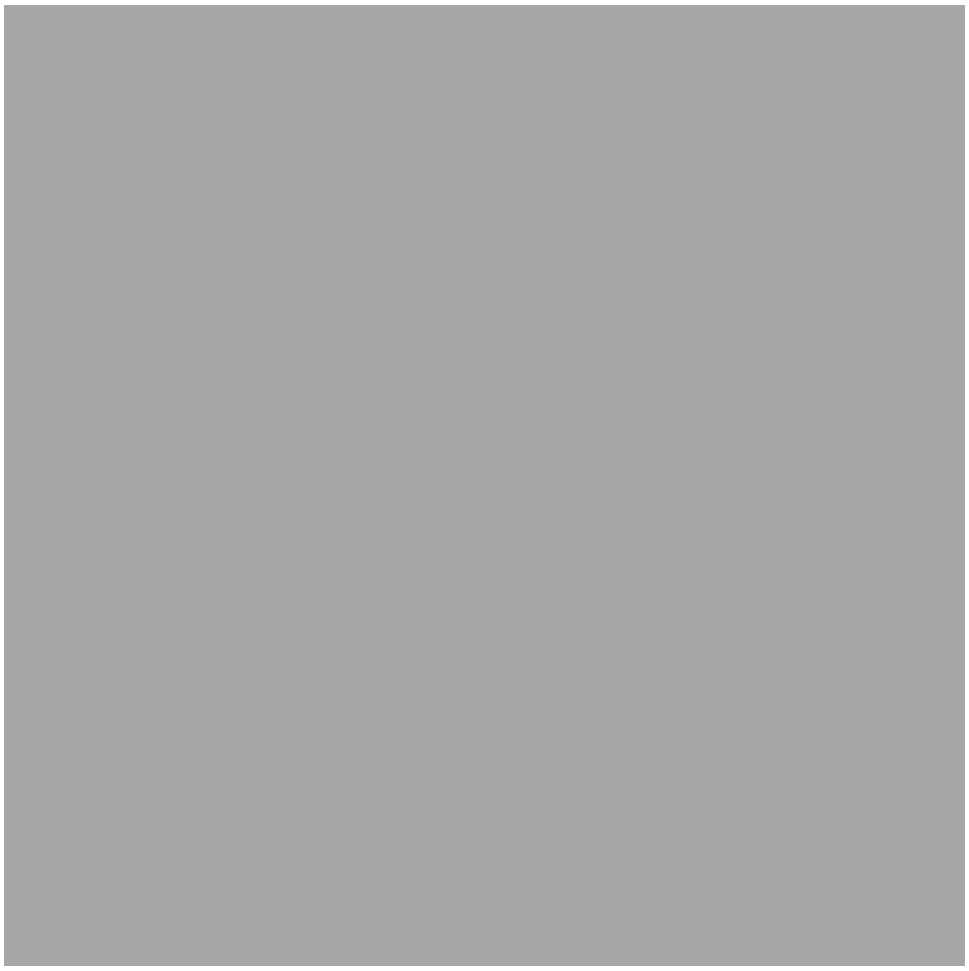

Working Paper No 1-S

NES Student

Paper series

LIQUIDITY PROVIDER POSITION ANALYSIS AND PRICING IN AUTOMATED MARKET MAKING SYSTEMS

Grigoriy Korolev

March 2026

Liquidity Provider Position Analysis and Pricing in Automated Market Making Systems

Grigoriy Korolev*

Abstract

Decentralized Finance (DeFi) has Automated Market Making (AMM) as a dominant. It allows decentralized exchanges to operate via liquidity pools. Crypto assets are traded in automatic way without any control by market makers with two types of agents: traders and liquidity providers (LPs). This paper develops a comprehensive framework for analyzing and pricing LP positions in protocol Uniswap. We demonstrate that an LP's position exhibits option-like characteristics and derive closed-form pricing formulas. Our empirical analysis of ETH-USDC returns reveals heavy-tailed distributions with significant asymmetry, indicating greater downside risk. Importantly, we show that concentrated liquidity in Uniswap V3 generates substantially higher returns compared to uniform liquidity in V2. The pricing models presented enable precise valuation of these novel financial instruments, allowing them to launch on the market.

Keywords: Automated market making, Decentralized Finance, cryptocurrency, asset pricing, liquidity provider, heavy tails.

*I would like to thank Sergey Kovbasyuk (New Economic School) and Vyacheslav Gorovoy (New Economic School) for their extraordinary support in this research project and general supervising. Also, I am grateful for the support provided by the New Economic School throughout the research process. The facilities and resources offered were extremely helpful and have been instrumental in completing this research project successfully.
Grigoriy Korolev, New Economic School, Russia. Email: gkorolev@nes.ru. ORCID: 0000-0003-1611-1277.

1 Introduction

Decentralized Finance (DeFi) has emerged as a disruptive innovation within the financial industry. It offers solutions that challenge traditional centralized systems, or Centralized Finance (CeFi). In traditional finance, centralized exchanges (CEX) act as trusted intermediaries. They use an order book, which is a list of buy and sell orders, to match traders and determine prices based on supply and demand. This market microstructure has been the standard for decades.

In particular, now transactions can be made without the control of a central counterparty with DeFi. Users can execute financial transactions without any side-monitoring. This shift is not just technological but also economic. Indeed, it changes how markets are structured and how liquidity is provided.

And one of the most transformative developments in the DeFi area is Automated Market Making (AMM). AMM-like protocols have novel liquidity provision. In this way, it enables smooth and efficient trading without the need for intermediaries.

AMM refers to the use of smart contracts and algorithmic trading to create and maintain liquidity pools within decentralized exchanges (DEX). Unlike traditional order book-based exchanges, AMM protocols rely on predetermined mathematical formulas and liquidity reserves to facilitate trades. This approach eliminates the need for centralized exchanges and allows anyone to become a liquidity provider by providing (depositing) their assets into these pools.

The majority of the implemented AMMs take the shape of a constant function market maker (CFMM), exemplified by platforms like Uniswap (Adams, 2020 [1] and 2021 [2]). The name CFMM stems from the constant product between the number of two assets in the liquidity pool: $X \times Y = const$. This mechanism has been well-established as a foundational DeFi primitive by Mohan, 2022 [4].

There are plenty of benefits of AMM. Just in case, it enables continuous liquidity. That is, it ensures that traders can execute transactions at any time without relying on counterparties. This 24/7 liquidity provision is a significant departure from traditional markets. In CeFi liquidity may be limited during certain hours or subject to intermediaries' problems.

Our analysis centres on Uniswap. Uniswap was launched in 2018 on the Ethereum blockchain. It has emerged as the dominant DEX protocol with around $4 billion in total value locked as of 2025 [3].

There exist only two types of agents who interact directly with the core protocol: traders and liquidity providers (LPs).

- Traders are the users of the exchange. Their role is simple: they come to the platform to swap one token for another (e.g., ETH for USDC). Of course, they pay a small fee for this service, which is added to the pool's reserves. These traders resemble classical traders in CEX (NASDAQ or Moscow Exchange).

- Liquidity Providers are the foundational actors that make trading possible. Instead of placing individual orders like in a traditional exchange, they collectively provide the capital that forms the liquidity pool. The most popular scheme is the following. They do this by depositing an equivalent value of two tokens into a pool (e.g., $100 of ETH and $100 of USDC). In return, they receive "LP tokens," which represent their share of the total pool. This type of agent is like a major market maker on a stock exchange.

This structure is slightly different from a centralized exchange, as illustrated in Figure 1. In CEX, liquidity comes from many individuals posting independent buy and sell orders into an order book. In a DEX like Uniswap, liquidity is pooled together into a single, automated reservoir that everyone trades against.

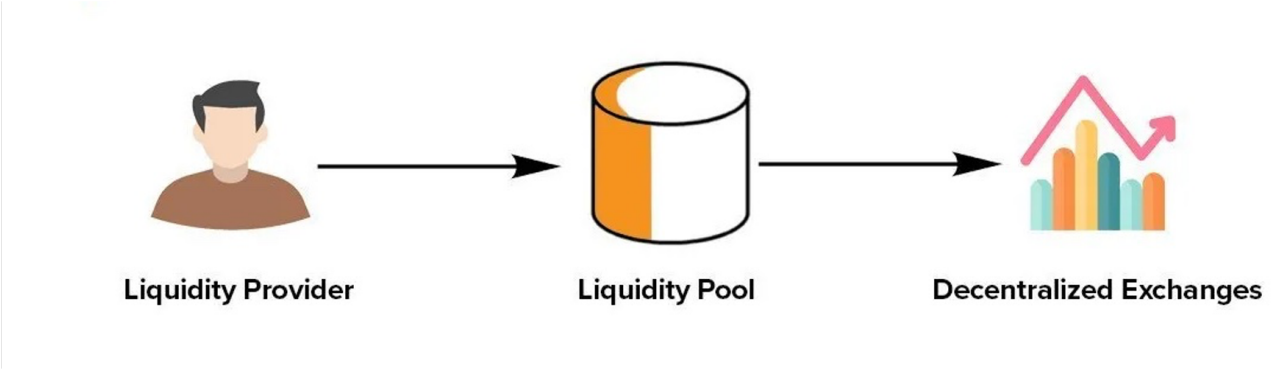

Figure 1: Simplified architecture of a DEX. Traders interact with a single, shared liquidity pool funded by LPs.

From a financial econometrics perspective, it is crucial to analyze the returns and risks for LPs in this new system. Their role is entirely passive yet exposed to significant financial engineering risks due to holding a position (will be shown later). Unlike a traditional bank deposit that earns a fixed interest rate, LP's earnings are highly variable and are driven by two main factors.

Firstly, LPs earn a small percentage of every trade that happens in their pool (trading fees). The more trading volume, the more fees they accumulate. This is their source of profit. The web interface available for LPs with current fee APR (Annual Percentage Rate) can be found in Appendix (Figure A1 and A2).

Secondly, Impermanent Loss (IL) is the critical risk. It is an unrealized loss that occurs when the market price of the deposited assets changes compared to when they were deposited. This loss is directly driven by the volatility of the pool's assets. The more the prices of the two tokens diverge, the greater the IL. Essentially, an LP's position behaves like a short volatility position. They profit from high trading volume in a stable market but can lose money relative to simply holding the assets if the prices are volatile. Therefore, the profitability of being an LP is a complex trade-off between earning fees and managing the risk of IL caused by asset volatility.

The rapid evolution of AMM protocols has accelerated research in this field. Recent studies by Lehar and Parlour, 2025 [5] formally model the microstructural aspects of Uniswap. Vlasov, 2025 [6] provide new insights into fee structures and IL. Concurrently, innovative approaches like the predictive AMM architecture using deep reinforcement learning demonstrate the dynamic nature of this research area. That is shown in Lim, 2024 [7]. This paper contributes to this fast-moving discourse by developing a comprehensive framework for analyzing and pricing LP positions.

In this article, we explore the concept of AMM and its underlying math mechanisms. The objective of this research paper is to gain a comprehensive understanding of the returns associated with liquidity provision in AMM and investigate the overall situation in the liquidity pool. We use tools from financial econometrics to model these returns and economic theory to price the LP's position.

For these purposes, we use the Black-Scholes model of option pricing from Black et al., 1973 [8] and the Feynman-Kac stochastic framework from Kac, 1949 [9]. The Black and Scholes model reveals the relationship between option returns and factors such as the underlying stock's return and volatility. Therefore, the model is convenient for enabling the pricing of options by substituting specific values for the model parameters. Along with it, the Feynman-Kac formula allows us to solve the partial differential equations that often arise in stochastic finance.

This paper makes three primary contributions. Firstly, we provide a closed-form derivation for IL and empirically validate it using historical on-chain ETH-USDC data. It shows that

our model's predictions closely match realized returns. Secondly, we develop a closed-form financial model for pricing liquidity provider positions in Uniswap. We demonstrate that an LP's payoff exhibits option-like characteristics and derive explicit valuation formulas. Third, we incorporate a heavy-tailed, asymmetric return model into the LP risk framework. We reveal that standard Gaussian risk metrics significantly underestimate tail risk. In addition, we offer practical guidance for position sizing and extreme event management. Together, these contributions bridge traditional option-pricing theory with DeFi practice. It provides LPs and protocol designers with a quantitative toolkit for position valuation and risk assessment.

The rest of the paper is organized as follows. In Section 2 we present a literature review. Section 3 describes the LP position analysis for two versions of the protocol Uniswap. Section 4 delivers heavy-tailedness analysis of returns in the liquidity pool. In Section 5 we present the pricing of options that emerge in the market for LPs. Section 6 concludes the paper with a summary.

2 Related works

AMM has its roots in the field of prediction markets. As surveyed by Pennock and Sami, 2007 [10], these early mechanisms explored decentralized information aggregation. The Constant Function Market Maker underpins most modern AMMs. CFMM can be viewed as an extension of the utility-based market making framework developed in this context. This framework defines a bonding function based on utility indifference conditions for binary payoff Arrow-Debreu securities [11].

The transformative application of CFMMs as decentralized exchange mechanisms sparked significant recent interest. This shift was initially proposed by Buterin, 2016 [12] and rapidly developed, with early architectural discussions presented by Lu and Kppelmann, 2017 [13]. A comprehensive overview of these mechanisms as DeFi primitives is provided by Mohan, 2022 [4]. A formal and comprehensive analysis of the mechanics and properties of Uniswap-like markets was subsequently provided by Angeris et al., 2019 [14], establishing a rigorous mathematical foundation for the field.

The evolution from Uniswap V2 to V3 marked a paradigm shift by introducing concentrated liquidity. This innovation, which allows liquidity providers to allocate capital within specific price ranges, is thoroughly explained by Aigner and Dhaliwal, 2021 [16], who also derive the risk profile of an LP.

The mathematical intricacies and numerical calculations for managing positions within these discrete ticks are further detailed by Elsts, 2021 [17]. The profound implications of this design are analyzed in a seminal work by Lehar and Parlour, 2025 [5]. They formally model the Uniswap AMM from a microstructural perspective. They demonstrate how LPs effectively sell option-like volatility swaps to traders. It directly links LP returns and IL to the volatility of the underlying pool assets. That is a central theme in our analysis.

A critical part of research focuses on quantifying the costs and risks for LPs. The concept of impermanent loss is the primary metric for this. It is the loss incurred by a market making position versus keeping the initially allocated amounts fixed. Recent work by Vlasov et al., 2025 [6] provides a nuanced analysis of how exchange fees and pool parameters impact IL for "conservative" AMMs, offering a framework to assess the fee levels required to compensate LPs for their risk. Our paper contributes to this discourse by deriving closed-form expressions for expected Impermanent Loss in both V2 and V3 and simulating its distribution under realistic market conditions.

We want to emphasise that our paper is the first to provide a comprehensive framework for analysing and pricing these novel crypto instruments and options, just in case.

The recognition that LP positions replicate financial options has opened new frontiers for DeFi innovation. The Panoptic protocol, introduced by Lambert and Kristensen, 2022 [15], leverages this insight to enable perpetual, oracle-free options trading directly on Uniswap V3 pools. The emergence of such sophisticated instruments naturally raises the question: how to price it robustly? That addresses our research.

Finally, we should acknowledge that the field is moving from static models towards dynamic and AI-driven approaches. Indeed, Lim, 2024 [7] explores a predictive AMM architecture using deep reinforcement learning. He highlights the potential for future AMMs to dynamically adjust trading parameters. For instance, fees can be adjusted in response to market conditions to optimize for LP returns or trader efficiency.

It is worth noting that this area of knowledge is very dynamic and has great research potential. We position our paper within this vibrant and rapidly evolving literature. Existing works have excelled at describing AMM mechanics, modeling LP risks and innovating new protocols. Meanwhile, our contribution is to provide a unified framework for analyzing and pricing the option-like financial instruments that are inherent to LP positions.

We combine the theoretical option-pricing framework of Black-Scholes and Feynman-Kac with empirical heavy-tailed analysis of crypto returns. It aims to deliver a practical valuation model. It will bridge a critical gap between traditional financial theory and decentralized finance practice.

The next section is devoted to the position analysis of LP with the help of that mathematical finance theory.

3 Position analysis

In this chapter, we analyze the financial position of a LP in the Uniswap protocol. The core problem an LP faces is a classic economic trade-off: weighing potential rewards against inherent risks. The decision to deposit assets into a liquidity pool is an investment decision. And like any other, it requires a clear understanding of its risk-return profile. Again, the main risk is Impermanent Loss. It is the loss incurred by a market making position versus keeping the initially allocated amounts fixed.

Consequently, our first research question is to derive the Impermanent Loss for both versions of the AMM protocol Uniswap and compare it with different benchmarks. The most relevant benchmark is the "hold" or "passive" strategy. It represents the outcome if the LP had never provided liquidity in the first place. This comparison allows us to quantify the true cost of providing liquidity.

We use the pair ETH-USDC for our analysis (see Figure 2) due to its high trading volume, deep liquidity and status as a benchmark pair within the DeFi ecosystem [3]. To understand the evolution of AMM design, we consider both versions of Uniswap: Version 2 (V2) and Version 3 (V3). V2 allows uniform liquidity across all prices. Version 3 presents concentrated liquidity, which allows LPs to specify a custom price range for their capital.

3.1 Impermanent Loss with uniform liquidity

The liquidity provider takes the divergence risk: tokens proportion in the pool may change with time. Let's assume that the pool had X USDC and Y ETH with AMM constant product:

$$X \times Y = L^2 = const. \tag{1}$$

Here we denote L as a constant product number used in future calculations.

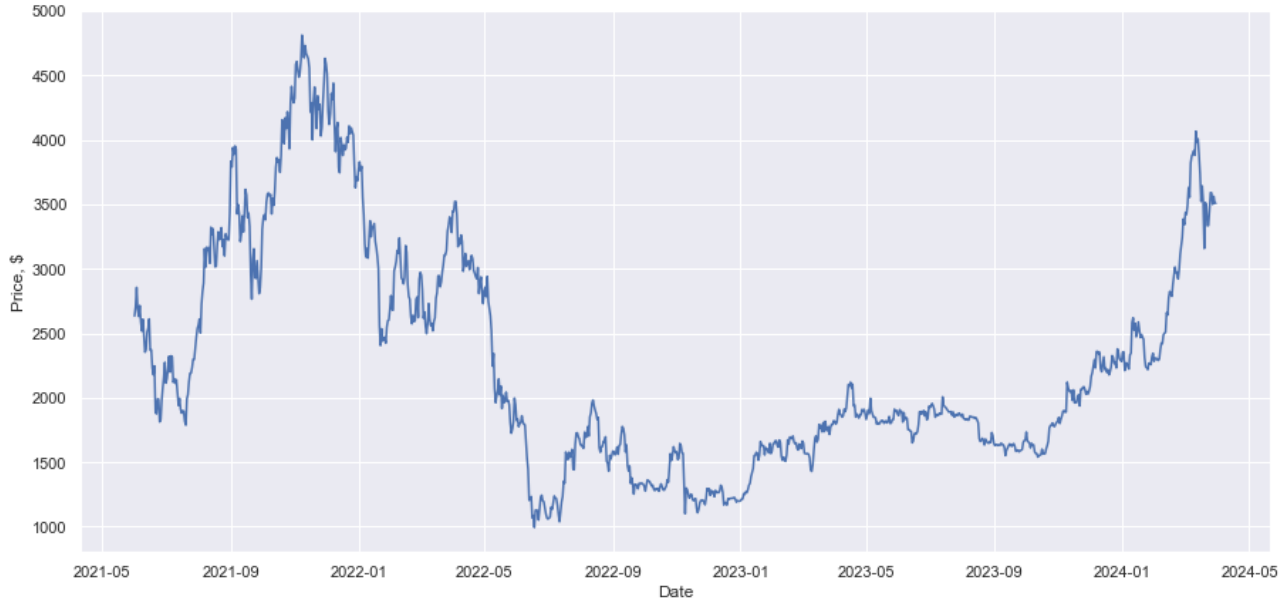

Figure 2: ETH-USDC price.

Suppose that USDC is a numeraire and ETH is a risky asset. Then the initial price is $P_0 = \frac{X}{Y}$, and the number of tokens $X = \sqrt{L^2 P_0}$.

At the start, the liquidity provider has tokens $X + Y = 2\sqrt{L^2 P_0}$, i.e. the value of the pool. The price changes exogenously to P. Now an arbitrageur trades with the pool until the new price in the pool becomes P.

The new liquidity providers portfolio is $2\sqrt{L^2 P}$, but in the cold wallet it would be $X + P\sqrt{\frac{L^2}{P_0}}$.

In other words, the value in the pool is:

$$V_2 = 2L\sqrt{P},$$

buy & hold portfolio refers to:

$$V_{2,hold} = L\sqrt{P_0}\left(1 + \frac{P}{P_0}\right).$$

Define the Impermanent Loss as the relative loss on paper until pool tokens withdrawal. It is computed as follows:

$$IL_2 = \frac{E[V_2] - E[V_{2,hold}]}{E[V_{2,hold}]}. \tag{2}$$

Assume Geometric Brownian Motion for the price movement. We perform all analysis in the risk-neutral measure. If W_t is a Wiener process with instantaneous variance 1 and drift 0, r and σ are drift and volatility, respectively, then the asset price at time t is:

$$P = P(t) = P_0 \exp\left[\left(r - \frac{\sigma^2}{2}\right)t + \sigma W_t\right].$$

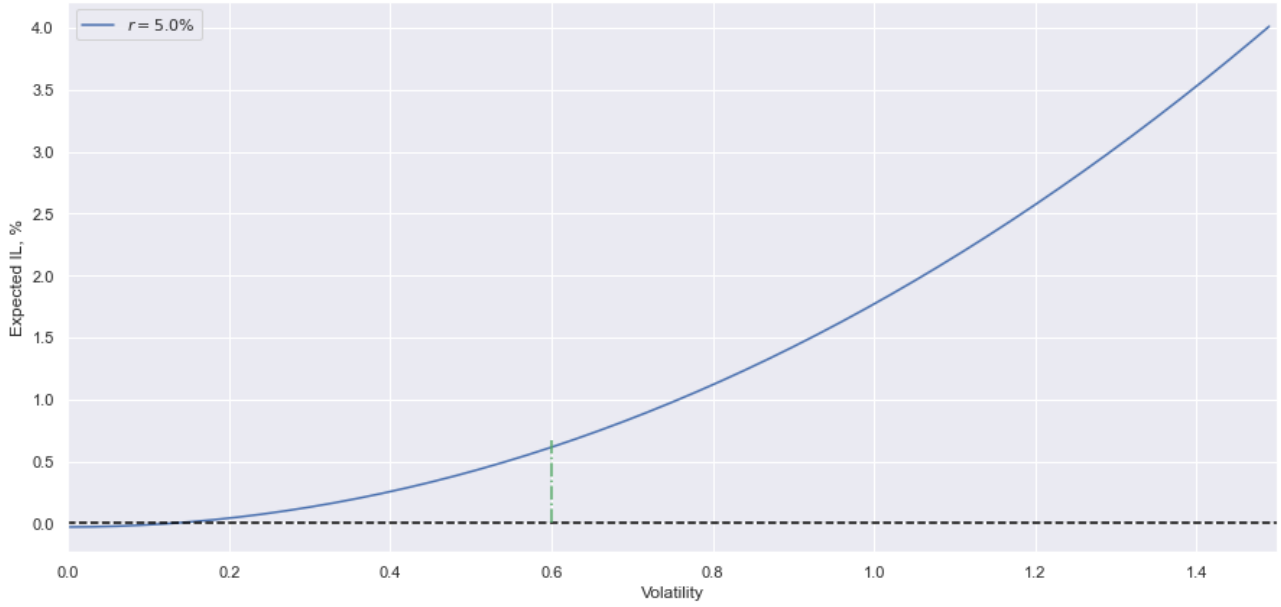

Figure 3: Impermanent Loss for V2 LP position in one year.

The expected value in the pool is computed using stochastic calculus.

$$E[V_2] = E\left[2L\sqrt{P_0}\exp\left(\left(\frac{r}{2}-\frac{\sigma^2}{4}\right)t+\frac{\sigma W_t}{2}\right)\right] =$$
$$= 2L\sqrt{P_0}\exp\left(\left(\frac{r}{2}-\frac{\sigma^2}{4}\right)t\right)\int_{-\infty}^{+\infty} e^{\frac{\sigma W_t}{2}}\frac{1}{\sqrt{2\pi t}}e^{-\frac{W_t}{2t}}dW_t =$$
$$= 2L\sqrt{P_0}\exp\left[\left(\frac{r}{2}-\frac{\sigma^2}{8}\right)t\right].$$

Also, assume that the fee paid to the LP α is a continuous compounding interest. Hence, the expected value in the pool that includes the rewarding fee α for the LP is:

$$E[V_2] = 2L\sqrt{P_0}\exp\left[\left(\frac{r}{2}-\frac{\sigma^2}{8}+\alpha\right)t\right]. \tag{3}$$

Expected buy & hold portfolio:

$$E[V_{2,hold}] = L\sqrt{P_0}(1+e^{rt}). \tag{4}$$

Combining the expressions (2), (3) and (4), we get Impermanent Loss for V2:

$$IL_2 = \frac{\exp\left((\alpha-\frac{\sigma^2}{8})t\right)}{ch(\frac{rt}{2})} - 1. \tag{5}$$

This equation elegantly captures the core financial trade-off for a V2 liquidity provider. The profitability of the position depends on the balance between three key parameters: rewarding fee, volatility and risk-free rate. While rewarding fee accumulates as income, risk-free rate represents the opportunity cost of capital in the benchmark "hold" strategy. Higher volatility directly increases the expected impermanent loss.

From Figure 3 under our model assumptions we can see that it is expected to be barely profitable to enter the position. This is due to $IL_2 > 0$ which means a positive profit (no more than 0.5% APR). This preliminary analysis was conducted using historical data for the ETH-USDC pair: a risk-free rate $r \approx 5\%$ and volatility $\sigma \approx 60\%$.

Figure 4: Uniswap V2 smoothed historical APR for ETH-USDC liquidity pool.

To validate our theoretical framework, we compare model predictions with historical data from the ETH-USDC liquidity pool on Uniswap V2. We obtain the realized APR for liquidity providers using historical on-chain data (Figure 4). The smoothed historical data is derived using moving average with the 30-days window. Average APR is 0.56% with fluctuations between approximately 0.4% and 0.7%. This volatility reflects varying market conditions over the period. So, this empirical result aligns closely with our theoretical prediction from Equation (5).

The result suggests that for an asset with high volatility like Ethereum or any other cryptocurrency, the fee income in the V2 model only marginally compensates for the significant risk of impermanent loss and the opportunity cost of capital. This finding immediately raises the question of whether the concentrated liquidity model of Uniswap V3 can offer a more attractive risk-return profile. V3 promises higher capital efficiency and fee earnings. Therefore, we need to compare these results for Version 3 against the same risk-free rate benchmark. That is, to determine if the new protocol design provides a substantive improvement for LPs.

3.2 Rewarding fee

The analysis of Impermanent Loss naturally leads to a critical question for any potential liquidity provider: what level of fee income is required to make providing liquidity worthwhile? To answer this, we calculate the fair or break-even fee rate α_{fair}. This is the minimum continuous fee rate that compensates the LP for both the impermanent loss risk and the opportunity cost of their capital, making the expected return of providing liquidity equal to the expected return of the passive buy-and-hold strategy.

Find the sufficient level of rewarding fee α to satisfy $IL_2 = 0$ from equation (5) in the risk-neutral world:

$$\alpha_{fair} = \frac{\sigma^2}{8} + \frac{1}{t}\left[\ln\left(1 + e^{rt}\right) - \ln 2 - \frac{rt}{2}\right].$$

This formula has a clear economic interpretation. The fair fee compensates for two main components. Firstly, the volatility term represents the compensation required specifically for

the risk of impermanent loss induced by the asset's volatility. It shows that the fee must be higher for more volatile pools, which is intuitively clear as LPs take on more significant rebalancing risks. This aligns with the findings of Lehar and Parlour, 2025 [5], who model LPs as sellers of volatility swaps.

The second term accounts for the opportunity cost of not earning the risk-free rate on the entire capital. It also incorporates a time-dependent adjustment. As the investment horizon t extends to infinity, this complex term simplifies, and the formula converges to a long-term equilibrium:

$$\alpha_{fair} \to \frac{\sigma^2}{8} + \frac{r}{2} \text{ when } t \to \infty.$$

This long-term approximation provides a practical rule of thumb. For instance, using a risk-free rate $r \approx 5\%$ and the historical volatility of the ETH-USDC pair $\sigma \approx 60\%$ we find that $\alpha_{fair} \approx 7\%$.

Now perform a more precise calculation for a one-year horizon. We obtained the break-even level of fee paid to the LP. Let the risk-free rate r be 5%. Then $\alpha_{fair} \approx 12.5\%$. This higher value reflects the greater short-term impact of opportunity cost and the convexity effects captured in the full formula. This result is economically significant and consistent with the fee structures observed in practice. Also, this number is consistent with the protocol Uniswap V3 [2]. Our calculated $\alpha_{fair} \approx 12.5\%$ is not a direct annual percentage fee but a continuous rate. When translated into a more familiar percentage, it justifies the need for higher fee tiers for highly volatile pairs. It was also discussed in the context of compensating LPs for impermanent loss by Vlasov, 2025 [6].

Note that the fair level of fee depends on time t. It implies that the required compensation is not static but varies with the LP's investment horizon. Shorter-term positions require a higher annualized fee to break even. This is due to the heightened impact of price swings over a brief period. Longer-term positions can tolerate a lower annualized fee as the law of large numbers and the compounding of fee income start to smooth out the volatility. This dynamic is essential for LPs to consider when formulating their investment strategies.

3.3 Impermanent Loss with concentrated liquidity

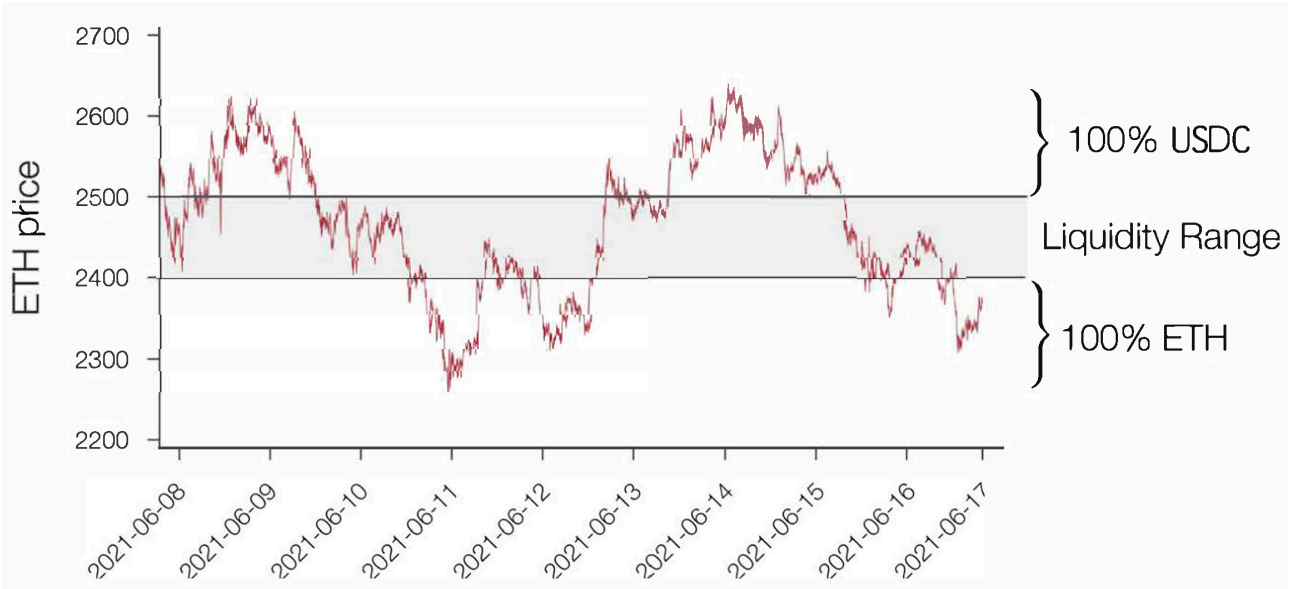

Figure 5: Uniswap V3 concentrated liquidity example.

In Uniswap V2 liquidity is distributed uniformly along the price curve from zero to infinity. However, V3-concentrated liquidity allows LPs to allocate liquidity within a custom price range $[P_{low}, P_{up}]$ and on a set of price ticks (Figure 5). Concentrated liquidity thus improves capital efficiency and offers better control over where liquidity is deployed within a pool. This design was detailed in the Uniswap V3 whitepaper [2] and slightly analyzed by Lehar and Parlour, 2025 [5].

The derivation of Impermanent Loss for V3 is more complex due to this segmented liquidity provision. Now the AMM constant product is:

$$\left(X + \frac{L}{\sqrt{P_{up}}}\right)\left(Y + L\sqrt{P_{low}}\right) = L^2, \tag{6}$$

where the LP decides in what price range $[P_{low}, P_{up}]$ to provide liquidity.

Introduce the notation. Let the strike $K = \sqrt{P_{up}P_{low}}$, range factor $\rho = \sqrt{\frac{P_{up}}{P_{low}}}$, $A = \frac{P}{K}$.

Then traders exchange tokens, the number of X and Y changes, the asset price moves, but the proportion in the pool remains. For instance, if the price $P \in [P_{low}, P_{up}]$, the number of tokens in the pool becomes:

$$X = L\frac{\sqrt{P_{up}} - \sqrt{P}}{\sqrt{P}\sqrt{P_{up}}},$$

$$Y = L\big(\sqrt{P} - \sqrt{P_{low}}\big).$$

We repeat the same steps for the cases $P < P_{low}$ and $P > P_{up}$ and get Impermanent Loss for V3:

$$IL_3 = \begin{cases} \frac{\sqrt{\rho}A-1}{A+1}, & A \le \frac{1}{\rho} \Leftrightarrow P \le P_{low} \ (100\%\ X); \\ \frac{\sqrt{\rho}}{\sqrt{\rho}-1}\left(\frac{2\sqrt{A}}{A+1} - 1\right), & \frac{1}{\rho} < A < \rho \Leftrightarrow P_{low} < P < P_{up} \ (Range); \\ \frac{\sqrt{\rho}-A}{A+1}, & A \ge \rho \Leftrightarrow P \ge P_{up} \ (100\%\ Y). \end{cases} \tag{7}$$

Now, we need to derive the similar expression as in (5) to compare the performance with V2. But taking the analytical expectation of the complex piecewise function in (7) under the assumption of Geometric Brownian Motion for the price P is mathematically complicated. So, the Impermanent Loss is computed by the Monte-Carlo method [18]. This method involves generating a large number of random price paths for ETH based on the historical parameters and calculating the average IL across all paths for a position with a specified range. The result for the price range of 20% ($\rho = 1.2$) for a one-year horizon is presented in Figure 6.

The key finding is starkly different from the V2 case. The expected Impermanent Loss is significantly negative, translating to a substantial positive profit for the LP. From Figure 6 we can see that it is profitable to enter the LP position. It means positive profit around 20% APR. Again, we used the historical data for $r \approx 5\%$ and for volatility $\sigma \approx 60\%$ to carry out the analysis.

Empirical validation of this theoretical prediction is provided by historical APR data from the ETH-USDC pool on Uniswap V3. It is shown in Figure 7. The smoothed historical APR for concentrated liquidity positions averages 20.13% over the period. There exist substantial variation reflecting the sensitivity of V3 returns to market conditions and price range selection. This empirical result closely matches our Monte-Carlo prediction. That confirms the superior profitability of concentrated liquidity.

Thus, we can conclude that Version 3 has a higher expected yield than Version 2. Also, the overall potential yield for Version 3 is greater than the risk-free rate ($r \approx 5\%$). This result underscores the revolutionary impact of concentrated liquidity. By concentrating capital

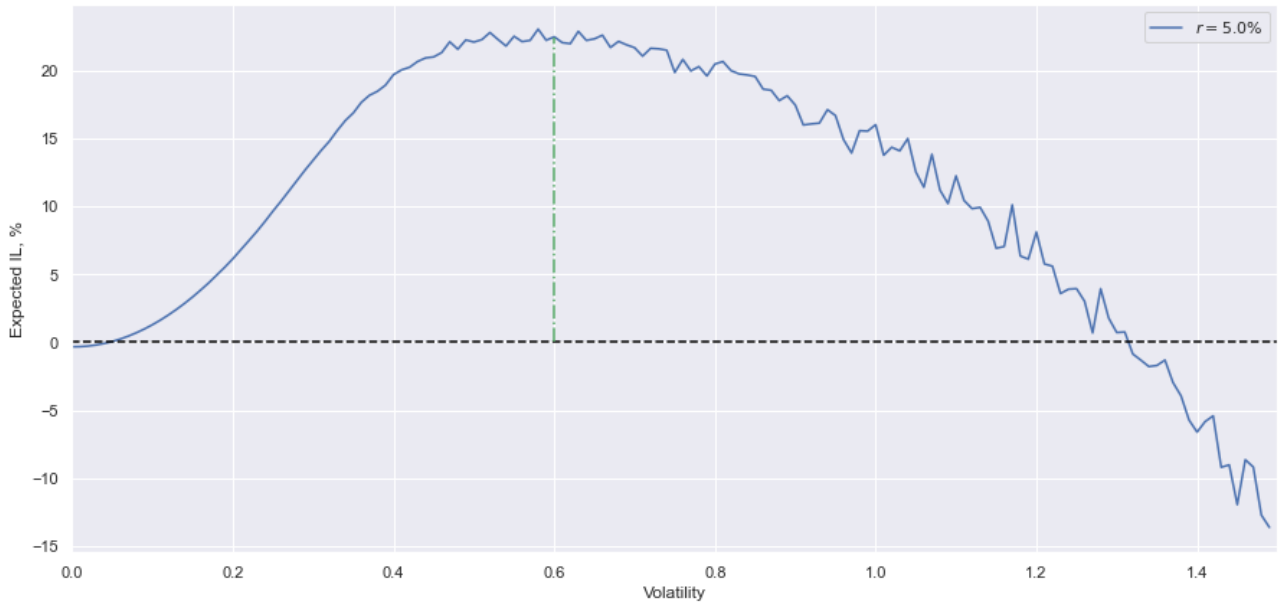

Figure 6: Impermanent Loss for V3 LP position in one year.

within a likely price range, the LP earns a much higher share of the trading fees relative to their committed capital. This increased fee revenue more than compensates for the impermanent loss risk. It leads to a superior expected yield compared to both the passive hold strategy and a V2 LP position.

So, the result tells us that the investor may consider the opportunity to enter the LP position along with the other instruments (deposits, etc.). However, the final decision must be made by considering specific factors. Among them are gas fees for rebalancing, personal risk tolerance and forecasts of future volatility and trading volume. That can significantly impact real-world returns. In particular, frictions (gas cost) reduce net APR.

The next Chapter is devoted to the econometric analysis of heavy-tailedness of returns in the liquidity pool.

4 Heavy-tailedness

The profitability and risk of a Liquidity Provider's position are fundamentally driven by the price movements of the underlying assets. Therefore, before proceeding to price the LP position, it is crucial to understand the statistical properties of returns within the ETH-USDC liquidity pool. In particular, focus on the statistical data on the ETH-USDC return in the Uniswap pool (combining Version 2 and Version 3). A key characteristic of financial returns, especially in crypto markets, is the presence of heavy tails, meaning that extreme price movements (large gains or losses) occur more frequently than predicted by a normal distribution.

The historical data is presented in Figure 8. We used the Uniswap API to get the data for the three years from 2021 to 2024. It can be seen that the graph experiences crises, extremes and large fluctuations. This motivates a formal investigation into the distributional properties of these returns.

4.1 Stylized facts

We begin by examining several well-established stylized facts of financial time series to see if they hold for the ETH-USDC pair. Firstly, check out some stylized facts applied to the financial

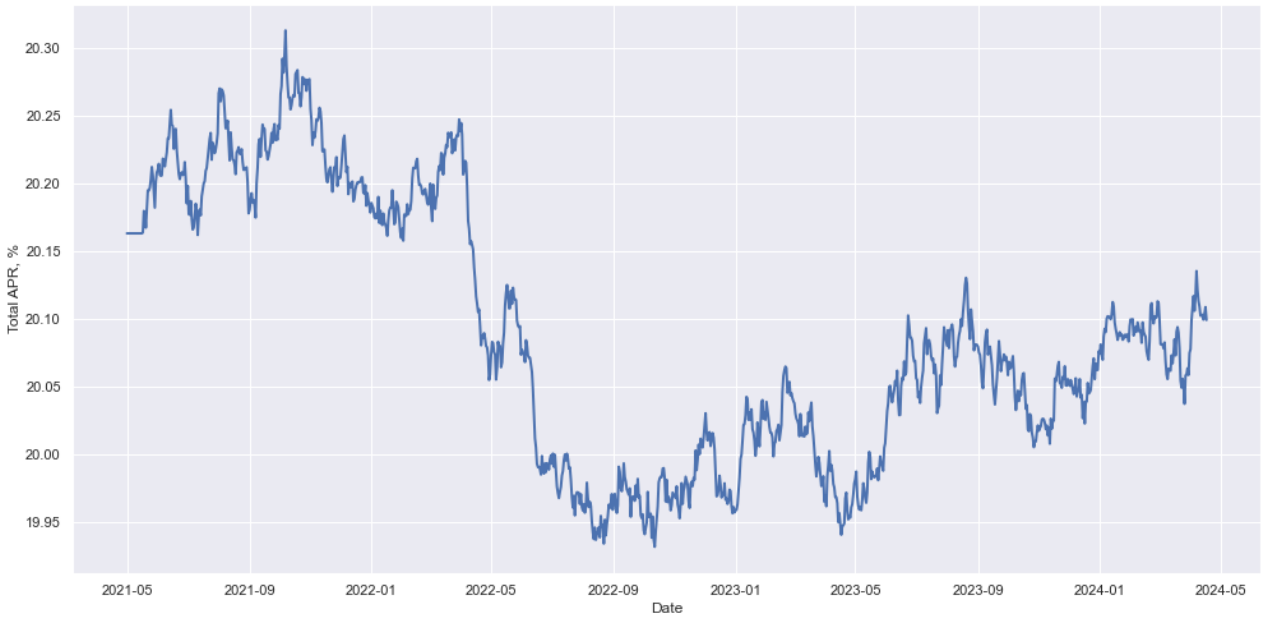

Figure 7: Uniswap V3 smoothed historical APR for ETH-USDC liquidity pool.

markets. These empirical regularities provide a baseline for understanding the dynamics of the crypto market.

Stylized fact	Expression	Value
Zero mean of returns	$\frac{1}{T}\sum_t^T R_t$	0.0009
No linear autocorrelation	$corr(R_t, R_{t-1})$	0.0029
Non-linear dependance (volatility clustering)	$corr(R_t^2, R_{t-1}^2)$	0.1637

Table 1: Stylized facts of ETH-USDC return.

So, from the Table 1 we can see that all stylized facts hold for the crypto world and the ETH-USDC pair in particular. Returns are approximately zero on average, show no significant linear correlation (efficient market hypothesis in the weak form), but display strong evidence of volatility clustering. The positive autocorrelation in squared returns indicates that periods of high volatility tend to be followed by more high volatility. This phenomenon is well-documented in traditional finance and equally relevant for risk management in DeFi.

Here arises the question about the exact distribution of returns (Figure 9). Even at first glance, outliers in the tails of the distribution become noticeable. Let us look at it in detail.

4.2 Tests for heavy-tailedness

Within the realm of probability theory, heavy-tailed distributions refer to probability distributions where the tails do not exhibit exponential bounds. It indicates that their tails are heavier than those of the exponential distribution.

Now, perform statistical tests for hypothesis H_0: the distribution of ETH-USDC returns in the liquidity pool is normal. That is, the first aim is to test whether the unconditional distribution of returns has heavier tails than normal.

From the Table 2 we can see that the there is indeed a heavy-tailed distribution of returns for the ETH-USDC pair in the Uniswap liquidity pool. Indeed, all tests unanimously reject the null hypothesis of normality at any conventional significance level. This provides strong

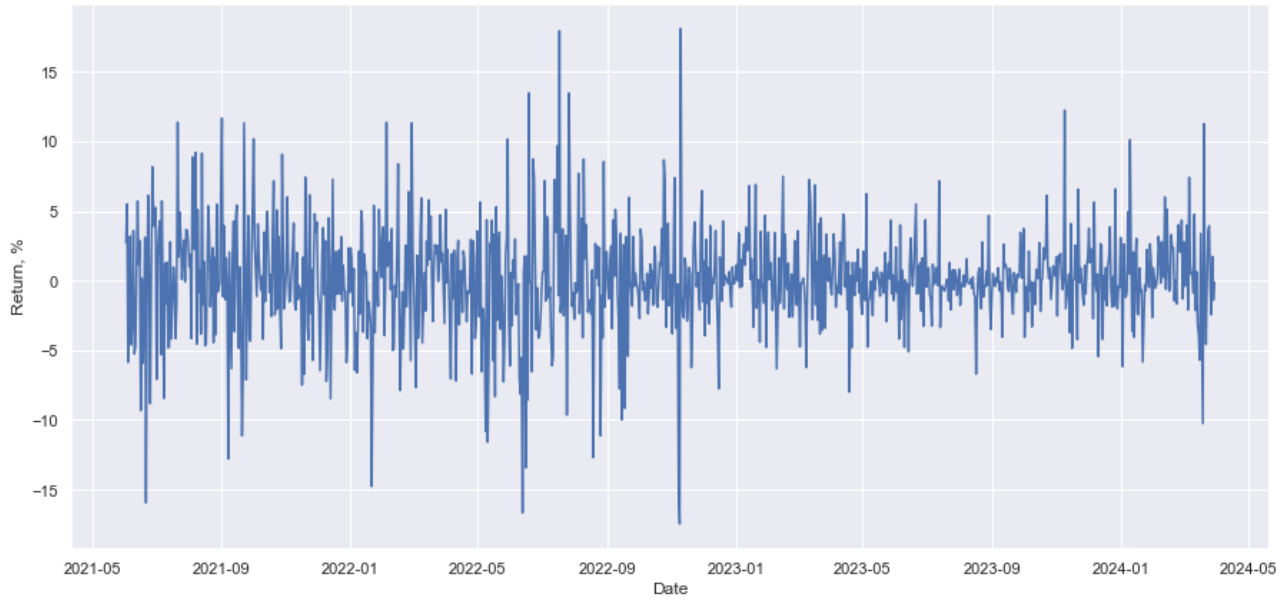

Figure 8: ETH-USDC return.

Test	p-value	Result
Kolmogorov-Smirnov	0.00001	Reject H_0 at 5% significance level
Anderson-Darling	0.00001	Reject H_0 at 5% significance level
Jarque-Bera	0.00000	Reject H_0 at 5% significance level
ShapiroWilk	0.00001	Reject H_0 at 5% significance level

Table 2: Tests for heavy-tailedness.

statistical evidence that the unconditional distribution of returns has heavier tails than a normal distribution.

Graphically, we can represent it by the Q-Q plot (see Figure 10). The Q-Q plot is generated by plotting two sets of quantiles against each other in a scatterplot. When both sets of quantiles originate from the same distribution, we anticipate observing the points aligning to form a nearly straight line. However, there exist clear deviations from the normal distribution on the edges. The pronounced S-shaped curve is a classic signature of heavy-tailed data. The deviations at the edges indicate that both large positive and large negative returns occur much more frequently than a normal distribution would allow.

4.3 Heavy-tails modelling

We have established the presence of heavy tails. Our second aim is to quantitatively model the distribution of returns. This allows us to quantify the risk of extreme price movements, which is crucial for accurately pricing the LP position and managing its risk. From the extreme value theory literature, a powerful approach for modelling the tails of a distribution is to assume they follow a power law. Often it is characterized by the Pareto distribution [19]. The key idea is that the probability of observing a return larger than a high threshold x decays polynomially:

$$Pr(|R_t| > x) \approx \frac{const}{x^{\zeta}}.$$

Degree of "heaviness" is determined by the parameter ζ. A smaller ζ indicates heavier tails, meaning extreme events are more likely. For instance, for returns on many stocks and

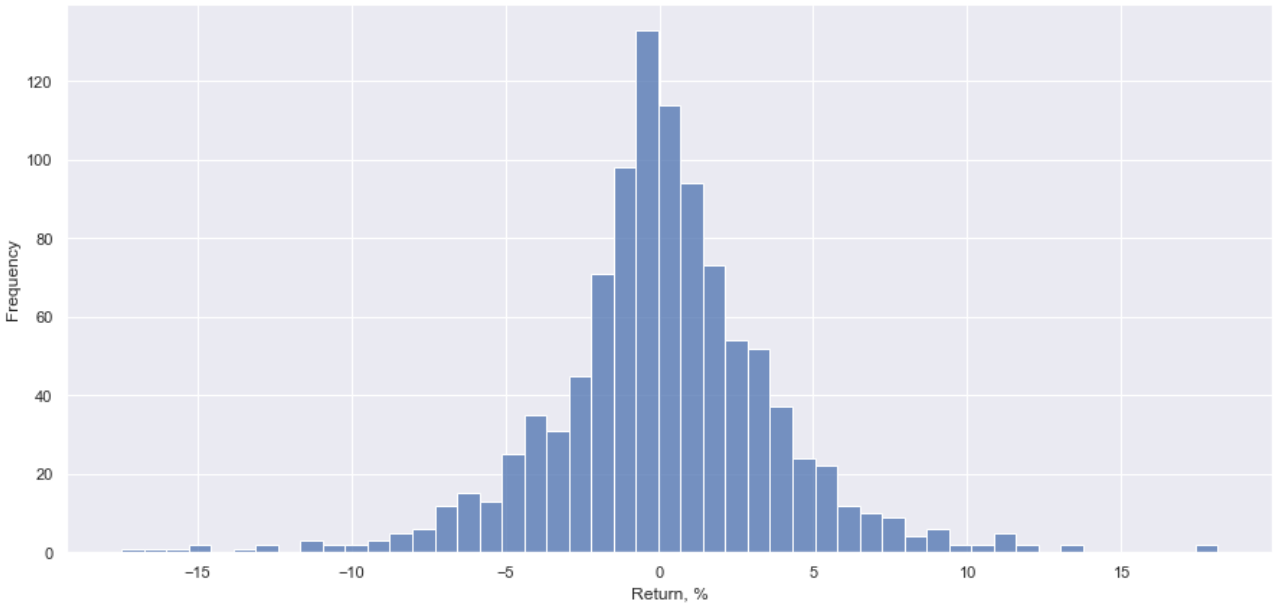

Figure 9: Distribution of ETH-USDC returns.

stock indices $2 < \zeta < 4$ (finite variance, infinite fourth moment). In contrast, technological innovations or emerging market assets can exhibit much smaller values, signifying even greater tail risk. Returns from technological innovations have $\zeta << 1$. Armed with this knowledge, find the parameter for ETH-USDC returns assuming Pareto distribution.

For our analysis, we exploit the method of Hill's estimator by Hill, 1975 [20]. Hill's estimator is a statistical method used to estimate the tail index of a distribution, particularly in extreme value theory. It is named after Hill, who introduced the method in the 1970s. This estimator is commonly used to gauge the tail of a distribution. It focuses on the extreme values or tails rather than the central part of the distribution.

The Hill estimator provides insights into the ends, helping to understand the likelihood of rare events. It is possible by analysing the order statistics of our returns.

The computation procedure is the following. If historical data has T observations, range them in descending order and take the first n large elements: $R_1 > R_2 > R_3 > ... > R_n > R_{n+1} > ... > R_T$. Here n is the truncation parameter. Then the formula for Hill's estimate of the parameter $\hat{\zeta}$ is:

$$\hat{\zeta}_{Hill} = \frac{n}{\sum_{k=1}^{n} [\log R_k - \log R_n]} \tag{8}$$

This estimator converges in probability to ζ. In addition, it is asymptotically normal when $k \to \infty$. It allows us to construct confidence intervals.

We examine both ends of the distributions. The left side represents losses, but we typically calculate tail-index estimators for positive data. This involves converting losses $R_i < 0$ to positive values by multiplying them by -1. Conversely, the right side $R_i > 0$ represents gains, and we retain these positive values.

As the number of order statistics k increases, the Hill estimator becomes more stable (Figure 11). The key question is which estimation is reasonable. A reliable approximation usually occurs in a stable region on the graph. While the graphs for both gains and losses do not show extensive stable areas, zooming in on the upper parts provides a clearer picture when $k = 10, ..., 0.1n$. Some estimates show stability in certain areas.

We use Python software to perform calculations and get the numbers. We obtain the results summarized in Table 3. For gains (right tail), a stable region might suggest an estimate around

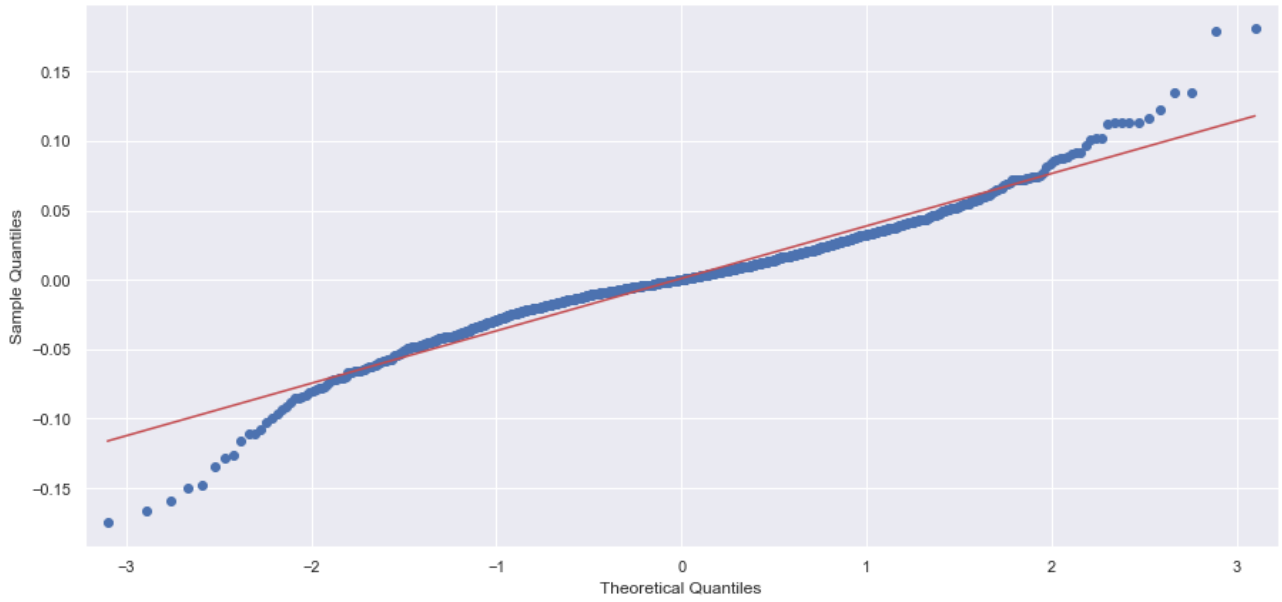

Figure 10: Q-Q plot of returns ETH-USDC against the normal distribution.

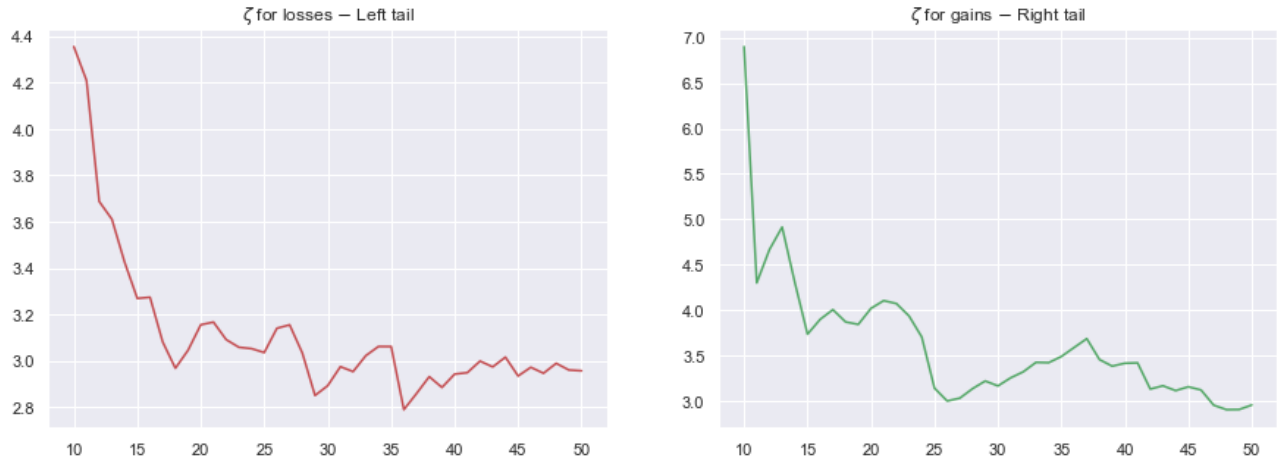

Figure 11: Hill estimator plot for the left (losses) and right (gains) tails.

$\hat{\zeta}_{gains} \approx 3.25$, but volatility still affects the plot. For losses (left tail), a somewhat reasonable approximation could be $\hat{\zeta}_{losses} \approx 2.9$, however, again, volatility makes it challenging to obtain a precise estimation.

Tail of distribution	Parameter	95% Confidence Interval
Left (losses)	$\hat{\zeta}_{losses}$	(2.649, 3.151)
Right (gains)	$\hat{\zeta}_{gains}$	(2.970, 3.530)

Table 3: Hill estimator modelling results for ETH-USDC return.

All in all, the results say that the left tail is heavier than the right one ($\hat{\zeta}_{gains} > \hat{\zeta}_{losses}$). In other words, losses experience more outliers than gains. However, the difference is not big enough but statistically significant. That is entirely consistent with our predictions because typically crypto markets are thought of as an emerging markets where negative shocks are more likely to happen, as shown in Ibragimov et al., 2013 [21]. The heightened risk of crashes directly impacts the LP's position. Sharp downward price movements can lead to significant

impermanent loss and increase the risk of the price exiting a concentrated liquidity range. This finding must be incorporated into any robust model for pricing the LP position.

4.4 Implications for LP

Our heavy-tailed model reveals critical vulnerabilities for liquidity providers. First, flash crashes pose a disproportionate threat to concentrated V3 positions. A flash crash is an abrupt, severe price decline. If a crash pushes the price outside an LP's specified range $[P_{low}, P_{up}]$, the position becomes 100% allocated to the depreciating asset (Eq. (7), case 1 or 3). Hence, temporary Impermanent Loss converts into an instantaneous, realized loss that may exceed typical VaR estimates. This estimate can be derived from the Pareto distribution. The general analysis of risk-management metrics applicable to this specific distribution is presented by He, 2022 [22]. The heavier left tail ($\hat{\zeta}_{losses} < \hat{\zeta}_{gains}$) implies that such adverse jumps occur more frequently and with greater severity than predicted by thin-tailed distributions. Thus, standard risk metrics like Gaussian-based Value-at-Risk (VaR) would significantly understate the true downside exposure.

Second, periods of low trading volume (liquidity droughts) compound this risk. During such phases, fee income α declines — the primary compensation for IL. When combined with heavy-tailed price movements, this creates scenarios where LPs bear significant downside risk without adequate fee-based compensation. These scenarios of low liquidity are presented in Figures 4 and 7. In particular, 2022 can be viewed as a "crypto winter" with low trading volumes on DEXs. Consequently, it reflects a low fee APR for LPs.

Therefore, these implications directly inform practical risk-management. Any robust risk-management framework for AMM liquidity must account for the heavy-tailed, asymmetric nature of crypto returns. LPs should consider hedging strategies for concentrated positions and monitor trading volume trends alongside price volatility. Since the market tends toward more frequent and severe negative shocks than positive ones, LP should carefully choose the price range. Indeed, tighter ranges ($\rho \approx 1.1$) offer higher fee capture but require more conservative capital allocation due to elevated crash risk.

Now we know the properties of the liquidity pool and are ready to price the position in the next Chapter.

5 Position pricing

In this chapter, we look at the position the LP holds. We transit from analyzing the risks and returns of a LP to formally pricing the LP position itself. Recall that now the LP provides liquidity (deposits tokens) and, therefore, may close the position (withdraw tokens) at any time.

So, an LP's share in a pool can be viewed as a financial derivative: its value is derived from the future price paths of the underlying assets (ETH and USDC). Assume that this position is a perpetual contract with no expiration. The LP deposits tokens at time $t = 0$ and has the right to withdraw their share of the pool at any future time t. For the purpose of valuation, we can model this as a perpetual contract with no expiration, as the position remains open until the LP decides to close it.

Our previous analysis demonstrated that providing concentrated liquidity in Uniswap V3 offers a superior expected yield compared to the uniform liquidity model of V2. Therefore, we will focus primarily on deriving a valuation model for the more complex and economically significant V3 position. However, we begin with the simpler case of Uniswap V2 to build intuition and establish a foundational pricing framework.

Our aim is to estimate the LP position in the pool and derive the explicit expressions. This involves discounting the expected future payoff of the position at the risk-free rate.

5.1 LP position with uniform liquidity

We start with the valuation of a position in Uniswap V2. Recall that the value of the LP's share in the pool, as a function of the current price P, is given by:

$$V_2(P) = 2L\sqrt{P}.$$

This value represents the total amount of USDC the LP would receive if they withdrew their liquidity at price P. To price the position at inception ($t = 0$), we are interested in the value of the future payoff relative to the initial investment.

The net payoff of the LP position at time t is the difference between the final and initial value:

$$Payoff_2 = 2L\big(\sqrt{P} - \sqrt{P_0}\big).$$

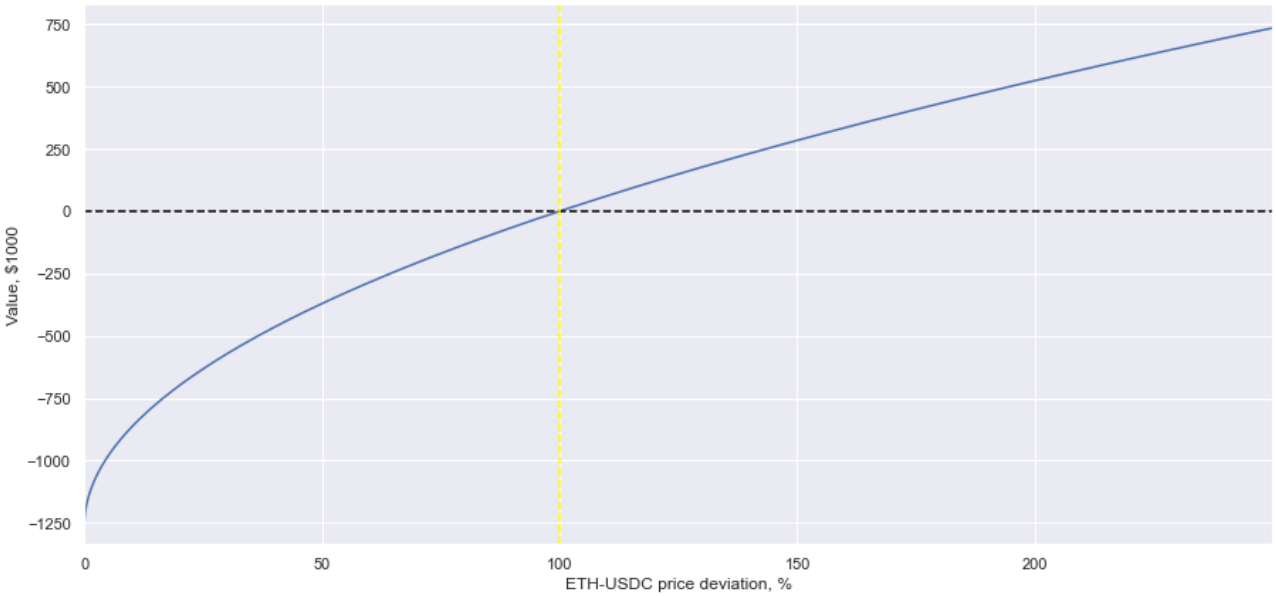

Figure 12: Payoff profile of a LP position with uniform liquidity.

This payoff function, illustrated in Figure 12, is a non-linear function of the price P. It is concave and resembles the payoff of a portfolio of volatility-sensitive instruments. It can be intuitively understood as being long a "square root forward" on the asset price.

To find the present value of this position at time $t = 0$, we calculate the expected value of the payoff under the risk-neutral measure and discount it at the risk-free rate r. The risk-neutral expectation accounts for the market price of risk and allows us to value the derivative without estimating the asset's real-world drift.

The expected value in the pool at the beginning in the risk-neutral measure is $V_2(P,t) = e^{-rt}E[Payoff_2] = e^{-rt}E[2L(\sqrt{P} - \sqrt{P_0})]$.

Assume the price P follows a Geometric Brownian Motion with volatility σ. Include the continuous fee income at rate α. And solve the expectation using properties of the log-normal distribution. This yields the following closed-form expression:

$$V_2(P,t) = e^{-rt}2L\sqrt{P_0}\left(\exp\left[\left(\frac{r}{2} - \frac{\sigma^2}{8} + \alpha\right)t\right] - 1\right). \tag{9}$$

That is the current price of the Uniswap V2 position for the LP. Several key insights emerge from this formula We observe a positive relationship between the rewarding fee and the price.

Also, without any fees, the price is not positive for historical data. This mathematically confirms that providing liquidity in V2 is unprofitable without compensation. That coincides with the concept of Impermanent Loss, since the investor would just lose money providing liquidity with no any compensation.

This V2 pricing formula provides a benchmark and a methodological foundation. We now proceed to the more complex and relevant task of pricing the Uniswap V3 LP position.

5.2 LP position with concentrated liquidity

Now focus on the most complicated part of our analysis. In particular, look at the core of our pricing analysis: valuing a LP's position in Uniswap V3.

The introduction of concentrated liquidity fundamentally changes the payoff structure, making it more complex but also more capital-efficient than its V2 counterpart. Following the framework established in the Uniswap V3 whitepaper [2], we can express the payoff of a V3 LP position as a piecewise function. It depends on the final price P relative to the chosen range $[K/\rho, K\rho]$, where K is the strike (typically the entry price P_0) and ρ is the range factor. The payoff of the Uniswap V3 LP position can be written as follows:

$$Payoff_3(P) = \begin{cases} P, & P < \frac{K}{\rho} \\ \frac{2\sqrt{PK\rho} - P - K}{\rho - 1}, & \frac{K}{\rho} < P < K\rho \\ K, & P > K\rho. \end{cases} \tag{10}$$

Range factor ρ determines the "sharpness" of the transition between holding ETH and USDC (i.e., money). That is, this factor controls the capital efficiency and risk profile. A smaller ρ (tighter range) leads to a "sharper" transition between the asset domains but requires more active management. This payoff profile is visualized in Figure 13. It is intrinsically non-linear and option-like.

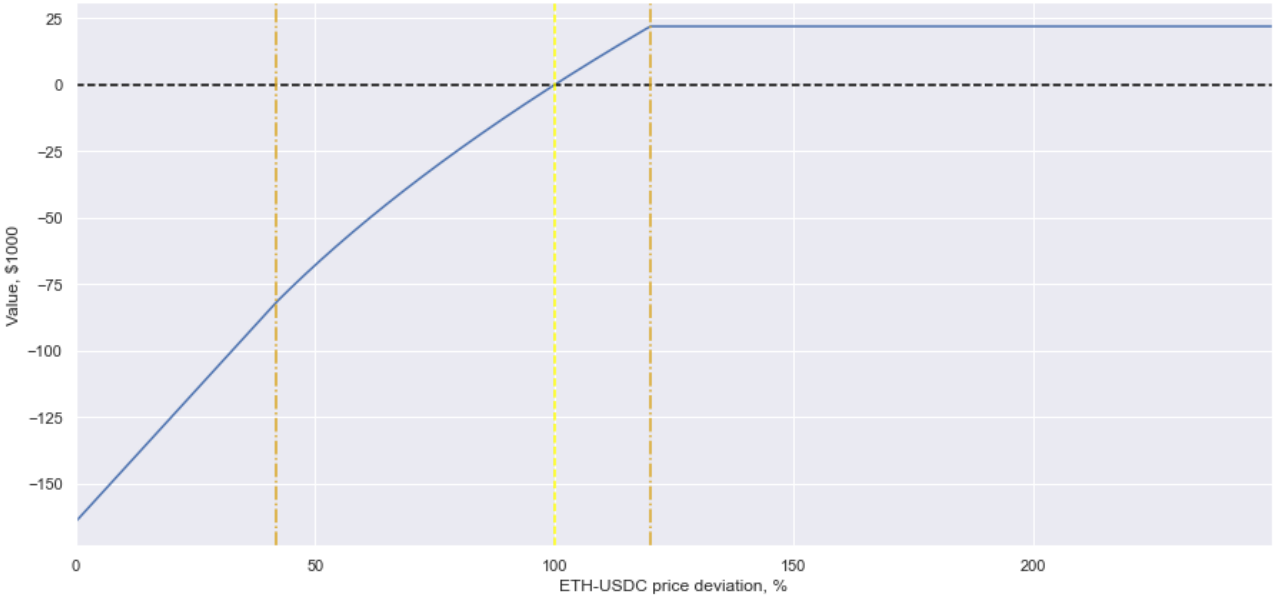

Figure 13: Payoff profile of a LP position with concentrated liquidity.

A powerful method to price this complex payoff is to decompose it into more familiar financial derivatives. As illustrated in Figure 14, the payoff can be separated into two components:

1. A short put option component, that corresponds to a single-price position;

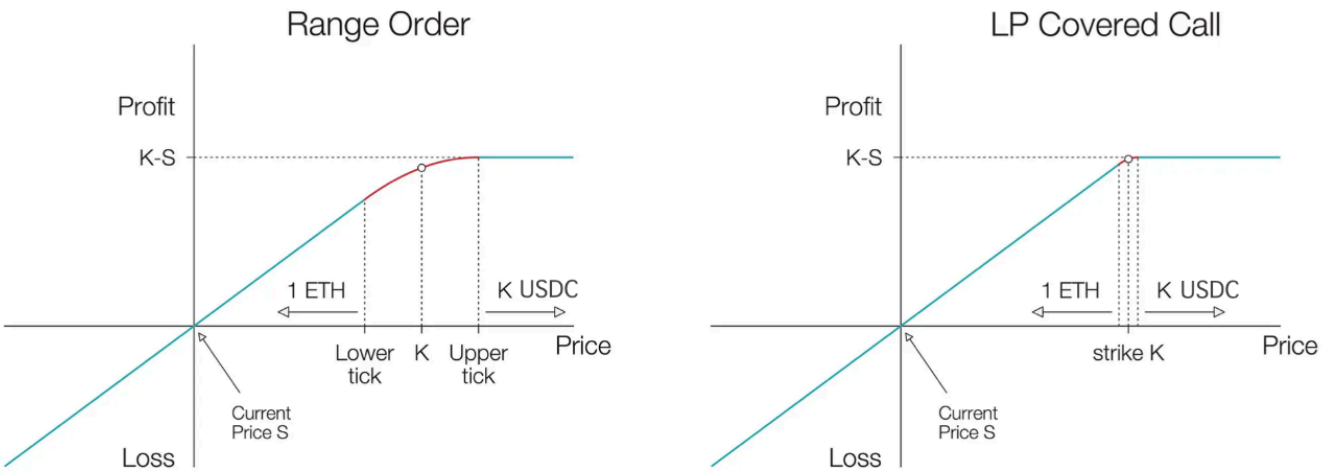

Figure 14: Concentrated liquidity payoff is similar to the covered call payoff.

2. A range-specific payoff that is active only between the lower and upper price bounds.

This decomposition is expressed mathematically as:

$$Payoff_3(P) = -\max(K-P,0) + \begin{cases} K, & P < \frac{K}{\rho} \\ \frac{2\sqrt{PK\rho}-P\rho-K}{\rho-1} + K, & \frac{K}{\rho} < P < K \\ \frac{2\sqrt{PK\rho}-P-K\rho}{\rho-1} + K, & K < P < K\rho \\ K, & P > K\rho. \end{cases} \tag{11}$$

That is, $Payoff_3(P) = Payoff_{Put}(P) + Payoff_{Range}(P)$.
This decomposition directly leads to the present value of the V3 LP position:

$$V_3(P,t) = V_{Put}(P,t) + V_{Range}(P,t). \tag{12}$$

Here $V_{Range}(P,t)$ is the value of the ranged part of the LP position. $V_{Put}(P,t)$ can be priced using the standard Black-Scholes put option formula, and $V_{Range}(P,t)$ requires a more specialized valuation.

Viewing the LP position through the lens of option theory provides direct risk-management insights. The decomposition in (11) shows that a V3 LP position is exposed to the following factors. Firstly, directional exposure (Delta) is vital. LP can hedge it by taking offsetting positions in the underlying assets when the price approaches range boundaries. Secondly, since LPs position has a negative volatility exposure (Vega), the position can be hedged via options on ETH. In addition, the short put component exposes the LP to downside tail risk (flash crashes) as were discussed in Section 4.

Remember that the LP enters the position when $P_0 = K$. Here arises a question: how can we approximate the value of the ATM option using existing financial instruments? Therefore, we need to look at the pricing of at-the-money (ATM) options.

5.3 Approximation for ATM options

A key insight for practical valuation is that the value of the range component near the ATM point ($P \approx K$) behaves similarly to a short-dated option. Use the Black-Scholes formula to approximate the pricing of ATM options. The approximate formula is due to Brenner et al., 1988 [23].

For the at-the-money call option, we have $P = Ke^{-rT}$. Plug this into the standard Black-Scholes formula:

$$V_{Call}(P,T) = N(d_1)P - N(d_2)Ke^{-rT}. \tag{13}$$

Recall that $d_1 = \frac{\ln \frac{P}{K} + (r + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T}$.
Then we get that

$$V_{Call}(P,T) = \left[N\left(\frac{1}{2}\sigma\sqrt{T}\right) - N\left(-\frac{1}{2}\sigma\sqrt{T}\right)\right]P. \tag{14}$$

Now, Taylor's formula implies for small x that

$$N(x) = N(0) + N'(0)x + N''(0)\frac{x^2}{2} + O(x^3). \tag{15}$$

Combining two equations, we will get some obvious cancellations that

$$V_{Call}(P,T) = P\left(N'(0)\sigma\sqrt{T} + O(\sigma^3\sqrt{(T)^3})\right).$$

But

$$N'(0) = \frac{1}{\sqrt{2\pi}}.$$

So, for the ATM options for small $\sigma\sqrt{T}$ we have:

$$V_{Call}(P,T) \approx P\sqrt{\frac{\sigma^2}{2\pi}T}. \tag{16}$$

Now note that the shape of the Uniswap V3 LP position at a given scaling factor ρ is very similar to that of a covered call option at a time T_{eff} away from expiration (Figure 14). Thus, equating (10) and intrinsic value (16) when $P = K$ (ATM option) we obtain the effective time to expiration of the LP option. It captures the option-like characteristics of the LP position. That is, if $Value_{Range}(P \to K, t \to 0) \approx Payoff_{Range}(P = K) \approx V_{Call}(P,T)$ the effective time is:

$$T_{eff} = \frac{2\pi}{\sigma^2}\left(\frac{\sqrt{\rho}-1}{\sqrt{\rho}+1}\right)^2. \tag{17}$$

For typical parameters, if the price range factor is $\rho = 1.2$ and $\sigma = 60\%$, $T_{eff} \approx 0.04$. It is equivalent to 2 weeks. This short duration makes intuitive sense, as the value of being within a specific range is highly sensitive to imminent price changes, similar to a short-dated option.

This equivalence allows us to construct a tractable approximation for the full position value. Decompose the value of the position and use the effective time to expiration from the equation (17):

$$V_3(P,t) = Put(P,t) + Value_{Range}(P,t) \approx V_{Put}(P, t + T_{eff}),$$

where $V_{Put}(P,t)$ is calculated by the Black-Scholes model. This approximation suggests that the value of a Uniswap V3 option at time t is approximately equivalent to a short put option that expires at the time T_{eff}.

For example, estimate the value V_3 for $t = 1$ year, initial price of ETH 2000\$, volatility $\sigma = 60\%$, risk-free rate $r = 5\%$ to validate our result: $V_3(P = K = 2000\$, t = 1) \approx 400\$$.

5.4 Exact solution

On the other hand, we can go deep into math. For maximum accuracy, we can derive an exact pricing formula. To obtain the exact closed-form solution for $V_3(P,t)$, it is possible to apply the Feynman-Kac formula to $V_{Range}(P,t)$. Then the value of the range component is given by the risk-neutral expectation of its discounted payoff:

$$Value_{Range}(P,t) = \int_0^{+\infty} f_{GBM}(x,P,t,r,\sigma)\ e^{-rt} Payoff_{Range}(x)dx. \tag{18}$$

Here $f_{GBM}(x,P,t,r,\sigma)$ is the Probability Density Function for GBM (i.e., for a lognormal distribution):

$$f_{GBM}(x,P,t,r,\sigma) = \frac{1}{x\sigma\sqrt{2\pi t}} exp\left[-\frac{\left(\ln\frac{x}{P} - (r-\frac{\sigma^2}{2})\ t\right)^2}{2\sigma^2 t}\right]. \tag{19}$$

Then we take the integral and obtain, albeit complex, closed-form solution:

$$\begin{aligned}
Value_{Range}(P,t) = -\frac{e^{-rt}}{2(\rho-1)}\Bigg[&- 2\sqrt{\rho PK}e^{\frac{1}{8}t(4r-\sigma^2)} erf\left(\frac{rt - \ln\frac{K\rho}{P}}{\sqrt{2}\sigma\sqrt{t}}\right) + \\
&+ 2\sqrt{\rho PK}e^{\frac{1}{8}t(4r-\sigma^2)} erf\left(\frac{rt - \ln\frac{K}{P\rho}}{\sqrt{2}\sigma\sqrt{t}}\right) + Kerf\left(\frac{2\ln K - 2rt + 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) + \\
&+ Kerf\left(\frac{2\ln\frac{K}{\rho} - 2rt + 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) - \rho Kerf\left(\frac{2\ln(K\rho) - 2rt + 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) + \\
&+ (\rho-1)Pe^{rt} erf\left(\frac{-2\ln K + 2rt - 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) - \\
&- \rho Pe^{rt} erf\left(\frac{-2\ln\frac{K}{\rho} + 2rt - 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) + \\
&+ Pe^{rt} erf\left(\frac{-2\ln(K\rho) + 2rt - 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right) - \\
&- Kerf\left(\frac{2\ln K - 2rt + 2\ln\frac{1}{P} + \sigma^2 t}{2\sqrt{2}\sigma\sqrt{t}}\right)\Bigg].
\end{aligned} \tag{20}$$

We used the notation for the Gauss error function:

$$erf(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-h^2}dh.$$

Then the value can be found using the equation (20) as:

$$V_3(P,t) = Put(P,t) + Value_{Range}(P,t).$$

Again, estimate the value V_3 for $t = 1$ year, initial price of 2000\$, volatility $\sigma = 60\%$, risk-free rate $r = 5\%$ to validate our result: $V_3(P = K = 2000\$, t = 1) \approx 500\$$. Here we can see a little discrepancy between our approximate result and the exact solution. It indicates that our simple ATM approximation tends to underprice the real value of the LP position. This is likely because the approximation does not fully capture the value contributed by the entire range of prices within the active band, especially those farther from the ATM point.

Thus, we have developed two methods for pricing the LP position with concentrated liquidity: a simple and intuitive approximation based on option equivalence and a precise closed-form solution. The exact formula, while more complex, provides a robust tool for the preliminary valuation of these novel financial instruments in the market. The significant positive value confirms the economic viability of providing concentrated liquidity in Uniswap V3 for the analyzed market conditions.

6 Conclusion

Decentralized Finance (DeFi) and decentralized exchanges particularly differ from centralized exchanges in that they do not rely on a central entity to oversee deposits. Rather, interactions between participants are facilitated by smart contracts on the Ethereum blockchain. Assets are exchanged through smart contracts that employ an Automated Market Making function (AMM), where each amount of token has a predetermined price. This is due to the constant function market making that uses liquidity pools with tokens, and traders wanted to exchange tokens.

Since the most popular AMM is presented in the protocol Uniswap, the great part of interest was to analyse the position held by liquidity providers (LP). They have the option to withdraw their tokens from the pool and, therefore, earn money. The protocol has two versions: V2 and V3. The first one presents the uniform liquidity provision for LP. The second one allows LPs to select the price range and exploit the concentrated liquidity mechanism. This paper provides a comprehensive financial and econometric analysis of the LP's position in AMM. It delivers insights into its risks, returns and fundamental valuation.

The analysis has several key contributions. Firstly, this paper presents an analysis of the LP position. We conduct a detailed position analysis for both uniform and concentrated liquidity. For uniform one, we derive a closed-form expression for the expected Impermanent Loss (IL). It is confirmed that IL is a direct function of asset volatility and the rewarding fee. The analysis reveals that for a highly volatile pair like ETH-USDC or other cryptocurrencies, providing uniform liquidity in V2 is only marginally profitable. This is due to the fee income barely compensates for the IL and opportunity cost. In contrast, for Uniswap V3 we demonstrate that concentrated liquidity significantly enhances capital efficiency. V3 enables substantially higher fee earnings by allowing LPs to allocate funds within a specific price range. Our Monte-Carlo simulations show that a concentrated position can generate an expected profit of around 20% Annual Percentage Rate. It outperforms both the uniform liquidity model and the traditional risk-free rate. So, the LP may participate in liquidity provision (depositing) into the pool.

Secondly, the author provides an in-depth econometric analysis of the returns of the ETH-USDC pair in the liquidity pool. It is confirmed the presence of well-known stylized facts, such as volatility clustering. More importantly, we establish that the return distribution is heavy-tailed, following a power law [19, 20]. A critical finding is the asymmetry in tail risks. Indeed, the left tail (losses) is heavier than the right tail (gains). This implies a greater likelihood of extreme negative price movements. Note that it is a characteristic of emerging markets [21] and must be factored into any robust risk management framework for LPs.

Finally, the pricing of the LP options is presented using the Black-Scholes model as a starting point. That is the most significant part of the research.

This paper introduces a novel pricing framework for the LP position. The key is to treat it as a derivative financial instrument. For uniform liquidity, we present a direct valuation formula. For the more complex concentrated liquidity position, we develop two pricing approaches. The first one is an approximate method based on decomposing the LP's payoff into a short put option and a range component. It includes deriving an effective time to expiration [23]. This method offers an intuitive link to traditional option pricing theory. The second approach is

an exact closed-form solution obtained by applying the Feynman-Kac formula that provides a precise valuation tool [8, 9].

The results from both methods are economically sensible and consistent with observed market dynamics. Moreover, it validates the concentrated liquidity model's potential for profitability as identified in our initial analysis. Our work thus bridges the gap between traditional financial theory and the novel instruments emerging in DeFi. Specifically, we work behind the Black-Scholes framework, contributing to the growing literature on the subject [2, 5].

While our empirical analysis centers on the benchmark ETH-USDC pair, the underlying pricing framework is derived from the universal mathematics of constant-product market makers. It therefore applies to any asset pair. Pair-specific outcomes are determined by three inputs: volatility and fee rate. For instance, stablecoin pairs would exhibit lower IL and more predictable returns. Meanwhile, highly volatile pairs would see an even greater relative advantage of concentrated liquidity. This scalability makes our model a versatile tool for LP position valuation across the diverse crypto assets.

It is worth noting the practical considerations and potential extensions of our results. In practice, several realworld factors require adjustment. Firstly, discrete fees (fees per trade) can be approximated by our continuous rate *alpha* over moderate horizons. Secondly, frictions (gas cost) reduce net APR. It can be incorporated as a fixed or proportional deduction from the models predicted APR. Thirdly, protocol upgrades may alter fee structures. But the fundamental relationship between volatility, range width and LP returns will remain. Thus, practitioners can use our formulas as a baseline. Then they can apply scenariospecific adjustments for frictions and protocol details. Future work could extend the model to formally optimize range selection and rebalancing frequency under gascost constraints.

In summary, this research provides LPs and protocol designers with a quantitative toolkit to assess the profitability and value of liquidity provision. It highlights the superior capital efficiency of concentrated liquidity. Along with it, it underscores the critical importance of understanding and pricing tail risks. The DeFi landscape continues to evolve. And the integration of traditional financial econometrics with blockchain-based market microstructure is essential.

References

[1] Adams, H., Zinsmeister, N., & Robinson, D. (2020). *Uniswap v2 core*.

[2] Adams, H., Zinsmeister, N., Salem, M., Keefer, R., & Robinson, D. (2021). *Uniswap v3 core*.

[3] DefiLlama. *DeFi Llama: Overview*. Retrieved from https://defillama.com.

[4] Mohan, V. (2022). Automated market makers and decentralized exchanges: a DeFi primer. *Financial Innovation*, 8(1), 20.

[5] Lehar, A., & Parlour, C. (2025). Decentralized exchange: The Uniswap automated market maker. *The Journal of Finance*, 80(1), 321374.

[6] Vlasov, R., Gorgadze, V., & Barger, A. (2025). The impact of the exchange fees on impermanent loss of liquidity providers for conservative automated market makers. *The Journal of The British Blockchain Association*.

[7] Lim, T. (2024). Predictive crypto-asset automated market maker architecture for decentralized finance using deep reinforcement learning. *Financial Innovation*, 10(1), 144.

[8] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637654.

[9] Kac, M. (1949). On distributions of certain Wiener functionals. *Transactions of the American Mathematical Society*, 65(1), 113.

[10] Pennock, D. M., & Sami, R. (2007). Computational aspects of prediction markets. In N. Nisan, T. Roughgarden, È Tardos, & V. V. Vazirani (Eds.), *Algorithmic game theory* (pp. 126). Cambridge University Press.

[11] Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3), 265290.

[12] Buterin, V. (2016). Ethereum: platform review. *Opportunities and challenges for private and consortium blockchains*, 45 (2016): 1-45.

[13] Lu, A., & Kppelmann, M. (2017, March). Building a decentralized exchange in Ethereum. *Gnosis Blog*.

[14] Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., & Chitra, T. (2019). An analysis of Uniswap markets. *Cryptoeconomic Systems Journal*.

[15] Lambert, G., & Kristensen, J. (2022). Panoptic: the perpetual, oracle-free options protocol. *Panoptic Whitepaper*. Retrieved from https://panoptic.xyz/whitepaper.pdf.

[16] Aigner, A. A., & Dhaliwal, G. (2021). Uniswap: Impermanent loss and risk profile of a liquidity provider. *SSRN Electronic Journal*.

[17] Elsts, A. (2021). Liquidity math in Uniswap V3. *Atis Elsts Blog*. Retrieved from https://atiselsts.github.io/pdfs/uniswap-v3-liquidity-math.pdf.

[18] Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, 44(247), 335341.

[19] Ahmed, H. H., & Almetwally, E. M. (2021). Generalizations of Pareto distribution with applications to lifetime data. *Mathematical Problems in Engineering*, 2021, 115.

[20] Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5), 11631174.

[21] Ibragimov, M., Ibragimov, R., & Kattuman, P. (2013). Emerging markets and heavy tails. *Journal of Banking & Finance*, 37(7), 25462559.

[22] He, Y., Peng, L., Zhang, D., & Zhao, Z. (2022). Risk analysis via generalized Pareto distributions. *Journal of Business & Economic Statistics*, 40(2), 852-867..

[23] Brenner, M., & Subrahmanyam, M. G. (1988). A simple solution to compute the Implied Standard Deviation. *Financial Analysts Journal*, 44(5), 8083.

Appendix

A Uniswap LP interface

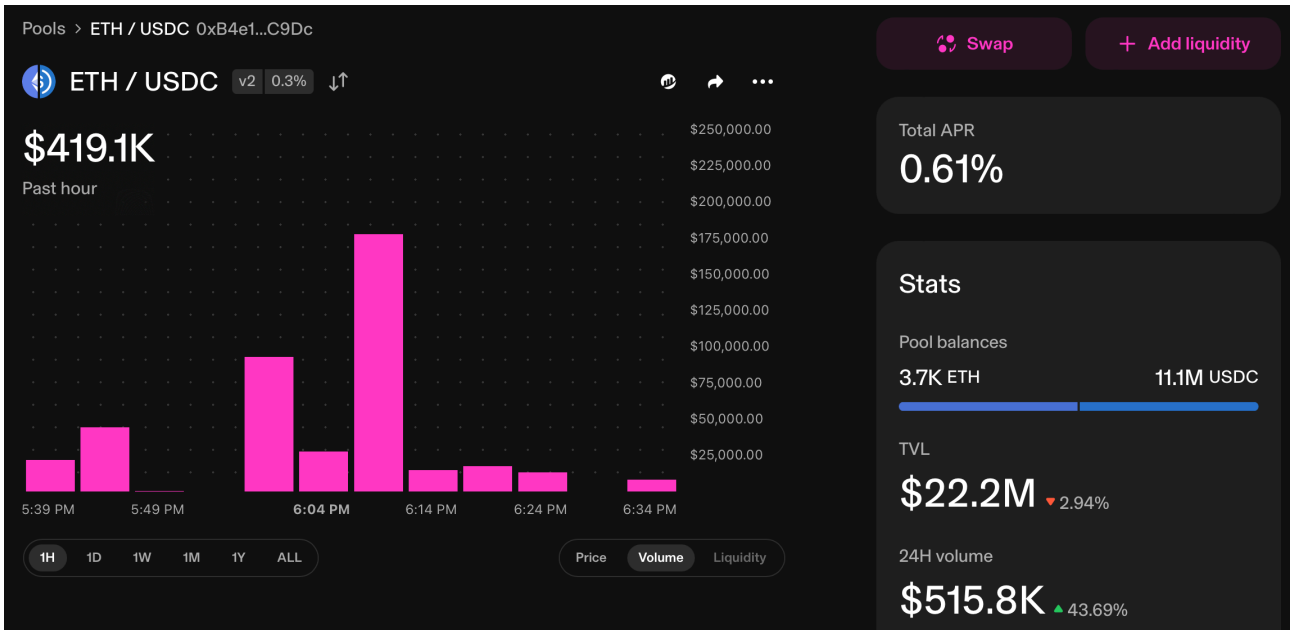

Figure A1: Uniswap V2 web interface for ETH-USDC liquidity pool with current APR.

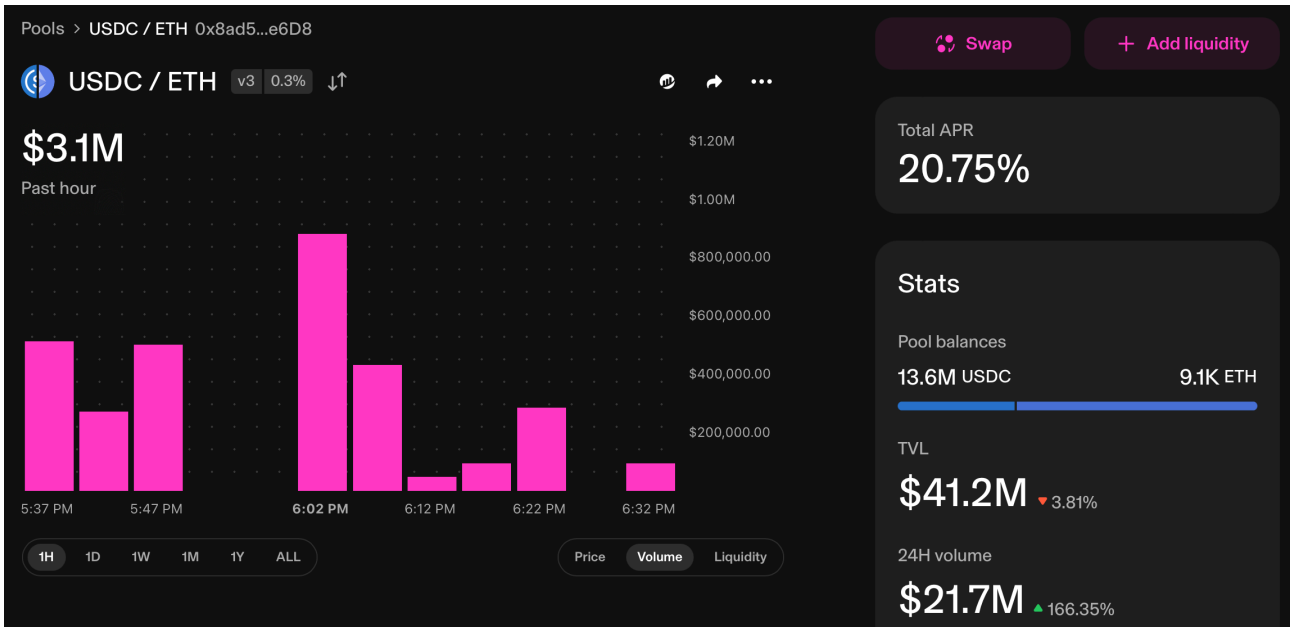

Figure A2: Uniswap V3 web interface for ETH-USDC liquidity pool with current APR.