Liquidity Providers Taking Position of Volatility: A Comparative Study of AMM Protocols

Anthias Labs

March 7, 2024

Abstract

Automated Market Maker (AMM) protocols are fundamental to the Decentralized Finance (DeFi) ecosystem. This study highlights their distinct approaches to liquidity provision and market making. Through a comprehensive analysis and simulation of prominent AMM protocols such as Uniswap versions 2 and 3, as well as Balancer, we examine their mechanisms, the risks and rewards for liquidity providers (LPs), and the impact of impermanent loss (IL) under different market conditions. A novel aspect of our research is the exploration of LPs' roles beyond mere providers of capital; we propose viewing their participation as strategic positioning against or in favor of market volatility. Our simulations reveal that while traditional 50/50 AMM models expose LPs to symmetrical risks from asset volatility, innovative features in AMMs like Balancer's variable weights and Uniswap v3's concentrated liquidity allow LPs to take nuanced positions on expected market movements. Specifically, we illustrate how a non-50/50 Balancer pool enables LPs to express views on directional volatility, and Uniswap v3's mechanism permits precision in liquidity provision, effectively allowing LPs to bet on specific price ranges. These capabilities introduce a dynamic where LPs can strategically position themselves to mitigate impermanent loss or capitalize on anticipated market trends, thereby elevating their role from passive participants to active managers of risk and return in the face of asset price volatility. Our findings contribute to the broader understanding of the evolving landscape of AMM protocols and offer valuable insights for LPs aiming to optimize their engagement in DeFi markets.

1 Introduction

In the traditional financial market, the trade of securities requires the matching of buyers and sellers, which is usually facilitated by an exchange. Such exchange must be centralized and large-scale in order to provide enough liquidity for both parties of the transaction and it serves the important role of price discovery in the financial markets. However, such centralization might lead to issues such as limited accessibility to retail market participants, high risk of a single point of failure, high transaction fees, and sometimes excessive regulations.

In the world of Web3 and blockchain, the idea of decentralization is a possible solution to the aforementioned questions. The idea of Decentralized Finance (DeFi) was invented to replace traditional financial institutions (TradFi) and to provide an alternative trading exchange besides centralized exchanges such as Binance. Automated Market Makers (AMM), smart contracts that are deployed on chains, allow asset holders to exchange one asset for another without a centralized agent. Unlike dealer markets, any agent can provide liquidity through an AMM, thus there is no barrier to entry for retail traders or asset holders. In addition, the available liquidity is always visible, and prices are automatically set by the protocols' pricing function.

While LPs are providing liquidity, they might also suffer from potential impermanent loss. One important feature of such trade-offs is LPs' positions on volatility. The goal of this project is to evaluate and compare different AMM protocols using simulations with Python in terms of impermanent loss and LPs' positions on liquidity. The paper will include a detailed introduction to several AMM's mechanisms, a description of the simulation method, as well as an analysis of the simulation result.

2 Uniswap v2: Constant Product

Uniswap is one of the first protocols for automated market maker (AMM) using the mechanism of constant product. The central idea of such a mechanism is just a simple equation:

$$x \times y = k$$

k is a constant, and x, y are the reserves of two tokens (Adams et al., 2020). Uniswap smart contracts keep liquidity reserves for ERC20 tokens listed on the exchange, serving the role of market maker. The major difference is that trading through Uniswap does not require a counter-party - trades are executed against these reserves, and prices are set automatically using the constant product mechanism, which helps to maintain a rough equilibrium for prices and reserves of different tokens. Under this simple model, the marginal exchange rate of token A for token B is simply the ratio between the amount of two reserves, $p_t = \frac{r_b^t}{r_a^t}$. When the supply for token A is higher, its value will therefore decrease. Similar to traditional exchanges, the AMM mechanism fulfills the price discovery function since arbitrage traders will correct the price if it deviates from the market price.

The liquidity providers (LP) are the suppliers of the reserves. They earn a portion of the transaction fee by depositing pairs of assets into the liquidity pool and earning a certain amount of liquidity tokens in exchange. The amount of liquidity tokens minted depends on their share of deposits in the liquidity pool.

3 Uniswap v3: Concentrated Liquidity

An important update of Uniswap v3 is the implementation of concentrated liquidity, which will be simulated and analyzed in the next section. In order to increase the depth of the liquidity pool and to facilitate the efficient use of tokens deposited, LPs are allowed to concentrate their liquidity to a smaller price range between two ticks, called a position, based on their own assessment of the market situation. The price of tick i is defined as $p(i) = 1.0001^i$. Each tick is an integer power of 1.0001, which gives the nice property of each tick being a 1 basis point price movement away from its neighboring ticks.

Using the concentrated liquidity model, the new constant product equation is:

$$(x + \frac{L}{\sqrt{P_b}})(y + L\sqrt{P_a}) = L^2$$

P_a and P_b denote the selected price range (Adams et al., 2021). This is similar to the original constant product equation except that the liquidity has boundaries on both sides.

From an LP's perspective, the benefits of depositing tokens depend on the difference between the value of the liquidity tokens plus transaction fees earned and the market price of the underlying assets. The opportunity costs of providing liquidity are defined as impermanent loss (IL), $IL = \frac{V_{\text{pos}} - V_{\text{hold}} + F}{V_{\text{hold}}}$. In order to benefit from providing liquidity, LPs should first have a prediction of the future price and volatility of the selected pool and choose the price range accordingly to maximize their return. Assuming a random process, the future price movement in the short term can be modeled with an Itô process without drift. Another important factor to consider is the overall liquidity distribution of the whole pool. In addition, LPs should promptly update their predictions about the market to actively adjust their price range, which requires a certain level of sophistication. This is analogous to active asset managers, whose ability to generate extra alpha is constantly being studied.

For most of the pools, IL generally surpassed the fees earned during the same period, which means an average LP of Uniswap v3 is financially harmed by providing liquidity, despite their choice of position, compared to simply holding the tokens (Loesch et al., 2021). This finding will hinder the entry of retail traders as LPs, resulting in less liquidity available in the ecosystem. Nonetheless, Uniswap v2 and v3 remain leading AMMs in terms of total value locked (TVL), which is the total liquidity available.

4 Balancer: Constant Mean

The Balancer Protocol is defined by a constant function of the pool's balances and weights (Martinelli and Mushegian, 2019). LPs can deposit up to eight different cryptocurrencies into Balancer pools at any value ratios they want. In the following value function, B_t is the balance of token t, and W_t is the weight of token t as a share of the total portfolio value. The product of each token's balance raised to the power of its corresponding weight is set to be equal to a constant V:

$$V = \prod_t B_t^{W_t}$$

The spot price between two tokens in an infinitesimal small trade can be calculated by the ratio between the balance and weight of the token going into the pool divided by the ratio between the balance and weight of the often going out:

$$\text{Spot Price}_i^o = \frac{B_i / W_i}{B_o / W_o}$$

Since the weights of tokens are held constant, the share of portfolio value for each token remains the same, and price changes only occur due to changes in the balance for each token. Whereas an index fund of tokens may require management fees, Balancer offers transaction fees to LPs as traders rebalance the pool through swaps.

5 Primitive: Dynamic Pricing Function

While constant function market makers (CFMM) provide some basic framework for AMMs, dynamic function market makers allow for more flexibility in the design of AMMs and more complexity in designing trading functions.

Primitive is an example of protocols that adopt a dynamic pricing function. Its AMM is called Replicating Market Maker (RMM-01), which approximates the payoff of a covered call to LPs through the aggregation of swap fees and changes over time. It follows the Black-Scholes option pricing method and sets K, σ, τ at the beginning of pool creation. Denote R_1, R_2 as the amount of reserves for two types of tokens. The trading function is $\varphi(R_1, R_2) = R_2 - K\Phi(\Phi^{-1}(1 - R_1) - \sigma\sqrt{\tau}$, where Φ is the CDF of a standard normal distribution and σ, τ, K are the volatility, time to expiry, and strike price, similar to an option (Angel & Czernik, 2021). Taking the derivative of the pricing function with respect to R_1, we can abstain the marginal price of asset 1 $S_1 = Ke^{\Phi^{-1}(1-x)\sigma\sqrt{\tau}-\frac{1}{2}\sigma^2\tau}$. All pools have these parameters available, which makes it possible to perform on-chain calculations.

Similar to other protocols, Primitives provides LPs with RMM-01 tokens (LPT). From the perspective of LPs, they could adopt a covered-call strategy to collect some premiums, which are in the form of transaction fees, or get some downside protection. According to Primitive's whitepaper, such strategies give them the opportunity to eliminate a part of impermanent loss and flexibility to adjust based on their risk preferences and prediction of the market volatility (Angel & Czernik, 2021).

Since the pricing function is dynamic with respect to time and volatility, arbitrage opportunities will almost surely arise, attracting arbitrageurs to participate in the process of aligning the price with external marketplaces while contributing to LP's revenues via swap fees. However, there exists a gap between the theoretical and actual payoffs in practice, which might be explained by the imperfect arbitrage conditions or the existence of noise traders who don't follow rational trading strategies. The assumption of constant implied volatility and a Gaussian distribution of returns makes RMM fragile for high-volatility assets too. Nonetheless, RMM-01 helps to pave the way for future implementation of on-chain derivatives using AMMs.

6 Catalyst: Units of liquidity

Catalyst is known as the Unit of Liquidity. One of its important features is its easy deployment on any new chain, allowing users of Catalyst easy access to assets across different chains. It is generally very difficult to perform cross-chain transactions, due to different architectures, smart contract standards, regulations, etc. However, Catalyst achieves cross-chain swaps by generating a unit of liquidity as a certificate of redemption that will be sent and validated on another chain. This improvement greatly increased the number of swap options for traders, the overall liquidity availability, and the possible diversification of portfolios.

Catalyst follows a model that is very similar to the constant product AMM. To facilitate cross-chain transactions, each asset on a different chain will have a separate price curve, $P(x)$, which must be decreasing and non-negative with respect to the amount of reserves of that asset. When an amount of Δx is deposited or withdrawn, the liquidity will move along the price curve. The area under the price curve $P(x)$ between x and $x + \Delta x$, is the change of liquidity on this chain, defined as a Unit of Liquidity (Lindgren & Sanmiguel, 2023).

$$U = \int_{x}^{x+\Delta x} P(w)dw$$

This Unit of Liquidity will be sent through a validated cross-chain message protocol to perform the opposite transaction on the other asset. We then find the amount of change Δy that makes the area of change equal to the Unit of Liquidity so that the overall liquidity remains constant.

A simple design of the price function is $P(w) = \frac{W}{w}$, where W is the relative weight to adjust the distribution of liquidity of different assets in the pool of assets. The total liquidity of the pool should remain constant after transactions, and the invariant is the sum of the integrals of the price functions for every asset. $K = \sum_{i=0}^{n} ln(w_i)W_i$. Taking the exponential of K, we can replicate the constant mean model of Balancer, showing that Catalyst is a generalization of different basic models of AMMs.

7 Swaap: Offchain Calculations

Swaap is an AMM that utilizes off-chain calculations to provide prices for traders, compared to other AMMs that only use on-chain data. This is defined as an off-chain pricing AMM. Swaap v2 incorporates external price sources, such as prices on centralized exchanges, into its quotation function, which allows them to access such information with very short latency.

The quotation module of Swaap aims to maximize LPs' PnL against a given benchmark, taking into consideration random factors such as risk profile, market movement, and volatility (Swaap Labs, 2023). It adopts Bergault et al.'s method to integrate exogenous centralized market exchange rates, such as Binance, Kraken, and Coinbase, as indicative signals for pricing (Bergault et al., 2023). This module is maintained off-chain, hoping that its dynamic nature will swiftly capture the up-to-date signals in market conditions to provide a more favorable quote. The market exchange rate movement is assumed to follow the Itô random process with drift and a Brownian motion. Swaap compares the return of providing liquidity to the benchmark of simply holding the asset and tries to maximize LP's benefits within a certain markup range. Following this model, The expected return of LPs for providing liquidity is a lot higher compared to Uniswap v2's average negative return.

However, Swaap still remains relatively small-scaled in terms of TVL, trading volume, and the number of LPs. The reliance on quotes in centralized exchanges does not completely align with the notion of DeFi and automated market making, whose goal is to provide quotes without an order book or a central organization. This might be most beneficial for retail LPs to earn extra returns that they could not have earned in a centralized market because of the fundamental design of Swaap's model.

8 UniswapX: Dutch Auctions

The UniswapX Protocol is characterized by Dutch auction-based decentralized trading (Adams et al., 2023). After traders create and sign an order, onchain agents such as market makers and MEV (maximal extractable value) searchers will enter a Dutch auction and compete for the right to fill the order. The order's offer increases within a specific range centered around the current market price, and fillers are incentivized to fill an order as soon as it allows for a profit. The order is then sent to and settled onchain using a reactor contract.

Fillers can lower the gas fee on transactions by batching orders. Batched orders also help prevent sandwich attacks (Canidio and Fritsch, 2023), which occur when trading bots profit off of transactions by placing buy and sell orders around a target order in quick succession, since multiple orders within a given period are simultaneously settled at a similar price.

9 Simulation Results

We have simulated Uniswap v2, v3, and Balancer with a focus on ETH and USDC pairing, and the results of the simulations are shown below. Each protocol's python class can handle an arbitrary `x` and `y` token. For each protocol, we instantiated the pool's liquidity with 3000 x tokens to simulate USDC and 1 `y` token to simulate ETH. Instantiating the pools as such meant that the starting price of 1 ETH is 3000 USDC. To get a price curve, we called the `swap_x_for_y` function on the classes, with a delta of 1. This function simulated iteratively trading 1 USDC for ETH, which allowed us to derive a price curve for each pool.

9.1 Uniswap v2

Figure 1 shows an example of the Uniswap v2 Price Curve, with the constant product model of $x \times y = 3000$. When trades are made, the amount of tokens will move along this curve to reach a new equilibrium where each asset makes up 50% of the value of the pool. The line pairs the quote of the `y` token in terms of the `x` token and vice versa. For example, at the start of the line, 1 `y` token, meant to be ETH, is worth 3000 x tokens, meant to be USDC. After 3000 more `x` tokens have been added to the liquidity pool, which is when there are

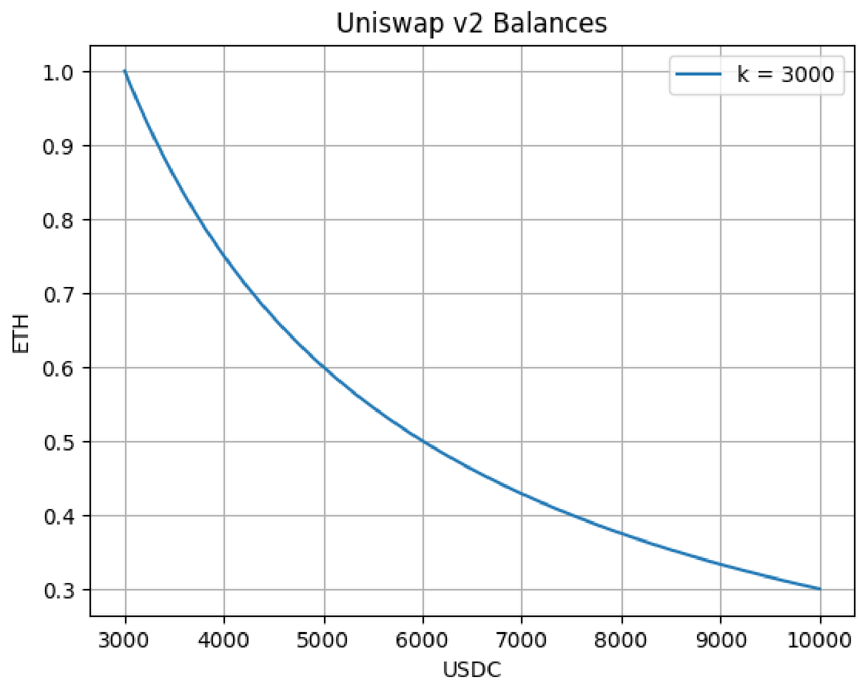

Figure 1: Uniswap v2 Price Curve

6000 `x` tokens, 0.5 `y` tokens are then worth 6000 `x` tokens.

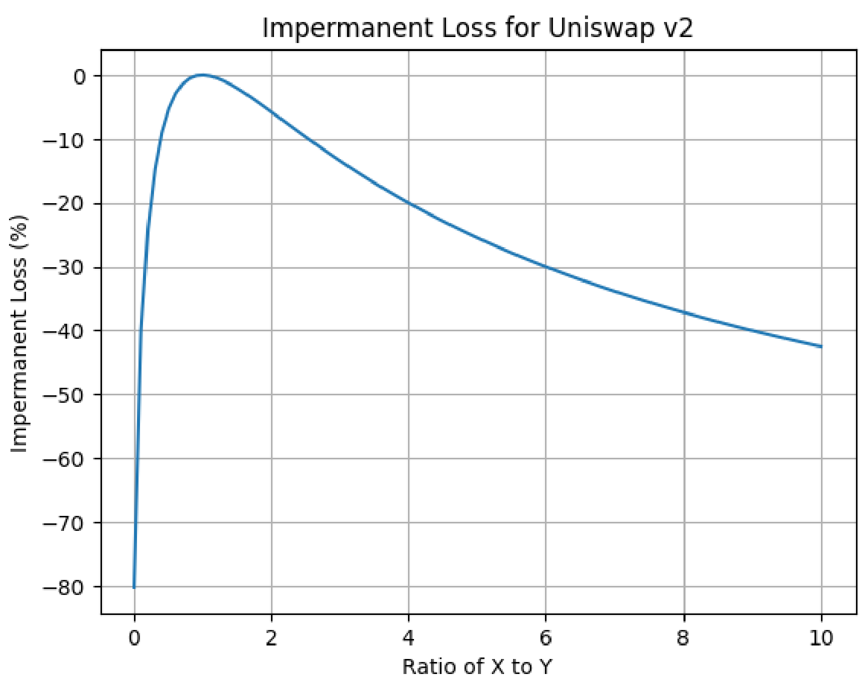

Figure 2: Uniswap v2 Impermanent Loss

Figure 2 shows the relationship between the amount of impermanent loss and the ratios of two tokens. The methodology for calculating impermanent loss (IL) in our simulations for both Uniswap v2 and Balancer protocols is grounded in the concept of price ratio change between assets in a liquidity pool. Impermanent loss occurs due to the divergence in the value of assets when compared to holding the assets outside the pool. Our calculation aims to quantify this loss as the price ratio of the assets changes.

For Uniswap v2, the impermanent loss is calculated based on the change in the price ratio of two assets (X and Y) in the pool. The price ratio change is defined as the new price ratio divided by the initial price ratio, which is typically 1:1 at the time of initial pool creation or at any point chosen as a baseline for comparison. The impermanent loss formula applied is:

$$IL = 2 \times \left(\sqrt{\frac{new\_price\_ratio}{initial\_price\_ratio}} \right) / \left(1 + \frac{new\_price\_ratio}{initial\_price\_ratio} \right) - 1$$

This formula captures the relative loss compared to holding the assets, expressed as a percentage. Our simulations covered a range of price ratios from 0.01 to 10 to illustrate how the impermanent loss varies as the price ratio between X and Y changes.

9.2 Balancer

Figure 3 shows the price curves of a two-asset Balancer pool of USDC and ETH, with different weights of the two assets. The simulation varied the weights of USDC and ETH in the pool, ranging from 10% USDC and 90% ETH to 90% USDC and 10% ETH, to analyze how varying asset weights within a pool affects the asset's exchange rates and the impermanent loss experienced by LPs. This was accomplished by iteratively performing swaps of 1 USDC for ETH.

The results reveal how Balancer's flexible weight mechanism influences the pool's behavior. A pool with a higher weight on USDC demonstrates a steeper exchange rate curve as USDC is swapped for ETH, as that a pool's sensitivity to price changes is directly correlated with the weight distribution between its assets. This flexibility allows liquidity providers to customize their risk and return profile, potentially reducing impermanent loss compared to a standard 50/50 weight distribution seen in Uniswap V2.

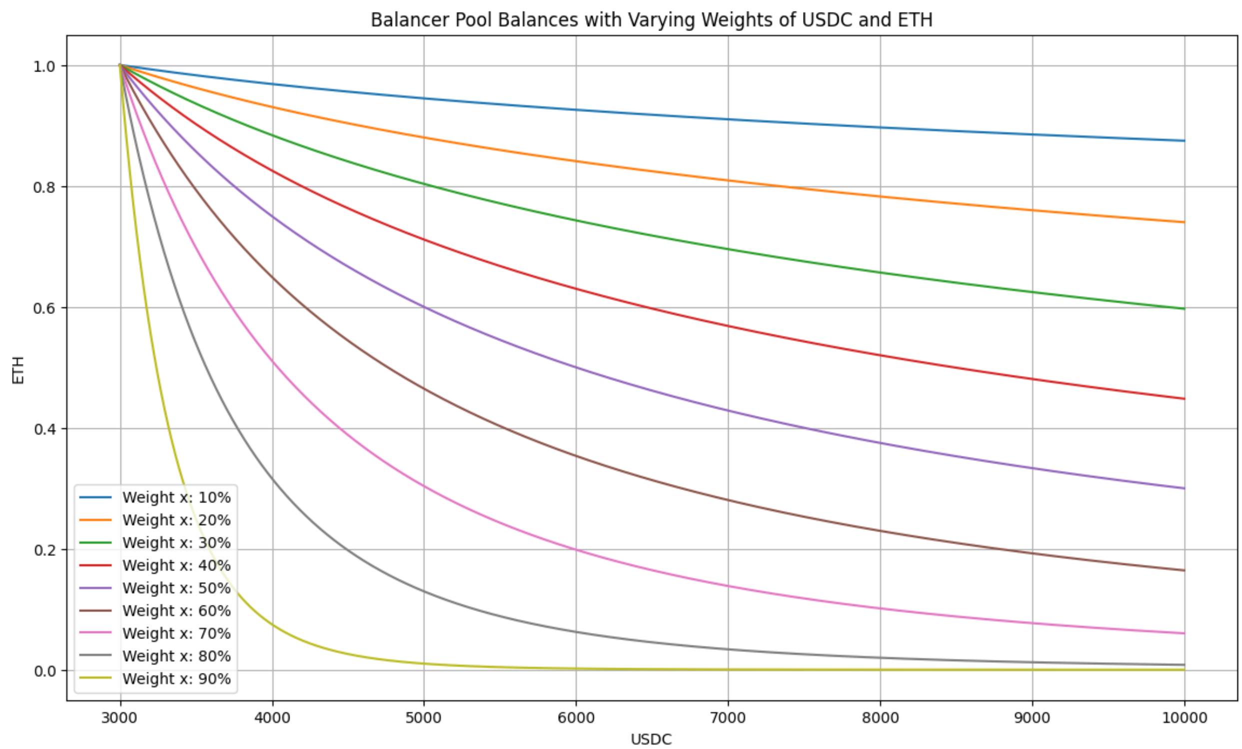

Figure 3: Balancer Price Curves

For the Balancer protocol, the impermanent loss methodology adjusts to account for the weighted nature of the pools. The initial price ratio is determined based on the weights assigned to the assets in the pool. Impermanent loss is then calculated for a variety of weight configurations and price ratios, providing insights into how different weight allocations influence the susceptibility to impermanent loss. Figures 5 and 6 display impermanent loss for the Balancer protocol at different pool weights in linear and logarithmic scales respectively.

9.3 Uniswap v3

When simulating the concentrated liquidity, Uniswap V3, we have carefully made some design choices to take into consideration computational resources and the accuracy of the data.

- First, we assumed a fixed fee rate of 0.03%, whereas Uniswap v3 allows LPs to freely choose a fee rate. We argue that it is not harmful to the result of the analysis, and

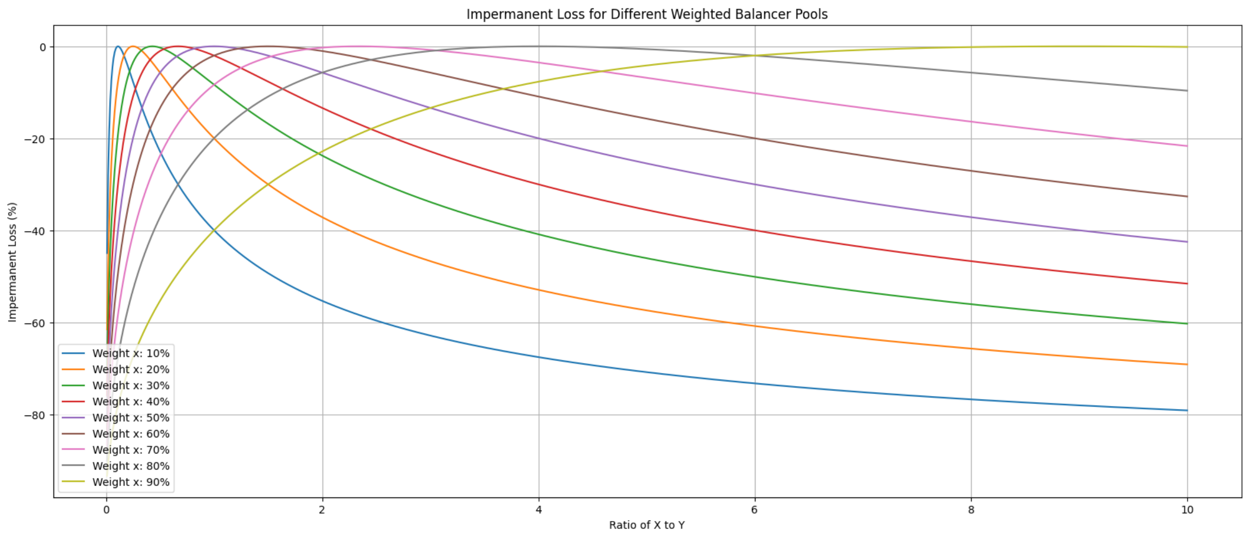

Figure 4: Balancer Impermanent Loss

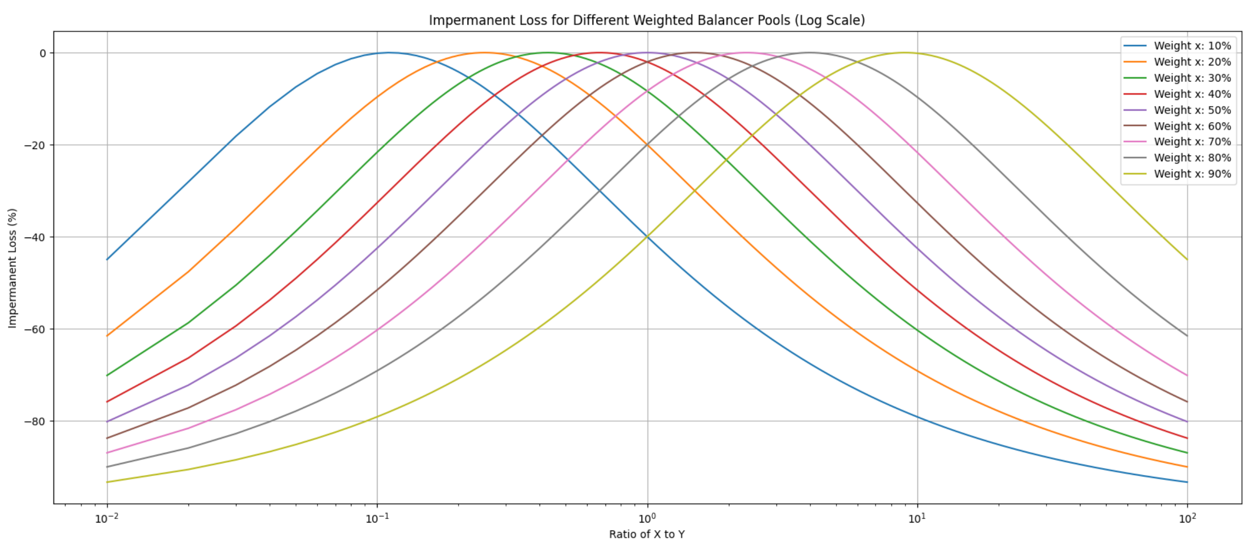

Figure 5: Balancer Impermanent Loss (Log Scale)

it also makes the comparison more reasonable by excluding the effect of fee structure when evaluating different AMM mechanisms.

- Second, instead of the 3000 to 1 ratio used in the previous graphs, we use a 3800 to 1 ratio to more closely capture the current price of ETH.

- Third, when generating the initial distribution, we mimic the decision-making process of an actual LP on Uniswap of providing liquidity: they will have a low price a and high price b based on their own judgment, and they will provide liquidity just in the range $[a, b]$. To approximate this, we assume that the low price and high price each follow a normal distribution with a mean of the 2-week low and a price that is 10% higher than the current price. Then we draw values from the two distributions for a fixed number of LPs, each providing a fixed amount of liquidity. Then we aggregate them to get an approximation of the initial liquidity distribution. The liquidity distribution is shown in Figure 6 below. Note that the x-axis is tick indices, and the corresponding price can be calculated by $P(t_i) = 1.001^i$. The central point of the distribution is around \$3811.
- Last but not least, Uniswap v3 utilizes intervals of integer powers of 1.0001 to capture 0.01% changes in price. In our simulations, we used integer powers of 1.001 instead to reduce runtime.

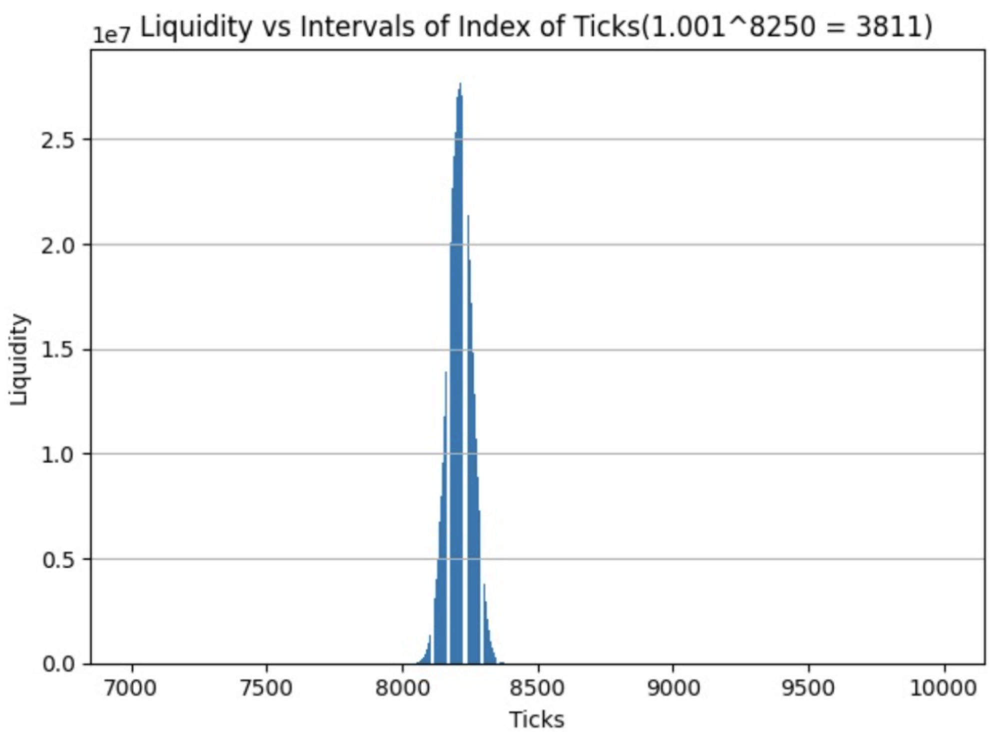

Figure 6: Uniswap v3 Liquidity Distribution

To compare the potential slippage after the same sequence of transactions on Uniswap v2 and v3, two experiments are conducted. In the first experiment, the total value locked (TVL) is fixed while the quantities of trades vary.

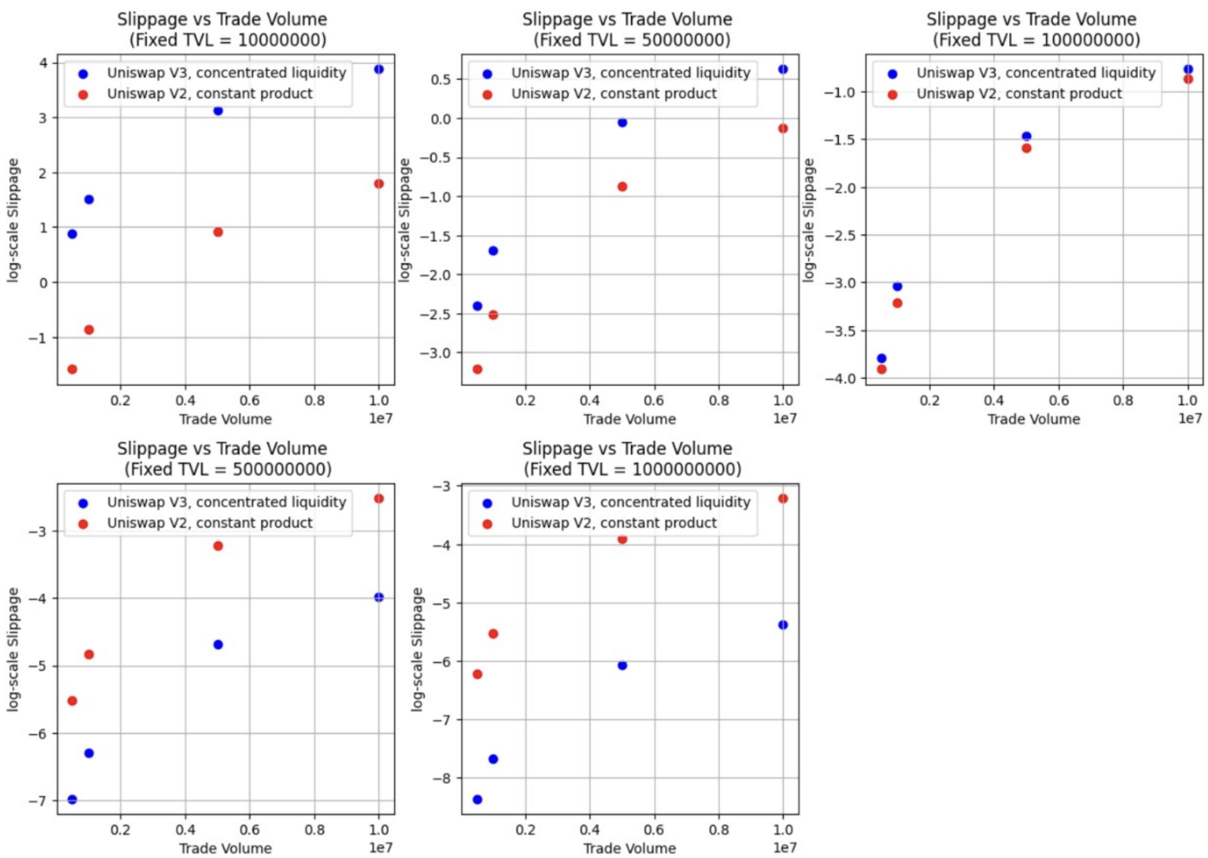

Figure 7: Uniswap v2/3 Slippage

In the simulations, TVL will gradually increase to 1×10^9, is $\frac{1}{3}$ of the actual TVL (Uniswap.com, 2024). Based on the result shown in Figure 7, it is clear that when TVL becomes higher than a certain value (1×10^8 in this experiment), Uniswap v3 will do significantly better than v2 in terms of slippage. When initializing the liquidity distribution for v3, the average liquidity provided per LP is quite high in order to save some computational efforts, which results in some intervals having less liquidity than they would actually have. When such liquidity is depleted, interval crossings will happen. As discussed before, our model captures a price change of 0.1% change instead of 0.01%. The price changes during interval swapping happen in fact 10 times more frequently than the actual Uniswap model.

The second experiment assumes that there are 10 random trades in random directions happening every day and explores where the price will be in half a year. TVL is a fixed large number for both models. Based on the result shown in Figure 8, Uniswap v3 does better than v2 in terms of slippage.

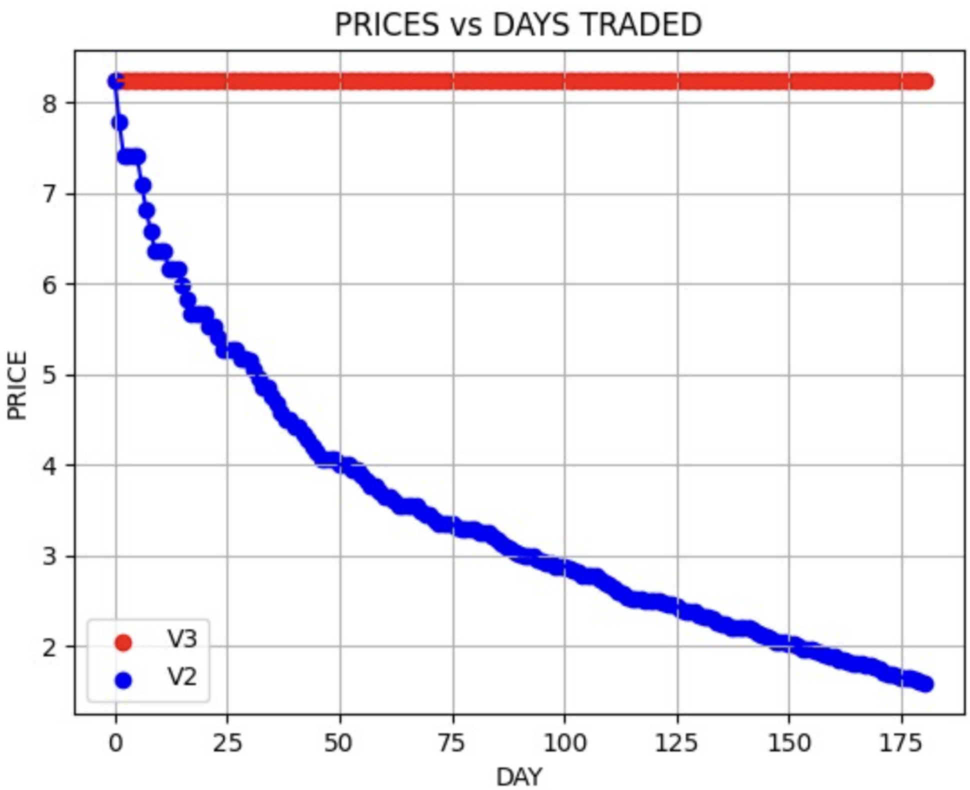

Figure 8: Uniswap v2/3 Slippage after 1/2 year

9.4 Impermanent Loss: A Deeper Dive

Impermanent loss (IL) represents a fundamental challenge and area of concern for liquidity providers (LPs) in the Automated Market Maker (AMM) ecosystem. As LPs deposit their tokens into liquidity pools, they become susceptible to IL, which occurs when the price ratio of deposited assets changes from the time of deposit. This change can lead to a situation where the dollar value of the LP's share in the pool is less than if they had simply held their assets outside the pool, despite earning transaction fees.

The issue of IL is particularly pronounced in volatile markets where asset prices can fluctuate widely. Protocols like Uniswap have shown that IL can significantly impact LP returns, as highlighted by Heimbach et al. (2022). Their analysis indicates that the profitability of liquidity provision is heavily dependent on the LP's ability to anticipate market movements and select appropriate fee tiers and price ranges for their liquidity. This highlights the importance of financial savvy and market insight in managing IL risk.

Moreover, the landscape of AMM protocols has evolved with mechanisms designed to mitigate the effects of IL. Beyond the fee earnings, some protocols have introduced features that provide LPs with more tools to manage their exposure to IL. For example, Swaap's quotation module is designed to optimize LP returns by adjusting liquidity provision strategies based on market conditions. This approach signifies a shift towards more dynamic and responsive AMM models that attempt to safeguard LP interests against volatile market movements.

Bancor's approach to IL mitigation represents a significant advancement in protocol-level protections for LPs. By establishing an insurance mechanism funded by protocol earnings, Bancor v2 aims to fully compensate LPs for IL experienced within their pools, subject to a vesting period. This model not only addresses the IL challenge head-on but also serves to boost LP confidence and encourage deeper liquidity provision by reducing the perceived risk of participation.

The concept of IL and the various strategies employed by AMM protocols to mitigate its impact underscore a critical aspect of liquidity provision in DeFi. LPs must navigate the trade-offs between potential returns from fee earnings and the risk of IL. The emergence of sophisticated AMM designs and protective measures is indicative of the evolving DeFi space, striving to balance the rewards of liquidity provision with the risks inherent in volatile cryptocurrency markets.

As the DeFi ecosystem continues to mature, the development of more advanced mechanisms to manage IL will be crucial in attracting and retaining LPs. This will not only enhance the stability and efficiency of AMM protocols but also contribute to the broader adoption and growth of DeFi as a viable alternative to traditional financial markets.

10 Liquidity Provision as a Short Position on Volatility

In the evolving landscape of Decentralized Finance (DeFi), Automated Market Makers (AMMs) have introduced innovative liquidity provision mechanisms, fundamentally altering the way liquidity providers (LPs) interact with market volatility. This section delves into the strategic positioning of LPs in relation to market volatility.

Providing liquidity in AMM protocols can be analogized to holding a short position on asset volatility. LPs primarily earn transaction fees as compensation for the risk of impermanent loss (IL), which escalates with increased asset price volatility. In a traditional 50/50 AMM model, such as Uniswap V2 or a balanced Balancer pool, LPs are equally exposed to IL regardless of the direction in which the asset prices move, provided there is significant volatility.

10.1 Balancer: Directional Volatility Positioning

Balancer allows LPs to custom-tailor their exposure to assets within a pool through flexible weight configurations. This unique feature empowers LPs to take positions on directional volatility. Consider a pool with 20% USDC and 80% ETH; in this scenario, the LP is expressing a stronger position on the volatility of ETH relative to USDC. Should ETH experience high volatility, the impermanent loss would be magnified due to the disproportionate weight. Conversely, this setup could potentially offer higher returns during periods of favorable market movements for ETH, showcasing Balancer's ability to allow LPs to manage their risk-return profile based on their market outlook.

10.2 Uniswap V3: Precision in Volatility Positioning

Uniswap V3 further refines the concept of volatility positions by introducing concentrated liquidity, allowing LPs to specify the price ranges in which they wish to provide liquidity. This mechanism enables LPs to take nuanced positions on expected asset price movements, essentially making a bet on the future price corridor of the assets. By concentrating liquidity within certain price ranges, LPs can mitigate exposure to impermanent loss outside these ranges but risk earning no fees if prices move beyond their specified bounds.

For example, an LP who anticipates that the price of ETH will oscillate within a specific range can allocate liquidity to capture fees from trades within this range, minimizing exposure to IL from drastic price movements outside this corridor. This approach allows for more strategic positioning against volatility, offering LPs the ability to customize their exposure based on their market predictions and risk tolerance.

10.3 Conclusion on Liquidity Positions and Volatility

The inherent nature of liquidity provision in AMMs posits LPs in a unique stance with regard to market volatility. While traditional 50/50 pools expose LPs to symmetrical risks of IL from volatility, platforms like Balancer and Uniswap V3 offer innovative tools for LPs to sculpt their exposure more deliberately. Balancer's variable weight pools allow for directional bets on asset volatility, whereas Uniswap V3's concentrated liquidity feature affords LPs unprecedented precision in targeting specific price ranges, thus tailoring their volatility exposure.

These advancements underscore a pivotal evolution in DeFi, highlighting the increasing sophistication with which LPs can navigate and position themselves within the volatile cryptocurrency markets. As AMMs continue to evolve, understanding and leveraging these mechanisms will become crucial for LPs aiming to optimize their returns while managing the risks associated with asset price volatility. This perspective challenges the conventional view of LPs' roles, positioning them not merely as passive participants but as active strategists capable of influencing their exposure to market movements.

11 Conclusion

In this study, we delved into the complexities of Automated Market Maker (AMM) protocols such as Uniswap versions 2 and 3, and Balancer, examining their operational mechanisms, the impact of impermanent loss (IL) on liquidity providers (LPs), and the strategic opportunities available to manage market volatility. Our simulations highlight how traditional AMM models expose LPs to symmetrical risks from asset price fluctuations, while innovative features in Balancer and Uniswap v3 empower LPs to tailor their market exposure and potentially mitigate IL. Specifically, Balancer's flexible weight configurations and Uniswap v3's concentrated liquidity offer nuanced strategies for LPs to position themselves favorably in volatile markets.

These findings suggest a significant evolution in the role of LPs within the DeFi ecosystem, from passive participants to active strategists capable of optimizing their risk and return profiles. This transition underscores the importance of continued innovation in AMM designs to enhance market efficiency and the resilience of the DeFi space. As the sector matures, the interplay between technological advancements and strategic liquidity provision will be crucial in shaping its future, promising a new era of financial democratization and efficiency in decentralized markets.

References

[1] Adams, H., Williams, E., Pote, W., Yang, Z., Zinsmeister, N., Wan, X., Lin, A., Campbell, R., Robinson, D., Toda, M., Leibowitz, M., Zhong, E., & Karys, A. (2023). *UniswapX.* https://uniswap.org/whitepaper-uniswapx.pdf

[2] Adams, H., Zinsmeister, N., & Robinson, D. (2020). *Uniswap v2 Core.* https://uniswap.org/whitepaper.pdf

[3] Adams, H., Zinsmeister, N., Salem, M., Keefer, R., & Robinson, D. (2021). *Uniswap v3 Core.* https://uniswap.org/whitepaper-v3.pdf

[4] Angel, A., Czernik, M., 0xEstelle, & experience. (2021). *Primitive RMM-01.* https://www.primitive.xyz/papers/Whitepaper.pdf

[5] Angeris, G., Chitra, T., Diamandis, T., Evans, A., & Kulkarni, K. (2023). *The Geometry of Constant Function Market Makers.* https://arxiv.org/pdf/2308.08066.pdf

[6] Bergault, P., Bertucci, L., Bouba, D., & Guéant, O. (2023). Automated market makers: mean-variance analysis of LPs payoffs and design of pricing functions. *Digital Finance.* https://doi.org/10.1007/s42521-023-00101-0

[7] Bouba, D. (2021). *Swaap.finance: Introducing the Matrix-MM.* https://www.swaap.finance/whitepaper.pdf

[8] Canidio, A., & Fritsch, R. (2023). *Arbitrageurs' profits, LVR, and sandwich attacks: batch trading as an AMM design response.* https://arxiv.org/pdf/2307.02074v3.pdf

[9] Heimbach, L., Zürich, E., Ch, S., Schertenleib, E., Switzerland, Z., & Wattenhofer, R. (2022). *Risks and Returns of Uniswap V3 Liquidity Providers.* https://arxiv.org/pdf/2205.08904.pdf

[10] Lindgren, A., & Sanmiguel, J. (2023). *Catalyst - Asynchronous Autonomous Market Making via a Unit of Liquidity.* https://catalystdao.github.io/papers/Catalyst%20-%20Asynchronous%20Autonomous%20Market%20Making%20via%20a%20Unit%20of%20Liquidity.pdf

[11] Loesch, S., Hindman, N., Welch, N., & Richardson, M. (2021). *Impermanent Loss in Uniswap v3*. `https://arxiv.org/pdf/2111.09192.pdf`

[12] Martinelli, F., & Mushegian, N. (2019). *Balancer Whitepaper*. `https://balancer.fi/whitepaper.pdf`

[13] Milionis, J., Moallemi, C. C., Roughgarden, T., & Zhang, A. L. (2023). *Automated Market Making and Loss-Versus-Rebalancing*. `https://arxiv.org/pdf/2208.06046.pdf`

[14] Mohan, V. (2022). Automated market makers and decentralized exchanges: a DeFi primer. *Financial Innovation, 8*(1). `https://doi.org/10.1186/s40854-021-00314-5`

[15] Swaap Labs. (2023). *Swaap v2: Optimal liquidity infrastructure*. `https://www.swaap.finance/v2-whitepaper.pdf`