The exact fees will depend on application. For example, an application allowing to trade using future 5-minute Uniswap 0.3% fee pool oracle price should have at least 0.25% fee. But if the application allows to trade using triple oracle price, then the fee should also be tripled and to be at least 0.75%.
2. Increase the oracle time interval. For example, increasing time interval from 5 minutes to 20 minutes will reduce the minimum fee required from 0.25% to 0.15%.
3. Use a similar Uniswap pool with a smaller fee (if available). For example, switching from 0.3% Uniswap pool 5-minute oracle to 0.05% pool 5-minute oracle will reduce minimal fee requirement from 0.25% to 0.025%.
4. Use a different oracle price provider. There are cases when Uniswap pool price oracle is not a viable option and the other oracle price provider should be used.

If the application requires less than 10% oracle price accuracy during extreme market conditions, the following measures can be used to prevent statistical vulnerabilities:

1. Take into account current volatility and pause application or significantly increase its fees when extreme market volatility is detected.
2. Use a different oracle price provider.

V. CONCLUSION

In this study we have analyzed the statistical properties of the Uniswap v3 price oracles using the real Uniswap v3 trading data from ETH/USDT 0.3% and 0.05% fee pools in comparison to ETH/USDT trading data from Binance exchange. Factors which influence the Uniswap v3 oracle price were analyzed in isolation and in overall impact based on trading data. A large deviation of geometric mean from arithmetic mean was determined in the extreme market conditions, and a very high autocorrelation of deviation between Binance and Uniswap v3 price oracle was determined in normal market conditions. The security implications of these findings were discussed and the procedures and measures to mitigate the statistical vulnerabilities appearing as a result of Uniswap v3 price oracle usage were offered.

The results of this study can be of practical use to blockchain software developers and smart contract auditors to prevent or mitigate the statistical vulnerabilities which can appear due to usage of Uniswap v3 oracle price and make the applications more secure.

REFERENCES

[1] Fabian Schär, "Decentralized Finance: On Blockchain- and Smart Contract-Based Financial Markets," Federal Reserve Bank of St. Louis Review, Second Quarter 2021, pp. 153-74. https://doi.org/10.20955/r.103.153-74

[2] Mohan, Vijay, Automated Market Makers and Decentralized Exchanges: a DeFi Primer (October 30, 2020). Financial Innovation 8, 20 (2022). https://doi.org/10.1186/s40854-021-00314-5

[3] 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 (AFT '20). Association for Computing Machinery, New York, NY, USA, 80–91. https://doi.org/10.1145/3419614.3423251

[4] Wu, Siwei, et al. DeFiRanger: Detecting Price Manipulation Attacks on DeFi Applications. ArXiv:2104.15068 [Cs], Apr. 2021.

[5] Adams, Austin and Wan, Xin and Zinsmeister, Noah, Uniswap v3 TWAP Oracles in Proof of Stake (October 27, 2022). Available at SSRN: https://ssrn.com/abstract=4384409 or http://dx.doi.org/10.2139/ssrn.4384409

[6] Binance Exchange. Available from: https://www.binance.com

[7] Binance Market Data. Daily Spot ETH/USDT, 1-second interval. Available from: https://data.binance.vision/?prefix=data/spot/daily/klines/ETHUSDT/1s/