Aalto University
School of Science

Master’s Programme in Mathematics and Operations Research

Automated Market Makers for Negative Prices: The Impact of Market Depth and Impermanent Loss on Liquidity Provider Profitability

Marlin Jarms

Master’s Thesis
2025

Aalto University
School of Science

Author Marlin Jarms

Title Automated Market Makers for Negative Prices: The Impact of Market Depth and Impermanent Loss on Liquidity Provider Profitability

Degree programme Mathematics and Operations Research

Major Mathematics

Supervisor Prof. Ahti Salo

Advisors Dr. Christian Finke, M.Sc. Dominik Zehe

Date 31 December 2025 **Number of pages** 59 **Language** English

Abstract

This thesis investigates the design of automated market makers (AMMs) for trading tokenized derivatives in decentralized finance. Motivated by the limitations of existing AMMs, which only permit strictly positive prices, we introduce an invariant that also allows for negative prices. As a primary use case, the thesis develops an AMM for trading an offset token against tokenized euros.

The thesis employs, on the one hand, a Monte Carlo–based simulation to evaluate the risk-adjusted returns of an AMM. The simulation includes two types of traders: arbitrage traders, who exploit price deviations between the AMM and the fair value, and noise traders, who represent demand for liquidity. On the other hand, we introduce KPIs such as impermanent loss and market depth.

The goal of the thesis is to analyse whether these KPIs can be used to predict the risk-adjusted returns of an AMM.

Keywords AMM, Impermanent Loss, Market Depth, Geometric Brownian Motion, Invariant Design, Arbitrage Trader, Noise Trader

Contents

Abbreviations and Notation

Abbreviations

AMM	Automated Market Maker
CI	Confidence Interval
DEX	Decentralized Exchange
DLT	Distributed Ledger Technology
GBM	Geometric Brownian Motion
IL	Impermanent Loss
KPI	Key Performance Indicator
LP	Liquidity Provider
MD	Market Depth
SR	Sharpe Ratio
Std	Standart Deviation

Notation

α	Offset token traded in the AMM, whose fair value may be positive or negative.
β	Tokenized euro used as the reference currency in the AMM.
$\mathcal{D}$	Set of days in the simulation horizon.
$\mathcal{T}$	Set of discrete time points at which trades may occur.
N	Time from emission to maturity of the offset token, measured in days.
(R_α, R_β)	Reserve state of the AMM, representing the quantities of α and β held by the pool.
$(R_\alpha(t), R_\beta(t))$	Reserve state of the AMM at time point $t \in \mathcal{T}$.
$(R_\alpha(v_\alpha), R_\beta(v_\alpha))$	Reserve state at which the AMM spot price of α equals v_α (i.e. $p(R_\alpha) = v_\alpha$).
$p(R_\alpha)$	Spot price of α quoted by the AMM as a function of the reserve level R_α.
$\bar{p}(R_\alpha, \Delta_\alpha)$	Average execution price of a trade of size Δ_α.
$(\Delta_\alpha, \Delta_\beta)$	Changes in AMM reserves induced by a single trade.
$(\Delta_\alpha(t), \Delta_\beta(t))$	Changes in AMM reserves due to the trade executed at time point t.
$\delta(t)$	Indicator function taking value 1 if a trade is executed at time $t \in \mathcal{T}$, and 0 otherwise.
κ	Fee parameter determining the transaction fee charged per traded offset token.
τ	Price sensitivity parameter used in the definition of market depth.
m	Mean trade size of noise traders.
s	Log-standard deviation of noise-trader trade sizes.
λ	Expected number of noise-trader arrivals per day.
ε	Price sensitivity of noise traders with respect to deviations from fair value.
$\mathcal{T}_{\text{arbitrage}}$	Set of time points at which the arbitrage trader executes a trade.
$\mathcal{T}_{\text{noise}}$	Set of time points at which noise traders execute trades.
$\mathbb{V}_\alpha$	Set of all possible fair-value time series of the offset token α.
X^{profit}	Random variable representing the (risk-adjusted) total profit of the AMM over a simulation path.
X^{revenue}	Random variable representing the total fee revenue earned by the AMM.
X^{cost}	Random variable representing the total trading cost or loss incurred by the AMM.
X^{IL}	Random variable representing the theoretical impermanent loss induced by the fair-value distribution of α.
$X^{\text{depth}(\tau)}$	Random variable representing the aggregated market depth of the AMM for price sensitivity parameter τ.

1 Introduction

1.1 Motivation, Market Context, and Blockchain Fundamentals

Distributed-ledger technology (DLT) has evolved into critical financial infrastructure, handling trillions of euros each year. Decentralized exchanges (DEXs) alone processed more than \$1.76 trillion in trading volume in 2024, an increase of 89 % year-on-year [1]. These figures illustrate that on-chain liquidity provision is no longer a niche activity but is rapidly becoming a core component of global capital markets.

A blockchain can be understood as a shared database that is maintained and updated by many computers simultaneously. All participants observe the same version, and once information is added, it cannot be altered. This enables strangers to agree on a common record without trusting a single authority.

On top of this, there are so-called *Smart Contracts*, which are programs that can modify the shared database. This makes it possible to create an on-chain exchange venue. When executed, multiple computers must run the program in parallel to ensure that changes to the shared database are correct. Consequently, executing data-intensive and computationally demanding programs is expensive and therefore undesirable.

Tokens are digital objects recorded on the blockchain. They can represent almost anything. Examples from finance include cryptocurrencies, fiat currencies such as USD or EUR, financial instruments such as bonds or shares, and derivatives such as swaps or futures. They can also represent non-financial items, for example the ownership of a picture, a flight ticket, or a physical book. Ultimately, tokens themselves are also smart contracts that encode the rules governing their own creation, ownership, and transfer. In chapter 2 we describe that in more detail.

In this thesis, we consider two types of tokens. The first, denoted by α, is an *offset token* that has properties similar to a futures contract in traditional finance (offset token, futures, and the similarities and differences between them are explained in greater detail in Section 2.4). The second, denoted by β, represents tokenised currency, in our case the euro. This notation (α, β) will be used consistently throughout the thesis.

1.2 From Order Books to Automated Market Makers

Early on-chain exchanges attempted to replicate the order book used in traditional finance. An order book is a list of outstanding buy and sell offers; whenever the highest bid meets the lowest ask, a trade occurs. While this approach works well in traditional finance, implementing it as a smart contract has significant drawbacks: it is highly data-intensive and computationally demanding, making it expensive and generally undesirable.

Automated Market Makers (AMMs) have been introduced as a more computationally efficient alternative. Instead of collecting and matching individual orders, an AMM holds reserves of two tokens, α and β. These reserves must always satisfy a predefined mathematical rule, known as the *invariant* of the AMM, if the reserves hold the invariant, they are said to be in a valid state. Anyone can trade against these reserves: when token α is deposited, the invariant determines the corresponding amount of token β that is exchanged. As a result of the trade, the reserves of tokens α and β are updated.

The corresponding amount of β the trader receives, and the number of β per α (i.e., the *AMM-quoted price* of α), is determined by the current reserves and the invariant (concepts such as price and fair value are discussed in Section 2.2). Each trade changes the reserves and adjusts the price for future trades. Figure 1 illustrates the process of how the AMM reserves change when two traders consecutively sell to the AMM.

This mechanism ensures that trades can be executed immediately without waiting for a matching order.

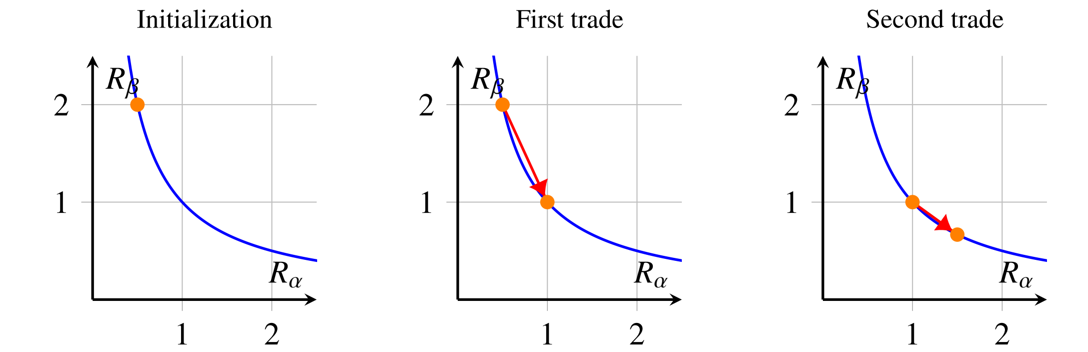

Start: The invariant (blue line) covering all possible states of the AMM. The initial reserve is $(0.5\alpha, 2\beta)$.

First trader sells 0.5α, he receives 1β. AMM reserves changed to $(1\alpha, 1\beta)$. The average price of α is 2β.

Second trader sells 0.5α, he receives only 0.33β. AMM reserves changed to $(1.5\alpha, 0.67\beta)$.The average price of α is 0.66β.

Figure 1: Visualization of AMM initialisation to trade α against β. It shows that the trader gets more per α when the AMM has fewer α tokens compared to when it has more α tokens.

1.3 Use Cases, Research Questions, and Methodology

The mechanisms described above illustrate how AMMs enable trading of tokens without relying on traditional order books or intermediaries. They form a central building block of today's decentralised finance (DeFi) ecosystem.

In recent years, decentralised finance has expanded beyond cryptocurrencies. One notable example is the issuance of a €60 million tokenised bond on a public blockchain in 2023 [2]. In this context, tokenisation means that investors did not hold a traditional bond certificate but instead received tokens representing it. Ownership of these tokens conveys the right to receive coupon payments and repayment of principal at maturity. This concept can be extended to derivatives. Owning tokens from a tokenised derivative means that, depending on the underlying's performance, the holder is either required to pay or eligible to receive money.

This thesis considers a financial setting in which not only bonds but also derivatives such as swaps, futures, and options are tokenised.

In this setting, it is desirable to enable trading of such tokenised derivatives on an entirely on-chain decentralised exchange, as this would allow market participants to transact without relying on intermediaries. Yet designing a suitable market mechanism faces two key challenges. First, a fully on-chain implementation of an order book is data-intensive and computationally expensive. Second, the current methodology of invariants restricts the average price of a trade to strictly positive values.

To address these challenges, this thesis introduces invariants that allow the price quoted by an AMM to be negative. Figure 2 shows how an invariant can produce negative prices.

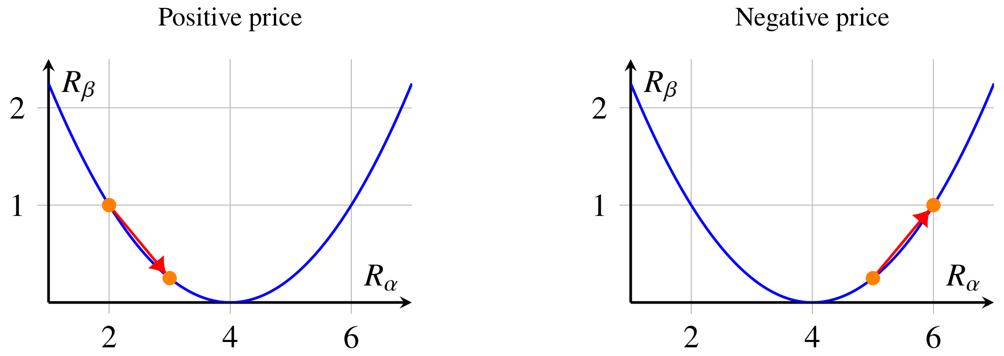

Initial reserves are $(2\alpha, 1\beta)$. First trader sells 1α, he receives 0.75β. AMM reserves changed to $(3\alpha, 0.25\beta)$. The average price of α in this trade is 0.75β.

Initial reserves are $(5\alpha, 0.25\beta)$. First trader sells 1α, additionally needs to pay 0.75β. AMM reserves changed to $(1\alpha, 1\beta)$. The average price of α in this trade is -0.75β.

Figure 2: Visualization of negative and positive prices.

As the primary use case, we develop an AMM that holds the two tokens: The offset token α and the tokenized euro β.

The objective of this thesis is to answer the research question: *How can Automated Market Maker (AMM) models be designed to support negative prices, and to what extent do aggregated market depth and impermanent loss determine the risk-adjusted profitability of liquidity providers of AMMs that allow negative prices?* The first part is already partially answered, but we will explore it further and introduce a class called the *Power Sum Invariant*.

We address the second part through a five-step process (as illustrated in Figure 3). First, we draw 256 invariants from the Power Sum Invariant class (introduced in section 3.3).

In the second step, we simulate 100 fair-value time series of α, where a fair-value time series is defined as a sequence of simulated fair-value estimates of an asset with one value per day over a one-year horizon. The dynamics are generated using a modified version of geometric Brownian motion (GBM), a stochastic process commonly used to model asset prices. In our setting, GBM is employed to simulate the fair value of the underlying asset and, through the offset, the fair value of the offset token.

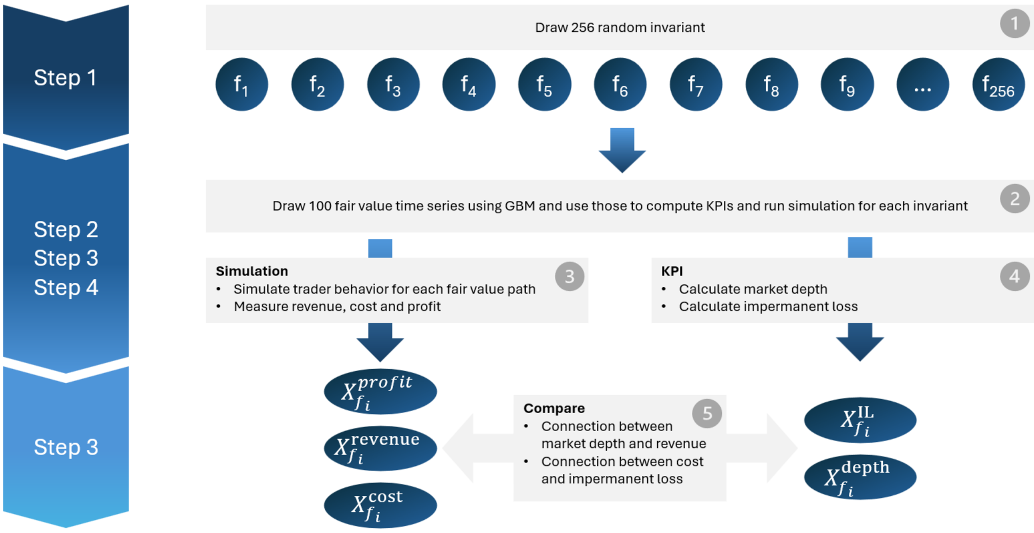

Figure 3: Process steps to evaluate the connection between KPIs and the simulation: Draw 256 invariants; run the simulation and calculate KPIs for each invariant; compare the results. Note: $X_{fi}^{profit} = X_{fi}^{revenue} - X_{fi}^{cost}$

In the third step, we set up an agent-based trading simulation in which traders interact with the AMM according to predefined rules by buying and selling α against β. The agent-based model consists of an arbitrage trader and noise traders. The arbitrage trader seeks to maximize profit by exploiting the difference between the AMM-quoted price and the fair value of α. The noise trader represents liquidity demand and trades for reasons unrelated to arbitrage.

In the fourth step, we use these Fair Value Time Series to estimate key performance indicators such as impermanent loss (i.e., the opportunity cost LPs face when providing liquidity compared to holding the offset token) and market depth (i.e., how many offset tokens a trader can buy or sell such that the average price of that trade remains within a given range), based on the market environment.

Finally, in the fifth step, we analyse the extent to which profit and risk-adjusted return are determined by impermanent loss and market depth.

1.4 Thesis Structure

The remainder of this thesis is structured as follows. Chapter 2 presents the fundamentals for Automated Market Makers (AMMs) that can support both positive and negative prices. It introduces the conceptual design of the offset token and illustrates the underlying mechanics through numerical examples. Chapter 3 formalises the AMM framework by defining the mathematical structure, including invariants, reserve dynamics, axioms, and the specific invariant family analysed in this work. Chapter 4 presents the simulation environment. It describes the market setup, trader models, and the stochastic process used to generate the fair value time series that drive the trading dynamics. Chapter 5 defines the key performance indicators used to evaluate AMM behaviour, with a focus on impermanent loss and market depth. The chapter also discusses their theoretical relationship to revenue and cost. Chapter 6 reports the simulation results. It compares simulated outcomes with the KPIs and proposes a method for predicting risk-adjusted profitability based on these indicators.

Finally, Chapter 7 summarises the contributions of the thesis, discusses limitations, and outlines potential directions for future research on AMM design for derivative tokens.

2 Fundamentals of Automated Market Makers

This chapter aims to provide intuition for AMMs and, in particular, AMMs that allow negative prices. Throughout this chapter, we neglect fees, as the goal is to provide intuition for how AMMs with negative prices work.

2.1 Motivation of the Invariant

To illustrate the behavior of an AMM, we use an invariant f that allows negative prices. Let A be an AMM with invariant f. Generally, it is clear from context that a function f is the invariant of a given AMM, if there is uncertainty, we write f_A to clarify that it is the invariant of AMM A. The trader buys or sells Offset Tokens (denoted by α) against euros (denoted by β). Hence, A holds reserves of α and β. Let R_α be the amount of α, and R_β the amount of β held by the AMM A, implying

$$R_\beta = f(R_\alpha).$$

The first challenge is designing an invariant that allows both negative and positive prices. An AMM is typically designed such that if the fair value of α decreases, the AMM ends up holding more α. To limit downside exposure, the invariant should ensure that the AMM has a theoretical maximum amount of α. Let this upper bound be 10. The invariant needs to be constructed so that it becomes infinitely expensive to move R_α to 10, just as it should be infinitely expensive to drain the pool (i.e., to move R_α to 0).

The invariant

$$f_1(R_\alpha) = \frac{1}{10 - R_\alpha}$$

ensures that R_α always remains below 10. The invariant

$$f_2(R_\alpha) = \frac{1}{R_\alpha}$$

ensures that R_α always remains above 0. Adding these two functions yields the invariant

$$f(R_\alpha) = f_1(R_\alpha) + f_2(R_\alpha) = \frac{1}{R_\alpha} + \frac{1}{10 - R_\alpha},$$

which combines both properties. We refer to the property that R_α can never reach the lower bound 0 nor the upper bound 10 as the *boundary behavior*.

Additionally, we want the price of α to be low when the AMM holds a large amount of α, and high when it holds only a small amount. Figure 4 illustrates this behavior, which we formalize in Chapter 3. One final remark regarding the invariant f: in the right-hand plot of Figure 4, we observe that f is defined so that the reserves always contain at least 0.4β, which cannot be removed. Therefore, we slightly modify the invariant by subtracting a constant $c = 0.4$, giving

$$f(R_\alpha) = \frac{1}{R_\alpha} + \frac{1}{10 - R_\alpha} - 0.4.$$

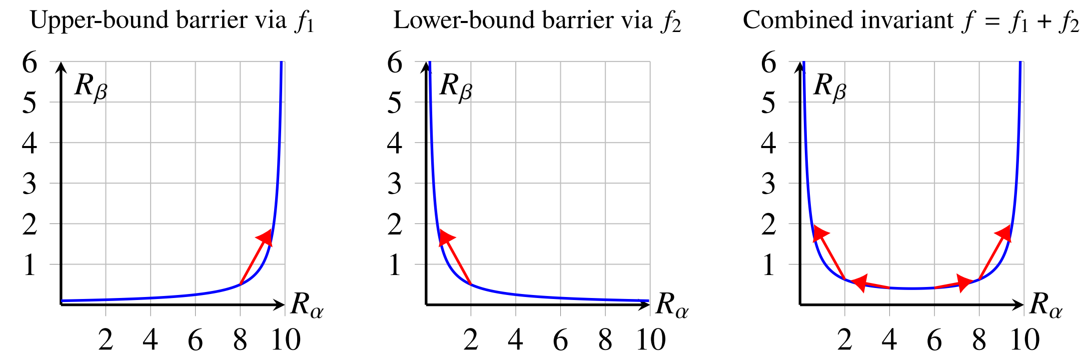

Idea: As R_α approaches 10 from the left, f_1 tends to ∞.
Effect: Selling α becomes prohibitively expensive near 10, so the AMM state stays < 10.

Idea: As R_α approaches 0 from the right, f_2 tends to ∞.
Effect: Buying α becomes prohibitively expensive near 0, so the AMM state stays > 0.

Idea: Summing f_1 and f_2 adds both barriers.
Effect: The α reserve is always in the interval $(0, 10)$ with increasing cost near both ends.

Figure 4: Boundary-enforcing invariants. The combined invariant on the right combines the properties of discouraging approaching $R_\alpha \to 0$ and $R_\alpha \to 10$.

2.2 Price, Fair Value, and AMM Price Discovery

In this thesis, we frequently refer to both the price of α and the fair value of α. Because the framework includes tokens that may have negative prices, this section introduces these concepts and explains how prices behave in AMMs compared with order-book markets.

Fair Value and Price

A definition of price, according to the IMF, is *"A price is the amount of money a buyer gives a seller in exchange for a good or a service."* [3]. Based on this, we define the fair value of a financial instrument (which is an economic good) as the price at which it would trade in a perfectly efficient market. In this thesis, if α has a fair value of 1, we assume that a trader can buy and sell α for 1 in an external market in arbitrarily large quantities. By external market, we mean a trading venue that is not the AMM. Price impact in that external market is ignored.

Throughout the thesis, when we refer to the *price* of a token, we mean the *AMM price*. This is the amount of money required to purchase a token from an AMM, determined solely by the AMM's invariant, its reserves, and the trade size. It does not depend on the fair value. However, as discussed in Section 4.2.2, arbitrage traders continually realign the AMM price with the fair value.

Prices are measured in $\frac{\beta}{\alpha}$. For example, we may say that *the price of α is 2β* or *the fair value of α is 2β*. When the context is clear, we omit β and simply write that the price or fair value is 2.

Negative Prices

Negative prices arise when holding an asset consists of both rights and obligations. If the obligation exceeds the expected revenue, the net value becomes negative; if the expected revenue exceeds the obligation, the value becomes positive.

Formally, let an asset generate an expected cash flow R and require fulfilling an obligation of size L. The fair value is

$$\text{Fair Value} = R - L.$$

If $R < L$, the value is negative, meaning holders must be compensated for taking on a net liability. Such situations occur when the obligations associated with an asset outweigh the rights it confers.

Spot Price

When discussing the AMM price so far, we have considered the *average trade price* for α, which depends not only on the current state of the AMM but also on the trade size. Larger purchases of α are more expensive due to the shape of the invariant. To obtain a unique metric that reflects the AMM price at a specific moment, we introduce the *spot price*—the average price of α in a trade in which an infinitesimally small amount of α is bought or sold (see Figure 5).

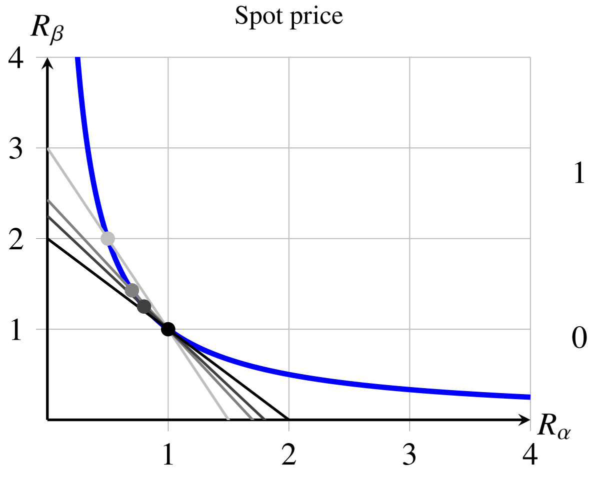

Figure 5: The average price of a trade can be interpreted as the negative slope of the secant line connecting the old state and the new state (grey). If we let the trade size approach 0, the secant becomes a tangent (black) and its negative slope is the spot price.

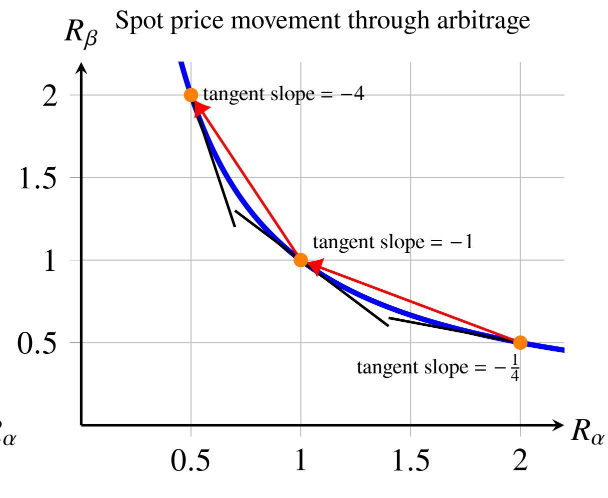

Figure 6: Fair value changes from $p_{\text{old}} = \frac{1}{4}$ to $p_{\text{new}} = 4$. Arbitrage moves the AMM state along $y = \frac{1}{x}$ from $(2, 0.5)$ to $(0.5, 2)$ until the instantaneous price (negative tangent slope) equals 4.

Price Discovery and Fair Value Adjustments

We now analyze how changes in the fair value of α affect the spot price in an AMM compared with an order-book exchange. The argument relies on rational trader behavior, as outlined in the introduction.

Suppose the fair value of α increases from p_1 to p_2. In an order-book exchange, rational sellers—observing the new fair value—will not accept prices significantly below p_2. All outstanding sell orders below p_2 are therefore canceled, and the spot price jumps immediately to p_2.

An AMM, however, does not observe the fair value. It quotes prices solely based on its reserves and invariant. A rational arbitrageur buys α as long as the AMM's spot price lies below p_2. Each arbitrage trade increases the spot price according to the invariant until it converges to p_2 (see Figure 6). At that point, further arbitrage profits disappear.

Thus, unlike order-book markets, AMMs cannot adjust instantaneously to changes in fair value. At least one arbitrage trade is required to realign the AMM spot price with the new fair value. This mechanism corresponds to the arbitrage-trader behavior introduced in Section 4.2.2.

2.3 Illustrative Example

To illustrate the mechanism of the invariant

$$f(R_\alpha) = \frac{1}{R_\alpha} + \frac{1}{10 - R_\alpha} - 0.4,$$

we go through the first few steps in the lifecycle of the AMM. Initially, the fair value of α (i.e., the offset token) is 0, so the LPs initialize the reserves such that the spot price is 0. This is the case for

$$(R_\alpha, R_\beta) = (5, 0).$$

Now a trader wants to sell 1α, i.e., $\Delta_\alpha = 1$. To determine the amount Δ_β of β that the trader pays or receives, we consider the state after the trade is executed. The state after the trade is $(R_\alpha^{\text{new}}, R_\beta^{\text{new}})$, with

$$R_\alpha^{\text{new}} = R_\alpha + \Delta_\alpha = 5+1 = 6, \quad R_\beta^{\text{new}} = f(R_\alpha^{\text{new}}) = \frac{1}{6} + \frac{1}{10-6} - 0.4 = \frac{1}{6} + \frac{1}{4} - 0.4 \approx 0.0167.$$

Hence,

$$\Delta_\beta = R_\beta^{\text{new}} - R_\beta \approx 0.0167 - 0 = 0.0167.$$

The trader gives 1α to the AMM and receives approximately -0.0167β, meaning they effectively pay 0.0167β to the AMM. The average price of α in this trade is therefore

$$\frac{-0.0167}{1} = -0.0167.$$

We see that the trader needs to pay more compared to selling all of α at the current spot price; this cost is called slippage.

Now the fair value of α changes to -1. The current state of the AMM is

$$(R_\alpha, R_\beta) = (6, 0.0167).$$

The effect described in Section 2.2 now takes place. The AMM state where the spot price equals -1 is approximately at $R_\alpha^{\text{new}} \approx 9$. Thus, an arbitrage trader would buy

$$\Delta_\alpha = R_\alpha^{\text{new}} - R_\alpha = 9 - 6 = 3$$

tokens of α from an external market. The arbitrage trader receives 3α and 3β from the external market. The amount Δ_β the arbitrage trader pays to the AMM is determined through

$$R_\beta^{\text{new}} = R_\beta + \Delta_\beta,$$

and R_β^{new} must satisfy

$$f(R_\alpha^{\text{new}}) = R_\beta^{\text{new}}.$$

Thus,

$$\Delta_\beta = f(R_\alpha^{\text{new}}) - R_\beta = \left(\frac{1}{9} + \frac{1}{1} - 0.4\right) - 0.0167 \approx 0.689.$$

This means the arbitrage trader makes a profit of 2.311β.

The LPs initializing the AMM held 5α and 0β. Had they simply held these tokens, their position would now be worth -5β (called the *HODL strategy*[1]). Instead, as liquidity providers, they now have 9α and 0.689β, which is worth -8.311β (the *LP strategy*). The difference between the HODL strategy and the LP strategy is called *impermanent loss*, previously described as the opportunity cost of providing liquidity. The LPs' portfolio is now worth 3.311β less than the HODL alternative. It is called impermanent loss because when the fair value of α returns to 0, the loss disappears (that is also the reason why it is called impermanent loss, as the loss is impermanent and can disappear).

2.4 Comparison between Traditional Finance and Decentralized Finance

We introduced the offset token as a token that can have a negative fair value and described it as a financial instrument similar to a future. We begin by explaining a futures contract and then introduce a detailed description of what an offset token is, highlighting similarities and differences between an offset token and a future. In the second part, we compare how Automated Market Makers relate to market makers in traditional finance and how AMMs compare with order-book-based trading venues in traditional finance.

[1] HODL: Originally a typo for "hold", now used as an acronym for "hold on for dear life", i.e., a buy-and-hold strategy

Futures

Futures are standardized derivative contracts traded on organized trading venues. They specify the obligation to buy or sell a particular underlying asset at a predetermined price on a defined future date [4]. This date is referred to as the maturity of the contract. A market participant who is obligated to buy the underlying at maturity holds a long position, while a participant who is obligated to deliver the underlying holds a short position. The underlying asset can be a commodity such as oil, silver, or gold; however, in this thesis, we focus on shares as the underlying asset.

At maturity, the long side is required to pay the predetermined price regardless of the current market price of the underlying, and the short side is required to deliver the underlying asset. For an underlying asset without dividends, a theoretical futures price is given by

$$F_t = S_t(1 + r_d)^{T-t},$$

where S_t is the price of the underlying at time t, r_d is the risk-free daily interest rate, $T - t$ denotes the remaining time to maturity in days, and F_t is the futures price at time t. This expression describes the theoretical pricing relation; in practice, the market futures price is determined by supply and demand.

In practice, however, futures are normally not settled by physical delivery of the underlying for most contracts. Instead, the contract is settled through cash payments. If the future is entered at time t, the payoff for a long position at maturity T is

$$S_T - F_t,$$

and for a short position, it is

$$F_t - S_T.$$

Thus, if $S_T > F_t$, the long position receives a payment from the short position, and if $S_T < F_t$, the short position receives a payment from the long position.

If futures were settled only at maturity, there would be counterparty risk: the gain of one party can become very large if the final underlying price differs substantially from the agreed futures price, potentially causing the losing party to default. To mitigate this, two mechanisms are used: first, a clearing house acts as the counterparty to all trades, and second, margin accounts with daily settlement are used.

A *clearing house* stands between buyers and sellers by becoming the counterparty to both sides of the trade. This eliminates direct counterparty exposure between participants and ensures that contractual obligations are fulfilled.

A *margin account* is a dedicated account that holds collateral for each futures position. When a position is opened, an initial margin is deposited, and the account is subsequently adjusted based on daily price movements.

Daily settlement, also known as *marking-to-market*, means that the futures position is revalued every trading day. For a long position opened at time t, the daily change to the margin account on day $t + 1$ is

$$F_{t+1} - F_t,$$

and for a short position it is

$$F_t - F_{t+1}.$$

If the future reaches maturity on day $t+1$ (i.e., $t+1 = T$), then $F_T = S_T$, so the futures price at maturity equals the fair value of the underlying asset.

Through this mechanism, gains and losses are realized continuously, which ensures that obligations are met and reduces counterparty risk.

Offset Token in Comparison to Futures

In the introduction, we presented the offset token as a token that can have a negative fair value. Owning an offset token is similar to buying a futures contract; however, the offset token is a hypothetical tokenized derivative. This section discusses how the infrastructure around an offset token could look. Throughout this thesis, it is sufficient to think of the offset token as an instrument whose fair value can be both positive and negative, with an oracle supplying the daily fair value. Arbitrageurs are assumed to be able to buy and sell the token at this fair value at another exchange venue. Nevertheless, to motivate the need for an AMM design that supports negative prices, it is useful to describe the mechanics of such a token more concretely.

The offset token has some underlying asset U and represents the performance of a 100€ investment into that underlying asset. We call the 100€ the notional of the offset token and let $v_U(t)$ be the fair value of that 100€ investment at time t, so $v_U(0) = 100$ is the notional of the offset token (we also denote it by v_0). Similar to futures contracts, the offset token has a predetermined date T at which the offset token reaches maturity. Once the offset token has matured, owners of offset tokens are entitled to a payout of

$$v_U(T) - v_U(0).$$

If the payout is negative, then owners are required to pay the difference. Owning an offset token is similar to owning the underlying asset and borrowing the money required to buy the underlying asset at day 0. The fair value of an offset token at time t is therefore given by

$$v_\alpha(t) = v_U(t) - v_U(0)(1 + r_d)^{T-t},$$

where T is the maturity date, t is the current day, and $v_\alpha(t)$ denotes the fair value of the offset token (we use α to denote the offset token). For simplicity, we assume in this thesis that the risk-free interest rate is 0.

For futures contracts, there is a long position, which is an obligation to buy the underlying at maturity. Holding that long position via a futures contract is similar to owning an offset token. Analogously, there also needs to be a short position: we call this the negative offset token. There is always a one-to-one connection between a negative offset token and an offset token; we say that a negative offset token has exactly one offset token as counterparty. The negative offset token has the same maturity as the offset token. The payout at maturity is

$$v_U(0) - v_U(T),$$

which is the negative of the offset token's payout. So if $v_U(T) > v_U(0)$, then the owner of the offset token (long position) receives a payment from the owner of the negative offset token (short position), and if $v_U(T) < v_U(0)$, the short position receives a payment from the long position.

The way we have designed the offset token so far, it involves counterparty risk. Similar to futures contracts, a clearing house would be used as counterparty for the long and short positions. The clearing house would need to design mechanisms such as margin accounts or other collateral requirements to ensure that owners of (negative) offset tokens have sufficient funds to cover the amounts they may need to pay.

Figure 7 shows how the fair value of the offset token behaves in comparison to the fair value of the underlying.

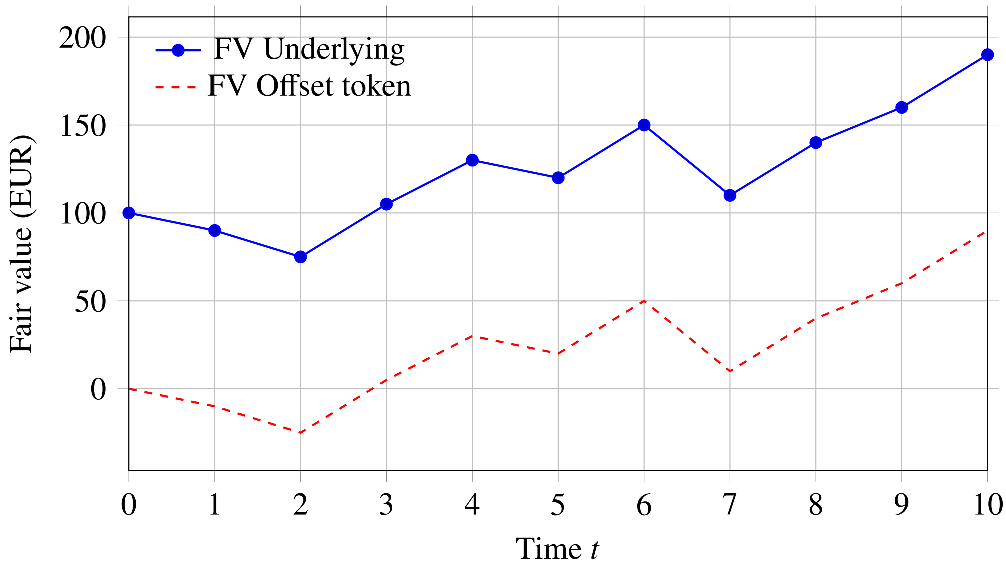

Figure 7: Example timeseries of the fairvalue of the underlying and the fairvalue of the offset token, which is exactly the fairvalue of the underlying minus the fairvalue of the underlying at $t = 0$ (i.e. 100 EUR).

Traditional Trading Venues and the Role of Market Makers

Traditional trading venues organize the interaction of buyers and sellers through a *Central Limit Order Book* (CLOB). The order book is the core mechanism for price formation and trade execution. It aggregates all outstanding buy and sell orders submitted by market participants and ranks them according to their bid and ask prices. Buy orders are sorted in descending order, with the highest bid at the top of the book, while sell orders are listed in ascending order, with the lowest ask at the top.

When a new order arrives, the matching engine checks whether it can trade with the best prices on the other side of the market. If it can, the order is executed immediately and removed from the book. If not, it remains in the book as a limit order until it is filled, cancelled, or expires.

Within this structure, *market makers* play a crucial role. These entities, often specialized trading firms or broker-dealers, continuously quote both buy and sell prices for selected financial instruments. By posting limit orders on both sides of the market, they provide liquidity to other participants and help maintain a narrow bid–ask spread. Their activity reduces execution uncertainty for traders who demand immediate liquidity and stabilizes prices by absorbing temporary imbalances in order flow. To manage their inventory and mitigate risk, market makers dynamically adjust their quotes in response to market conditions, volatility, and their own position exposure.

How Automated Market Makers Combine the Functions of Trading Venues and Market Makers

Automated Market Makers (AMMs) replace both the traditional trading venue infrastructure and the role of professional market makers through a fully algorithmic mechanism. Instead of organizing orders in a central limit order book, AMMs rely on liquidity pools that hold reserves of two or more assets. Traders interact directly with these pooled reserves, without the need for a matching engine or a counterparty submitting a corresponding order.

The pricing mechanism of an AMM is different from that of a traditional trading venue, as seen in Section 2.2. Nevertheless, it still allows traders to buy and sell assets immediately without waiting for a matching order. The reason is that the AMM also fulfills the role traditionally held by market makers. Instead of relying on specialized firms to post quotes, liquidity is supplied by users who deposit assets into the pool. These liquidity providers collectively serve as the market maker: the pool stands ready to buy and sell, with the price determined by the invariant of the AMM. In return for providing liquidity, liquidity providers earn compensation in the form of trading fees. Arbitrage traders play an indirect but essential role by ensuring that the AMM's price remains aligned with external markets.

Through this combination, AMMs unify two distinct functions—trade execution and liquidity provision—into a single, protocol-driven system that operates without order books, quote providers, or centralized matching infrastructure.

2.5 Distributed Ledger and Smart Contracts

Based on [5], we now examine how blockchains work and what it means to be a distributed ledger.

Blockchain as Distributed Ledger

A blockchain can be understood as a distributed ledger in which all transactions are stored transparently, permanently, and in a way that can be verified by all participating nodes. Instead of relying on a central authority to validate and maintain the ledger, this responsibility is shared among many independent participants in the network. Each node keeps its own copy of the ledger and verifies new transactions according to a common set of rules.

Once a collection of valid transactions has been gathered, it is bundled into a block. Each block contains a cryptographic reference to the previous block, forming a chain of interconnected records. This structure ensures that past data cannot be altered without breaking the cryptographic links between blocks, making manipulation easily detectable. As a result, the blockchain provides a decentralized and resilient method for maintaining a trustworthy ledger without the need for a single controlling entity.

Table 1 shows an example collection of transactions, and Figure 8 illustrates how the balances of the accounts change as a result of these transactions.

Number	Action
1	A transfers 10 Tokens to B
2	B approves A to transfer 10 Tokens from B to any address
3	A transfers from B 10 Tokens to A

Table 1: A minimalistic example of entries in a distributed ledger

Account	Balance
A	10
B	0

(a) Initial state of the balances

Account	Balance
A	0
B	10

(b) After execution of first statement

Account	Balance
A	0
B	10

(c) After execution of second statement

Account	Balance
A	10
B	0

(d) After execution of third statement

Figure 8: Reserves for different accounts are changing based on the transactions from Table 1, assuming the initial balance for account A is 10 and for account B is 0. The actual balances are not stored directly on the blockchain.

Smart Contracts

Up to this point, we have simply assumed the existence of a token that can be sent and received. Blockchains such as Ethereum have a native currency—Ether in the case of Ethereum—and, as shown in the previous section, the blockchain stores the actions (the distributed ledger) from which account balances can be derived.

Blockchains like Ethereum also allow users to create their own tokens by deploying smart contracts that specify the behavior of those tokens. Figure 21 shows a minimalistic example. ERC20 is a standard for fungible tokens on the Ethereum blockchain, defining a uniform interface for balance management, token transfers, and delegated spending. Fungible tokens are interchangeable units of equal value, such as tokens or currencies. Non-fungible assets, by contrast, represent unique items whose value cannot be exchanged one-to-one, such as artwork or a flight ticket with a specific seat. In this thesis, however, we consider only fungible tokens. Its fixed set of functions and events ensures interoperability between wallets, smart contracts, and decentralized applications [6]. Although many token standards exist, ERC20 is a simple standard illustrating how one might create a digital bond or a stablecoin.

Smart contracts can be used for far more than token creation. One important use case is the development of decentralized exchange venues. Figure 20 provides pseudocode describing how the code of an AMM could look. When a trader buys or sells some amount of token α, the AMM first checks its current reserves. It then computes the amount of β the trader must pay or receive according to the invariant, and also computes the applicable fee. Finally, the AMM collects or pays out the required amounts of α and β to or from the trader, and forwards the fee to the liquidity provider.

For example, suppose a trader wants to sell 5α, and according to the invariant, this requires the trader to pay 4β (i.e., the price of α is negative). Then the trader must call the `approve` function of α to allow the AMM to collect 5α, and must call the `approve` function of β to allow the AMM to collect 4β.

In the example code of the AMM we omitted several aspects to keep the presentation simple. The pseudocode is only intended to illustrate the structure of a trade. For completeness, the omitted elements include necessary security checks such as verifying whether $f(R_\alpha) = R_\beta$ (in the example code we used rAlpha for R_α and rBeta for R_β), the constructor in which the liquidity provider, the token addresses for α and β, and the fee parameter would typically be specified, as well as standard protections against overflow or underflow and other safeguards commonly used in production smart contracts.

3 Automated Market Makers

Building on the conceptual framework developed from Section 2, this chapter provides a formal foundation for Automated Market Makers (AMMs) capable of handling both positive and negative prices. The discussion transitions from qualitative intuition to precise mathematical structure, introducing the formal components that define an AMM and specifying the assumptions required for consistent price formation.

An AMM can be described by four fundamental elements: its invariant, the reserves of the traded tokens, a liquidity level, and a fee mechanism. Together, these determine the price dynamics, the shape of the trading curve, and the incentives for both traders and liquidity providers. In contrast to conventional AMMs, whose invariants implicitly restrict prices to positive domains, the framework developed here generalizes the structure to permit negative price regions.

The remainder of this section proceeds as follows. Section 3.1 introduces the core mechanics of an AMM, defines reserves and valid states, and establishes regularity and convexity assumptions for the invariant. Section 3.2 formalizes the roles of the liquidity provider and the fee structure, including the concept of liquidity scaling. Section 3.3 presents the specific invariant family used throughout this thesis, which is the *Power Sum Invariant*, and demonstrates that it satisfies the required axioms and scaling properties. And finally in Section 3.4 we summarize other literature introducing invariants, that allow negative prices.

3.1 Invariant, Reserves and Core Mechanics

An AMM is an exchange venue where traders can exchange token α against token β. In the setting of this thesis, α represents offset tokens, while β represents the euro.

Definition 3.1 (Reserves and State)**.** An AMM holds reserves of α and β, denoted by

$$(R_\alpha, R_\beta).$$

The pair (R_α, R_β) is called the *state* of the AMM. The *initial state*, i.e., the state before any trades have been executed, is denoted by

$$(R_\alpha^{(0)}, R_\beta^{(0)}).$$

An AMM cannot hold negative reserves. Furthermore, if a token can take on a negative fair value, its reserve must have an upper bound (as discussed in Section 2.1). In the setting considered here, the fair value of α can be negative, while the fair value of β is strictly positive. Let $k > 0$ denote the upper bound for the α reserves. Consequently, the reserves satisfy

$$(R_\alpha, R_\beta) \in (0, k) \times [0, \infty).$$

The constant k is referred to as the *liquidity level*.

Definition 3.2 (Invariant and Valid State)**.** An *invariant* of an AMM is a function $f : (0, k) \to [0, \infty)$ such that a state (R_α, R_β) is *valid* if and only if

$$R_\beta = f(R_\alpha).$$

When the current state of the AMM is (R_α, R_β), a trader may execute a trade $(\Delta_\alpha, \Delta_\beta)$ if and only if

$$R_\alpha + \Delta_\alpha \in (0, k) \quad \text{and} \quad R_\beta + \Delta_\beta = f(R_\alpha + \Delta_\alpha).$$

In this case, $(\Delta_\alpha, \Delta_\beta)$ is called a *valid trade*.

Axiom 1: Regularity

In Section 2.2 we established that the spot price is linked to the derivative of the invariant. To define the spot price rigorously, we therefore require a regularity assumption on f.

Axiom 1 (Regularity)**.** The invariant f is twice continuously differentiable on $(0, k)$, i.e., $f \in C^2((0, k))$.

Definition 3.3 (Average and Spot Price)**.** Given a valid state (R_α, R_β) and a valid trade $(\Delta_\alpha, \Delta_\beta)$, the *average price* of α in that trade is defined as

$$\bar{p}(R_\alpha, \Delta_\alpha) = -\frac{\Delta_\beta}{\Delta_\alpha} = -\frac{f(R_\alpha + \Delta_\alpha) - f(R_\alpha)}{\Delta_\alpha}.$$

The *spot price* is the infinitesimal limit:

$$p(R_\alpha) = \lim_{\Delta_\alpha \to 0} \bar{p}(R_\alpha, \Delta_\alpha) = -f'(R_\alpha).$$

Axiom 2: Strict Convexity

Intuitively, the average price for a small trade should be higher than for a large trade, analogous to classical exchanges. For all $R_\alpha \in (0, k)$ and $\delta_\alpha < \Delta_\alpha$, this requires that

$$\bar{p}(R_\alpha, \Delta_\alpha) < \bar{p}(R_\alpha, \delta_\alpha).$$

This property is equivalent to the strict convexity of f (as shown in [7]), because $\bar{p}(R_\alpha, \Delta_\alpha)$ is the negative secant slope going through the two points

$$(R_\alpha, f(R_\alpha)), \quad (R_\alpha + \Delta_\alpha, f(R_\alpha + \Delta_\alpha)).$$

Axiom 2 (Strict Convexity)**.** The invariant f is strictly convex, i.e., its derivative $f'(x)$ is strictly increasing on $(0, k)$.

Axiom 3: Boundary Behavior

The invariant should be designed such that it becomes infinitely expensive to move the reserve R_α to either boundary point 0 or k.

Axiom 3 (Boundary Behavior)**.** The invariant f satisfies

$$\lim_{x\uparrow k} f(x) = +\infty, \qquad \lim_{x\downarrow 0} f(x) = +\infty.$$

This ensures that

$$\lim_{x\uparrow k} p(x) = +\infty, \qquad \lim_{x\downarrow 0} p(x) = -\infty.$$

Since f is continuous, differentiable, and strictly convex, the derivative $-f'$ is invertible, and its inverse $(-f')^{-1}$ is well-defined on $\mathbb{R}$.

Corollary 3.4 (Summary of Properties)**.** *If f satisfies Axioms 1–3, then:*

1. *f is strictly convex and twice continuously differentiable,*
2. *$p(x) = -f'(x)$ is continuous, strictly monotonic, and bijective,*
3. *The inverse $u_f = (-f')^{-1}$ exists and is defined on $\mathbb{R}$.*

3.2 Fee and Liquidity Provider

In the previous section, we introduced the core mechanics of an AMM primarily from the trader's perspective. We now turn to the role of the liquidity provider (LP) and the fee mechanism.

The initial liquidity provider determines the invariant f (including the liquidity level k) such that it satisfies Axioms 1, 2, and 3. The LP also selects an initial valid state $(R_\alpha^{(0)}, R_\beta^{(0)})$ with $f(R_\alpha^{(0)}) = R_\beta^{(0)}$, ensuring that the spot price matches the current fair value $p(R_\alpha^{(0)}) = v_\alpha(0)$.

Scaled Liquidity

Suppose the initial LP creates an AMM A_1 and decides to use the invariant f. However, later, the LP realizes that he wants to double the liquidity level. One way the LP could achieve that would be to create a new identical AMM A_2 that also has f as invariant and then split every trade evenly across the two identical AMMs A_1, A_2. Rather than operating multiple AMMs, we construct a single AMM B whose invariant g replicates the combined effect. Let (R_α, R_β) denote the reserves of A_1, A_2. Then AMM B should have reserves $(2R_\alpha, 2R_\beta)$, implying that

$$g(2R_\alpha) = 2f(R_\alpha),$$

or equivalently, by setting $x = 2R_\alpha$,

$$g(x) = 2f\left(\tfrac{x}{2}\right).$$

By generalizing this argument, for any scaling factor $c \in \mathbb{Q}_{>0}$,

$$g(x) = cf\left(\frac{x}{c}\right).$$

Hence, we define the liquidity scaling operator as follows:

$$L(c, f)(x) := cf\left(\frac{x}{c}\right). \tag{3.1}$$

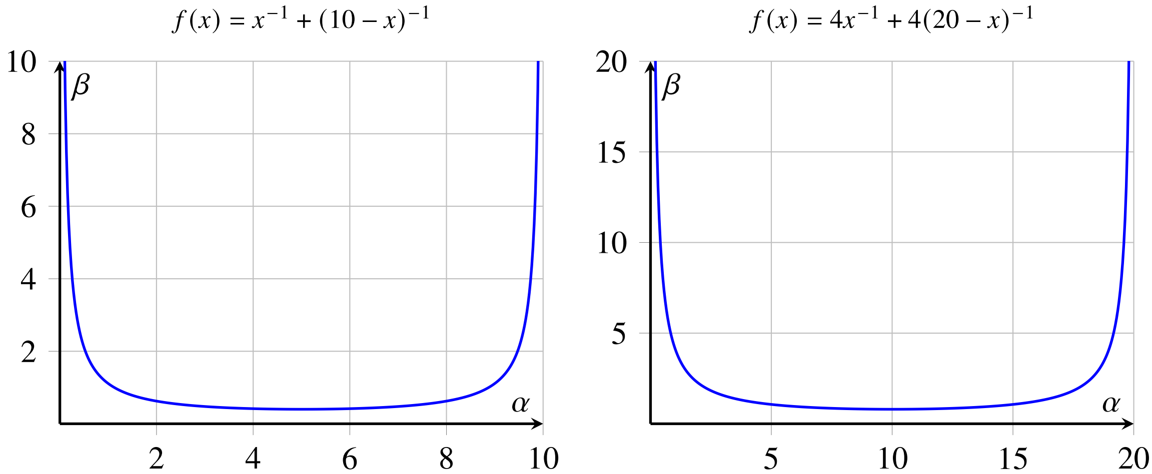

Initial invariant with liquidity level $k = 10$ Invariant with scaled liquidity by factor 2

Figure 9: Same invariant with different levels of liquidity; when the axes are scaled accordingly, the shape of the invariant remains unchanged.

When comparing two AMM which have the same invariant, but one has higher liquidity. For the one AMM with higher liquidity, a single trade has less influence on prices, while with lower liquidity, the same trade causes larger price changes, as seen in Figure 10.

Definition 3.5 (Scaling Property)**.** Let $(f_k)_{k\in\mathbb{Q}}$ be a family of invariants satisfying Axioms 1, 2, and 3 for each k. Then we say $(f_k)_{k\in\mathbb{Q}}$ fulfills the scaling property, if for all $k_1, k_2 \in \mathbb{Q}$ following holds:

$$f_{k_1} = L\left(\frac{k_1}{k_2}, f_{k_2}\right),$$

where $L(\cdot,\cdot)$ is the scaling operator as defined in equation 3.1.

Fees.

When a trader executes a trade of size Δ_α, the AMM charges a fee on behalf of the liquidity provider from the trader to execute the trade. The fee is paid in β. The fee is specified by a function

$$\gamma(R_\alpha, \Delta_\alpha) \geq 0,$$

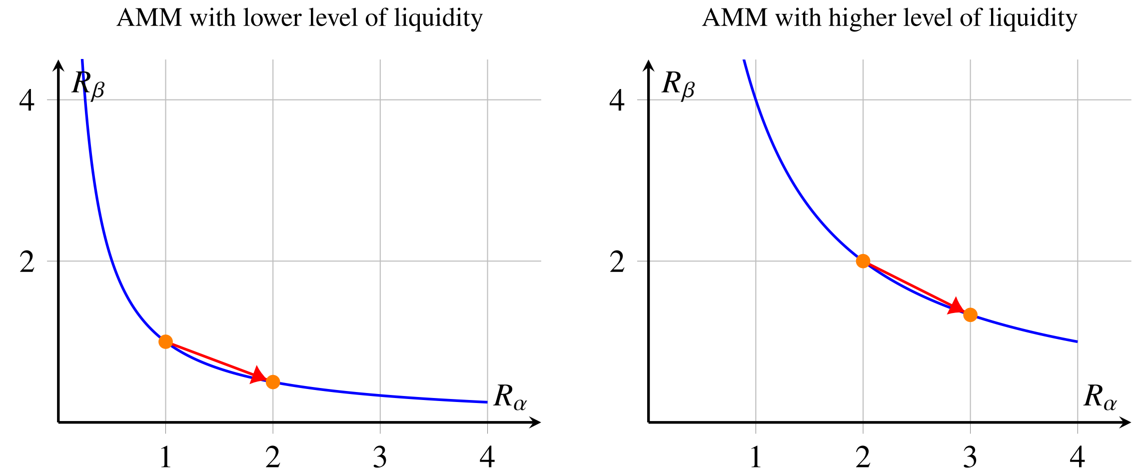

Current spot price of α is 1, therefore current state is $(1\alpha, 1\beta)$, a trader sells one α and receives 0.5β. AMM reserves changed to $(2\alpha, 0.5\beta)$. The average price of α is 0.5β, the new spot price of α is 0.25.

Current spot price of α is 1, therefore current state is $(2\alpha, 2\beta)$, a trader sells one α and receives 0.67β. AMM reserves changed to $(3\alpha, 1.33\beta)$. The average price of α is 0.67β, the new spot price of α is 0.44.

Figure 10: Visualization on how higher liquidity leads to lower changes of the spot price. And the average price of a trade is closer to the spot price before the trade is executed.

which depends on the current reserve level and the traded amount. We adopt a fee model where the fee per contract is constant. Let v_0 be the notional of the offset token, then

$$\gamma(R_\alpha, \Delta_\alpha) = \kappa \cdot v_0 \cdot |\Delta_\alpha|,$$

with $\kappa > 0$ as a proportionality constant. In figure 20 is highlighted by pseudo code, showing how the fee is charged and how it is distributed to the liquidity provider.

3.3 Chosen AMM Design

In Section 2.1, we introduced the invariant

$$f(R_\alpha) = R_\alpha^{-1} + (10 - R_\alpha)^{-1}$$

as a first potential candidate. The constant 10 represents the liquidity level k. For small values of R_α, this invariant resembles the constant-product AMM ($R_\alpha \cdot f(R_\alpha) = 1$), which in general takes the form $R_\alpha \cdot R_\beta = M$ for some constant $M > 0$. Generalizing this idea yields

$$f(R_\alpha) = MR_\alpha^{-1} + N(k - R_\alpha)^{-1},$$

where $M, N > 0$ are constants. Inspired by the BMM (*Better Market Maker*) introduced in [8], we further generalize to

$$f(R_\alpha) = MR_\alpha^{-a} + N(k - R_\alpha)^{-b},$$

for constants $a, b > 0$. However, this form does not satisfy the scaling property 3.5

$$f_{k_1} = L\left(\frac{k_1}{k_2}, f_{k_2}\right).$$

To ensure this property holds, we modify the invariant to

$$f(R_\alpha) = Mk^{a+1}R_\alpha^{-a} + Nk^{b+1}(k - R_\alpha)^{-b}.$$

Definition 3.6 (Power Sum Invariant)**.** The family of *Power Sum Invariant* invariants is given by

$$f_{a,b,M,N,k}(R_\alpha) = Mk^{a+1}R_\alpha^{-a} + Nk^{b+1}(k - R_\alpha)^{-b},$$

with parameters $a, b, M, N, k > 0$.

Lemma 3.7 (Scaling Property)**.** *For any $k_1, k_2 > 0$ it holds that*

$$f_{a,b,M,N,k_1} = L\left(\frac{k_1}{k_2}, f_{a,b,M,N,k_2}\right).$$

Proof. Let $c = \frac{k_1}{k_2}$. By definition of $L(c, f)$ we have

$$L(c, f_{a,b,M,N,k_2})(x) = c\, f_{a,b,M,N,k_2}\left(\frac{x}{c}\right).$$

Substituting the definition of f_{a,b,M,N,k_2} yields

$$\begin{aligned} L(c, f_{a,b,M,N,k_2})(x) &= c\left[Mk_2^{a+1}\left(\frac{x}{c}\right)^{-a} + Nk_2^{b+1}\left(k_2 - \frac{x}{c}\right)^{-b}\right] \\ &= c\left[Mc^a k_2^{a+1}x^{-a} + Nc^b k_2^{b+1}(ck_2 - x)^{-b}\right] \\ &= Mc^{a+1}k_2^{a+1}x^{-a} + Nc^{b+1}k_2^{b+1}(ck_2 - x)^{-b} \\ &= M(ck_2)^{a+1}x^{-a} + N(ck_2)^{b+1}(ck_2 - x)^{-b}. \end{aligned}$$

Since $ck_2 = k_1$, we obtain

$$\begin{aligned} L(c, f_{a,b,M,N,k_2})(x) &= Mk_1^{a+1}x^{-a} + Nk_1^{b+1}(k_1 - x)^{-b} \\ &= f_{a,b,M,N,k_1}(x), \end{aligned}$$

and the scaling property holds. □

Lemma 3.8 (Satisfaction of Axioms 1–3)**.** *For $a, b, M, N, k > 0$, the Power Sum Invariant*

$$f(x) = Mk^{a+1}x^{-a} + Nk^{b+1}(k - x)^{-b}, \qquad x \in (0, k),$$

satisfies Axioms 1 (Regularity), 2 (Strict Convexity), and 3 (Boundary Behavior).

Proof. Axiom 1 Regularity: On $(0, k)$, both functions $x \mapsto x^{-a}$ and $x \mapsto (k - x)^{-b}$ are C^∞. Hence, $f \in C^2((0, k))$.
Axiom 2 Strict Convexity: The first and second derivatives are

$$\begin{aligned} f'(x) &= -aMk^{a+1}x^{-a-1} + bNk^{b+1}(k - x)^{-b-1}, \\ f''(x) &= a(a+1)Mk^{a+1}x^{-a-2} + b(b+1)Nk^{b+1}(k - x)^{-b-2}. \end{aligned}$$

For all $x \in (0, k)$, each term in $f''(x)$ is positive. Thus $f''(x) > 0$, implying that f' is strictly increasing and f is strictly convex.

Axiom 3 Boundary Behavior: As $x \downarrow 0$, the first term diverges to $+\infty$, and as $x \uparrow k$, the second term diverges to $+\infty$:

$$\lim_{x \downarrow 0} f(x) = +\infty, \qquad \lim_{x \uparrow k} f(x) = +\infty.$$

Therefore, the boundary condition is satisfied. □

3.4 Other Concepts Allowing Negative Prices

In [9], [10], and [11], an invariant of a circular form is introduced and discussed. The basic form of the invariant is

$$(R_\alpha - z)^2 + (R_\beta - z)^2 = z^2.$$

This formulation has the advantage that both tokens α and β can have positive or negative prices. In our setting, we have so far assumed that only the asset α can have a negative price. In our invariant, we limited the amount of α held by the AMM through an upper bound k, while the amount of β is unlimited. As long as β has a positive fair value with respect to some reference currency, this poses no problem. However, if β can also have a negative fair value with respect to some currency, then a trader could get rid of negatively priced tokens almost for free.

4 Simulation

One key objective of this thesis is to quantify how *market depth* and *impermanent loss* shape the risk-adjusted return of an AMM that admits negative prices. To evaluate this return, we construct a simulation framework that combines a stochastic model for fair values with an intraday trading environment capturing arbitrage activity, noise-trader flow, fee accumulation, and inventory risk.

The structure of this chapter is as follows. Section 4.1 defines the market environment that generates fair-value time series. Section 4.2 introduces the trading environment. Section 4.3 defines simulation paths together with the associated random variables: fee revenue, cost, and liquidity-provider profit.

4.1 Market Environment

In financial modeling, Geometric Brownian Motion (GBM) is a commonly used process for representing the evolution of asset prices. As outlined in [4], GBM constitutes the underlying assumption of the Black–Scholes option pricing framework, which remains a standard reference point in quantitative finance.

The GBM framework, however, assumes that asset returns follow a normal distribution. Empirical studies indicate that this assumption is often violated in financial markets. For instance, [12] provide evidence that daily stock returns display noticeable deviations from normality, including heavy tails. Such findings illustrate known limitations of GBM and motivate the consideration of alternative models when capturing observed return dynamics.

4.1.1 Notation

We model the fair value of the token α on a discrete set of *trading days*. To formalize this, we first fix the number of trading days in a one-year horizon. Let $N \in \mathbb{N}$ denote the number of trading days per year (we use $N = 252$ throughout the thesis, as that is the average number of trading days per year). The set of days is

$$\mathcal{D} := \{1, 2, \ldots, N\}.$$

We refer to the fair value of α and of its underlying asset U on any given day. For $d \in \mathcal{D}$, we denote these fair values by

$$v_\alpha(d) \quad \text{and} \quad v_U(d),$$

respectively. A complete history of α's fair value over the year is called a *fair-value time series*. Formally, the sequence

$$\big(v_\alpha(d)\big)_{d \in \mathcal{D}}$$

is a fair-value time series of α, and we write $\mathbb{V}_\alpha$ for the set of all such time series.

4.1.2 Offset Geometric Brownian Motion

Because the price formation mechanism is not our main object of study, we use *Geometric Brownian Motion* (GBM) as a standard model from mathematical finance to produce realistic, strictly positive dynamics for the underlying U. Since α may be negative by construction, we obtain α's fair value via a deterministic *offset* of the underlying.

We now define the discrete-time GBM for the underlying and its offset mapping to v_α. We first specify daily log returns and the compounding rule for U.

Let $\mu \in \mathbb{R}$ be the annual drift, $\sigma > 0$ the annual volatility, N the number of trading days per year, and $v_0 > 0$ a reference level. We model daily log returns by

$$r_i \;=\; \left(\mu - \tfrac{1}{2}\sigma^2\right)\tfrac{1}{N} \;+\; \sigma\sqrt{\tfrac{1}{N}}\,Z_i, \quad \text{with } Z_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1).$$

Given an initial level of the underlying $v_U(0)$, the fair value of U after t trading days is obtained by compounding these log returns

$$v_U(t) \;=\; v_U(0)\,\exp\Big(\sum_{i=1}^{t} r_i\Big).$$

We obtain the fair value of α by subtracting the reference level $v_U(0)$, thereby allowing negative values for α

$$v_\alpha(t) \;=\; v_U(t) - v_U(0).$$

For notational convenience, we encapsulate this construction in a short definition.

Definition 4.1 (Offset GBM for α)**.** Let U_t follow a discrete GBM with parameters (μ, σ, N) and initial level $U_0 = v_0 > 0$. Define α's fair value by $v_\alpha(t) := U_t - v_0$ for $t = 0, \ldots, N$. We write

$$v_\alpha \sim \text{OffsetGBM}(\mu, \sigma, N, v_0).$$

Equivalently, if r_i are the GBM log returns, then $U_t = v_0 \exp(\sum_{i=1}^{t} r_i)$ and

$$v_\alpha(t) \;=\; v_0\left(\exp\Big(\sum_{i=1}^{t} r_i\Big) - 1\right).$$

Example (Path sampling). Drawing M independent fair-value paths,

$$v_\alpha^{(m)} \sim \text{OffsetGBM}(0.10,\ 0.10,\ 252,\ 100), \qquad m = 1, \ldots, M,$$

corresponds to simulating M paths with 10% annual drift, 10% annual volatility, 252 trading days, and $v_0 = 100$. Each path starts at $v_\alpha(0) = 0$ and may take negative or positive values thereafter.

4.2 Trading Environment

This subsection describes the intraday trading environment used in our AMM analysis. The structure is motivated by [13], which studies constant-product (Uniswap-like) markets, explains how arbitrage keeps AMM prices close to an external reference price, and uses an agent-based simulation to test stability under different market conditions.

We adapt these ideas to our setting. First, we model each day as an ordered sequence of *timepoints* so that trade execution and state updates are easy to track (opening arbitrage, a fixed number of noise-trader intents, and closing arbitrage). Second, we include an execution filter that skips noise trades when their average execution price differs too much from the fair value, preventing trades that would be clearly uneconomic.

4.2.1 Timepoints and AMM State

Trading activity occurs within each day as a sequence of *trade intents*. To index these events and the AMM state after each event, we now introduce *timepoints*.

Fix a maximum number of within-day trade intents $T_{\max} \in \mathbb{N}$ and define the within-day index set

$$S := \{1, 2, \ldots, T_{\max}\}.$$

A *timepoint* is an ordered pair consisting of a day and an intent index,

$$t = (d, s) \in \mathcal{T} := \mathcal{D} \times S.$$

At any timepoint $t = (d, s)$, a proposed trade size in α-units is denoted by

$$\Delta_\alpha(d, s).$$

An intent either executes or is skipped. We therefore introduce an indicator

$$\delta(d, s) \in \{0, 1\},$$

where $\delta(d, s) = 1$ signifies execution of the intent and $\delta(d, s) = 0$ otherwise. The AMM reserves after processing the s-th intent on day d are written as

$$\big(R_\alpha(d, s),\ R_\beta(d, s)\big).$$

If $\delta(d, s) = 0$, then the state does not change relative to the previous timepoint.

4.2.2 Arbitrage Trader

The arbitrage trader exploits discrepancies between the AMM's spot price and the fair value $v_\alpha(d)$ observed externally. We assume the arbitrageur is active twice per day, at the first and last intents. The corresponding timepoints are

$$\mathcal{T}_{\text{arbitrage}} := \mathcal{D} \times \{1,\ T_{\max}\} \subseteq \mathcal{T}.$$

At an arbitrage timepoint $t = (d, s) \in \mathcal{T}_{\text{arbitrage}}$, the trader chooses a trade size to maximize profit, measured in β-units. We first relate a candidate Δ_α to the implied β-flow via the invariant curve.

Let f denote the AMM's trading function (invariant). If the state before the trade is $(R_\alpha(t), R_\beta(t))$ and the trader submits a size Δ_α, then the corresponding β-amount exchanged is

$$\Delta_\beta \;=\; f\big(R_\alpha(t) + \Delta_\alpha\big) \;-\; R_\beta(t),$$

and the trade must keep reserves in the feasible domain, which we write generically as

$$R_\alpha(t) + \Delta_\alpha \in (0, k).$$

We now express the arbitrage profit for a candidate size. The trader values α at $v_\alpha(d)$ and pays the AMM's path price (captured by Δ_β) plus a fee $\gamma(R_\alpha(t), \Delta_\alpha)$. The profit function is therefore

$$\Pi(\Delta_\alpha; t) \;=\; \Delta_\alpha\, v_\alpha(d) \;-\; \big(f(R_\alpha(t) + \Delta_\alpha) - R_\beta(t)\big) \;-\; \gamma(R_\alpha(t), \Delta_\alpha),$$

subject to $R_\alpha(t) + \Delta_\alpha \in (0, k)$. The arbitrageur's optimal size at t solves

$$\Delta_\alpha(t) \;\in\; \underset{\Delta_\alpha :\, R_\alpha(t)+\Delta_\alpha \in (0,k)}{\arg\max} \;\Pi(\Delta_\alpha; t).$$

Arbitrage intents are always succesfull, therefore $\delta(t) = 1$ for all $t \in \mathcal{T}_{\text{arb}}$, however $\Delta_\alpha(t)$ might be 0. The intuitive effect of an arbitrage trade is that whenever the AMM's spot quote deviates sufficiently from $v_\alpha(d)$, the optimal arbitrage trade moves the state toward the fair value.

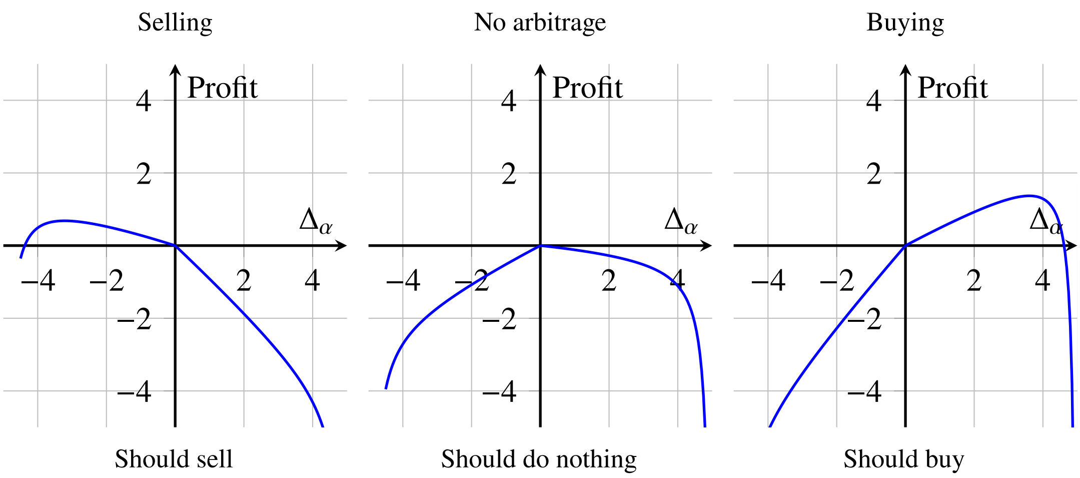

Figure 11: Visualization of arbitrage optimization problem.

4.2.3 Noise Traders

Noise traders generate baseline demand for liquidity. We assume a fixed number $\lambda \in \mathbb{N}$ of noise-trader intents per day. Their timepoints are

$$\mathcal{T}_{\text{noise}} \;=\; \mathcal{D} \times \{2, 3, \ldots, \lambda + 1\},$$

so that each day consists of an opening arbitrage intent ($s = 1$), followed by λ noise intents, and a closing arbitrage intent ($s = \lambda + 2$). Consequently,

$$T_{\max} = \lambda + 2.$$

Each noise intent draws a direction (buy or sell) with equal probability and a random trade size. We model the absolute trade size with a lognormal distribution. Let m and $s > 0$ parameterize the underlying normal distribution. We set

$$|\Delta_\alpha(t)| \sim \text{Lognormal}\big(m - \tfrac{1}{2}s^2,\ s^2\big), \qquad \mathbb{P}\big(\Delta_\alpha(t) > 0\big) = \mathbb{P}\big(\Delta_\alpha(t) < 0\big) = \tfrac{1}{2}.$$

Large trades face slippage. To prevent executing uneconomic executions, we impose an *execution filter*: an intent only executes if its average execution price (including fees) stays within an ε-band around fair value. Let $p_{\text{avg}}(R_\alpha(t), \Delta_\alpha(t))$ denote the AMM's average price for size $\Delta_\alpha(t)$. The execution rule is

$$\delta(t) = \begin{cases} 1, & \text{if } v_\alpha(d) - \varepsilon \ \le\ p_{\text{avg}}\big(R_\alpha(t), \Delta_\alpha(t)\big) \ \le\ v_\alpha(d) + \varepsilon, \\ 0, & \text{otherwise.} \end{cases}$$

Hence, noise-trader behavior is governed by the parameters $(m, s, \lambda, \varepsilon)$.

4.3 Simulation Paths and Profit Components

4.3.1 Simulation Paths

A simulation path consists of one fair-value time series of α together with all noise-trader intents over the horizon. Formally,

$$\omega = \Big((v_\alpha(i))_{i=1}^{N},\ (\Delta_\alpha(i,j))_{i=1,\dots,N}^{j=2,\dots,\lambda+1} \Big),$$

and we use the projection

$$V(\omega) := (v_\alpha(i))_{i=1}^{N}$$

to refer to the fair-value component alone.

4.3.2 Fee Revenue, Cost, and Risk-Adjusted Profit

Liquidity providers earn fees on executed trades. Our fee rule as described in 3.2 is that the fee per trade is proportional to the amount of α bought or sold

$$\gamma(R_\alpha(t), \Delta_\alpha(t)) = v_0 \kappa\, |\Delta_\alpha(t)|, \quad \kappa \ge 0,$$

and v_0 is the notional value of the offset token. The total fee revenue on a path ω is therefore

$$X^{\text{revenue}}(\omega) = \sum_{t \in \mathcal{T}} \gamma\big(R_\alpha(t), \Delta_\alpha(t)\big)\, \delta(t) = \sum_{t \in \mathcal{T}} \kappa\, |\Delta_\alpha(t)|\, \delta(t).$$

The only source of cost is the opportunity cost of providing liquidity instead of holding the initial liquidity. Let the initial reserves be

$$(R_\alpha^{(0)}, R_\beta^{(0)}),$$

and the terminal reserves after the final intent on day N be

$$(R_\alpha^{\text{fin}}, R_\beta^{\text{fin}}).$$

To estimate the cost on path ω, we calculate, what the fairvalue of the initional investment $(R_\alpha^{(0)}, R_\beta^{(0)})$ would be now, if it has not been used to provide liquidity and compare it with the fair value of the terminal reserves. The fair value of $(R_\alpha^{(0)}, R_\beta^{(0)})$ would be

$$\left(R_\alpha^{(0)} \, v_\alpha(N) + R_\beta^{(0)}\right)$$

and the fair value of the terminal reserves are

$$\left(R_\alpha^{\text{fin}} \, v_\alpha(N) + R_\beta^{\text{fin}}\right),$$

therefore the liquidity provider has cost of

$$X^{\text{cost}}(\omega) = \left(R_\alpha^{(0)} \, v_\alpha(N) + R_\beta^{(0)}\right) - \left(R_\alpha^{\text{fin}} \, v_\alpha(N) + R_\beta^{\text{fin}}\right).$$

The liquidity provider's profit on a path ω is defined as

$$X^{\text{profit}}(\omega) = X^{\text{revenue}}(\omega) - X^{\text{cost}}(\omega).$$

The risk-adjusted return is estimated as the ratio of expected profit to the standard deviation of profit,

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})}.$$

A closely related measure is the Sharpe ratio as introduced in Sharpe [14],

$$\frac{\mathbb{E}\left(\frac{X^{\text{profit}}}{I}\right) - R_f}{\sigma\left(\frac{X^{\text{profit}}}{I}\right)},$$

where R_f denotes the risk-free rate, typically derived from short-term government bond yields. The Sharpe ratio is widely used because it evaluates returns relative to the risk taken and allows comparisons across investment strategies.

In our setting, providing liquidity to an AMM requires no initial investment, therefore $I = 0$. Taking the limit $I \to 0$ yields

$$\lim_{I\to 0} \frac{\mathbb{E}\left(\frac{X^{\text{profit}}}{I}\right) - R_f}{\sigma\left(\frac{X^{\text{profit}}}{I}\right)} = \lim_{I\to 0} \frac{\frac{1}{I}\left(\mathbb{E}(X^{\text{profit}}) + I\,R_f\right)}{\frac{1}{I}\sigma(X^{\text{profit}})} = \frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})}.$$

5 Key Performance Indices

In the previous Chapter 4, we introduced the simulation framework used to compute the risk-adjusted returns of an AMM. While this approach provides detailed insights into the realized performance of liquidity provision, it is also computationally intensive and depends on numerous assumptions about trading behavior. This motivates the search for *a model-independent key performance indicators* (KPIs) that can serve as reliable predictors of AMM performance.

In particular, we investigate whether *impermanent loss* and *market depth* can act as meaningful explanatory variables for risk-adjusted returns. To be suitable for this purpose, such KPIs should be independent of specific trading activity and depend only on the fair value of the underlying asset α or its expected fair-value time series. Our approach is conceptually related to the framework introduced by Engel and Herlihy [15], who proposed the notion of *load* as a combined measure of *slippage* (the price impact of executing trades) and impermanent loss, and extended it to the *expected load* by incorporating a probability distribution over future fair values of α. Building on this idea, we generalize the concept to a set of KPIs derived from the distribution of possible future price paths and explore whether these theoretical measures can approximate or predict simulation-based performance outcomes.

In the following, we formally define the two main KPIs considered in this study:

1. **Impermanent Loss**, which quantifies the opportunity cost of liquidity provision relative to a simple buy-and-hold strategy, given a distribution of possible future prices; and

2. **Market Depth**, which measures an AMM's capacity to absorb trading volume around the fair value without causing significant price impact.

5.1 Notation

We introduced R_α and R_β as variables that denote the current state of an AMM. Given this state, the spot price can be expressed as $p(R_\alpha)$. We now invert this relationship: let v_α denote the fair value of α. Then

$$(R_\alpha(v_\alpha), R_\beta(v_\alpha))$$

represents the state such that the spot price at this state satisfies $p(R_\alpha(v_\alpha)) = v_\alpha$.

As defined in 4.1.1, we denote by

$$\mathbb{V}_\alpha = \{(v_\alpha(d))_{d\in\mathcal{D}} : v_\alpha(d) \in \mathbb{R}\}$$

the set of all random walks of α. Let $v = (v_\alpha(d))_{d\in\mathcal{D}} \in \mathbb{V}_\alpha$, then

$$v_{[-1]} := v_\alpha(\max(\mathcal{D}))$$

denotes the fair value of α on the last day of that random price walk. For a set of random price walks $V \subseteq \mathbb{V}_\alpha$, we denote

$$V_{[-1]} := \{v_{[-1]} : v \in V\}.$$

Finally, let $X \in \mathbb{V}_\alpha$ be a random variable taking values in $\mathbb{V}_\alpha$.

5.2 Impermanent Loss

Impermanent loss measures the opportunity cost of providing liquidity to an AMM. Loesch et al. [16] find that "Impermanent Loss (or 'IL') is the dominant factor in determining the financial impact of liquidity provision on Uniswap v3."

Hence, we use impermanent loss as a KPI responsible for predicting potential losses. We define it as follows: assume we know that the fair value of the token α is p at the end of the AMM's lifetime. The impermanent loss is the loss incurred by providing liquidity instead of simply holding the tokens $(R_\alpha^{(0)}, R_\beta^{(0)})$. At that point, the liquidity provider holds (R_α, R_β) such that the spot price equals p.

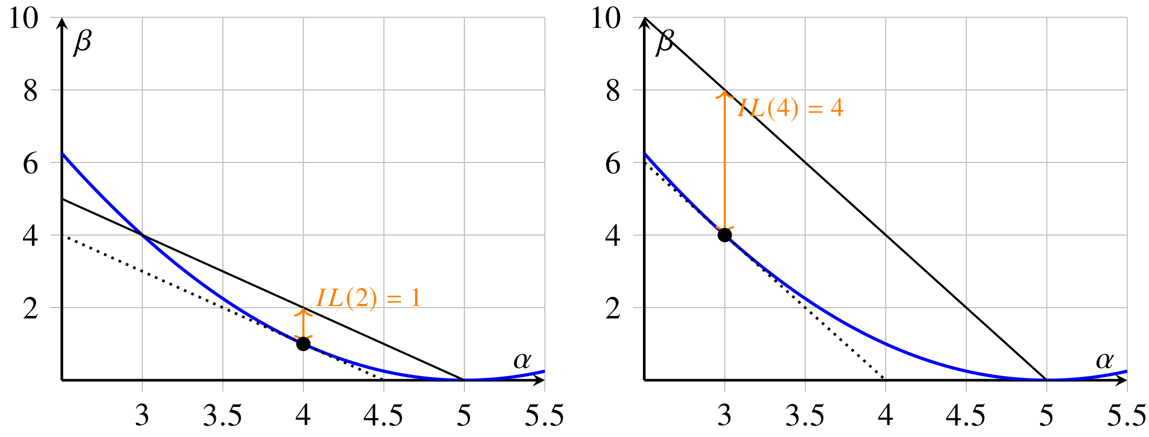

For $p = 2$, the AMM deviates slightly from the ideal path, causing a small impermanent loss of 1β.

For $p = 4$, the deviation is stronger, resulting in an impermanent loss of 4β.

Figure 12: Illustration of impermanent loss. The black curve shows states where $IL = 0$. The dotted line is parallel to it and indicates the final AMM state at the fair price. The vertical arrow shows the impermanent loss in β needed to restore the zero-IL condition.

The initial investment would be worth

$$V_{\text{hold}}(p) = R_\alpha^{(0)} \cdot p + R_\beta^{(0)},$$

whereas the AMM reserves are worth

$$V_{\text{LP}}(p) = R_\alpha(p) \cdot p + R_\beta(p).$$

The difference between these two values defines the impermanent loss

$$IL(p) = V_{\text{hold}}(p) - V_{\text{LP}}(p).$$

However, we are often interested in the impermanent loss over a random price path of α. For a fair-value path v of α, we denote its terminal fair value by $v_{[-1]}$. So, if we know the fair-value path v of α, we denote it as

$$IL(v) = IL(v_{[-1]}).$$

In most cases, we do not know the exact fair-value path of α. If we assume that it follows a probability distribution described by the random variable V, we can define the theoretical impermanent loss as

$$X^{\mathrm{IL}} = IL(V).$$

5.2.1 Comparison Between Cost and Impermanent Loss

When comparing the definitions of cost and impermanent loss, we see that they differ only in the evaluation of the fair value of the α and β tokens when retracting liquidity. In the impermanent loss case, we assume that the amounts of α and β are such that the spot price equals the fair value of α. For the simulated impermanent loss, we do not know if his is the case. However, we know that

$$R_\alpha \in \{R_\alpha : p(R_\alpha) \in [p - \kappa,\ p + \kappa]\}.$$

Therefore, simulated impermanent loss and theoretical impermanent loss are close to each other.

5.3 Market Depth

Market depth captures the ability of an Automated Market Maker (AMM) to absorb trading volume without causing large price movements. Empirical evidence suggests that higher market depth is essential for supporting greater trading activity: Liao and Robinson [17] find that "high market depth is needed to support high transaction volumes with reliable execution." Hence, market depth can be viewed as a key determinant of execution quality and trading capacity in AMMs, as high market depth allows large trades to be executed without having large changes in the price, as market depth describes how much of the token α can be bought or sold before the average price exceeds a certain threshold.

At a specific fair value p of α, assume the reserves of the AMM are set such that the spot price equals p, and let τ denote the price sensitivity. Then the market depth is the difference between the largest and smallest order such that the average price of that trade is still greater than or equal to $p - \tau$ and less than or equal to $p + \tau$. Therefore, we define $MD(p)$ as

$$MD_\tau(p) = \sup\{d : \bar{p}(R_\alpha(p), d) \le p + \tau\} - \inf\{d : \bar{p}(R_\alpha(p), d) \ge p - \tau\}.$$

Now assume we know the entire fair-value path of α, denoted by

$$v = (v_\alpha(d))_{d \in \mathcal{D}}.$$

In this case, it is more informative to measure the mean market depth over the duration of the path, defined as

$$MD_\tau(v) = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} MD_\tau(v_\alpha(d)).$$

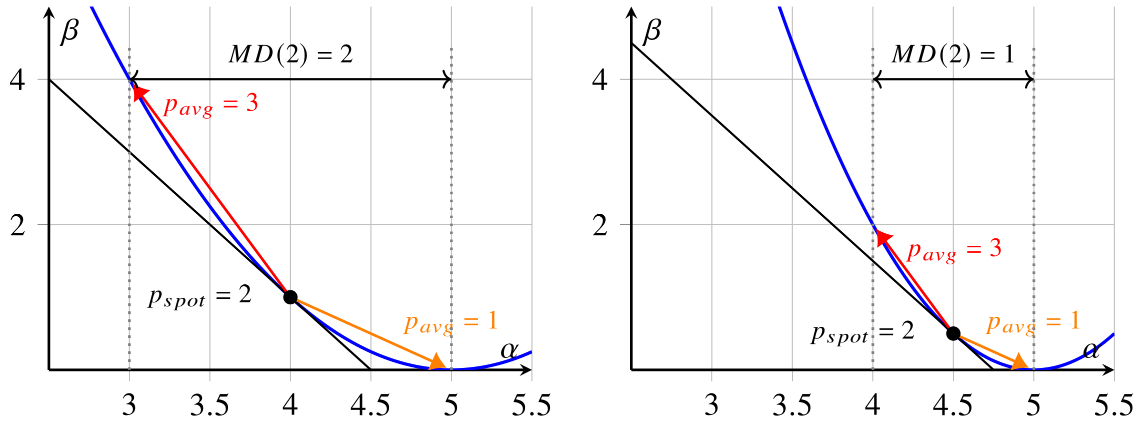

Invariant $f(x) = (x-5)^2$ and a market depth of 2 when fair value of α is 2.

Invariant $f(x) = 2(x-5)^2$ and a market depth of 1 when fair value of α is 2.

Figure 13: Visualization of the market depth: Current fair value is $p = 2$, price sensitivity is $\tau = 1$, black dot shows the state of the AMM, where the spot price is 1. The red (orange) arrows are highlighting the largest buy (sell) orders, such that the average price of those trades are still within $[p - \tau, p + \tau]$

In general, we do not know the exact fair-value path but can assume that it follows some random distribution V, which takes values in the set of all fair-value time series. Then we define the random variable

$$X^{\text{depth}(\tau)} = MD_\tau(V).$$

5.3.1 Common Definition of Market Depth and τ Comparison

In the literature (for example, in Kempf and Korn [18]), market depth is commonly defined as the required trade size to achieve a certain price change, which in our case would translate to the spot price. It is defined as

$$\text{Market Depth} = \left(\frac{dP}{dQ}\right)^{-1},$$

where dP is the change of price and dQ is the amount of the asset that got traded. Using our notation, assuming the current spot price is p_0 and it changes by dP, we can compute dQ as

$$dQ = R_\alpha(p_0 + dP) - R_\alpha(p_0),$$

and therefore the market depth would be defined as

$$\text{Market Depth} = \frac{R_\alpha(p_0 + dP) - R_\alpha(p_0)}{dP}.$$

The common definition and the definition we introduced look quite similar. The differences are: first, the common definition is with respect to the spot price and not the average price of a trade, and second, the common definition considers the limit as

$dP \to 0$. We have introduced a measure of market depth depending on the average price of a trade because the buy/sell decision of the noise traders is based on the average price of the trade. Based on the displays in Figure 14 we can see that when

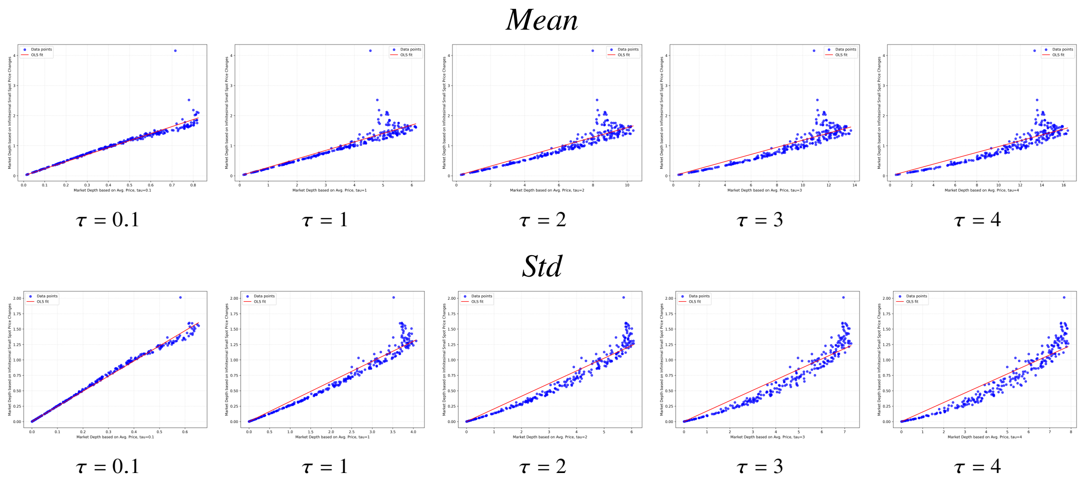

Figure 14: Comparison between market using common definition and using average price based definition of market depth for $\tau \in \{0.1, 1, 2, 3, 4\}$, by plotting the (i) expected value and (ii) standard deviation of the common definition on the vertical axis and $X^{\text{depth}(\tau)}$ on the horizontal axis.

choosing a small τ (e.g., $\tau = 0.1$), there is an almost linear relationship. However, for larger τ it becomes curved and less precise.

5.3.2 Market Depth and Revenue

In our simulation, an AMM generates more revenue when a larger number of trades execute successfully, and when larger trades are executed successfully more frequently. A trade is not executed if its average price deviates too much from the fair value. This is precisely where market depth becomes crucial: it measures the range of trade sizes for which the average execution price stays within a deviation of at most τ from the fair value. If this range is large, more potential trades pass the execution filter, and larger trades are also more likely to be executed.

6 Numerical Results

We investigate whether market depth and impermanent loss (IL) are sufficient to predict the risk-adjusted profitability of an AMM. Concretely, we aim to identify a function g such that

$$\frac{\mathbb{E}\left(X^{\text{profit}}\right)}{\sigma\left(X^{\text{profit}}\right)} \approx g\left(\mathbb{E}\left(X^{\text{IL}}\right), \sigma\left(X^{\text{IL}}\right), \mathbb{E}\left(X^{\text{depth}}\right), \sigma\left(X^{\text{depth}}\right)\right). \tag{6.1}$$

We proceed in several steps to understand the functional form of g.

6.1 Simulation Setup

Throughout, we consider AMMs with invariants of the form

$$f(x) = M\, k^{a+1} x^{-a} + N\, k^{b+1} (k-x)^{-b}. \tag{6.2}$$

For each experiment, we specify how many invariants are drawn, the parameter ranges for a, b, M, N, k, and the simulation-path parameters. Let $\Omega = \{\omega_1, \ldots, \omega_l\}$ denote the set of simulation paths described in Section 4.3. The market environment is parameterized by μ, σ, and the noise-trader behavior by m, s, λ, ϵ (see Chapter 4). We choose $l = 100$; therefore, we draw 100 simulation paths to test each invariant.

For every AMM A, we simulate fee revenue, cost, and profit along full price paths with trader intents, and compute impermanent loss and market depth from price paths only (i.e., using the projection $\{V(\omega_1), \ldots, V(\omega_l)\}$). We use the shorthand

$$\begin{aligned} E^{\text{profit}}(A) &= \mathbb{E}\left(X_A^{\text{profit}}\right), & \Sigma^{\text{profit}}(A) &= \sigma\left(X_A^{\text{profit}}\right), \\ E^{\text{revenue}}(A) &= \mathbb{E}\left(X_A^{\text{revenue}}\right), & \Sigma^{\text{revenue}}(A) &= \sigma\left(X_A^{\text{revenue}}\right), \\ E^{\text{cost}}(A) &= \mathbb{E}\left(X_A^{\text{cost}}\right), & \Sigma^{\text{cost}}(A) &= \sigma\left(X_A^{\text{cost}}\right), \\ E^{\text{IL}}(A) &= \mathbb{E}\left(X_A^{\text{IL}}\right), & \Sigma^{\text{IL}}(A) &= \sigma\left(X_A^{\text{IL}}\right), \\ E^{\text{depth}}(A) &= \mathbb{E}\left(X_A^{\text{depth}}\right), & \Sigma^{\text{depth}}(A) &= \sigma\left(X_A^{\text{depth}}\right). \end{aligned}$$

We set up two experiments. Across both experiments, we draw AMM invariants of the form (6.2) using identical parameter ranges. The parameters a, b are drawn randomly from the interval $[10^{-0.7}, 10^{0.7}]$ and the parameters M, N are drawn randomly from the interval $[10^{-2}, 10^{0}]$. The parameter k is fixed at $k = 30$. The market environment is held constant across experiments, with drift $\mu = 0.1$ and volatility $\sigma = 0.1$.

Noise-trader behavior is also largely identical between the two experiments. In both cases, we set the average trade size parameter to $m = 10$ and the average number of trade intents to $\lambda = 8$. The experiments differ only in the dispersion (s) and tolerance parameters (ϵ) of the noise traders. In Experiment 1, we use $s = 2.5$ and $\epsilon = 0.7$, whereas in Experiment 2 these parameters are changed to $s = 4$ and $\epsilon = 1.5$.

6.2 Simulation Results

As a reminder, the goal is to find a function g as described in 6.1. As a first step, we can reformulate the risk-adjusted profit as

$$\frac{\mathbb{E}\left(X^{\text{profit}}\right)}{\sigma\left(X^{\text{profit}}\right)} = \frac{\mathbb{E}(X^{\text{revenue}}) - \mathbb{E}(X^{\text{cost}})}{\sigma(X^{\text{revenue}} - X^{\text{cost}})}.$$

However, we want to rewrite this expression in terms of X^{IL} and $X^{\text{depth}(\tau)}$. As a first step, we determine the appropriate choice of τ.

6.2.1 Step 1: Finding the Optimal Market Depth Parameter τ

Description: The market depth $X^{\text{depth}(\tau)}$ depends on the parameter τ. The goal of this step is to analyze whether the choice of τ is relevant and, if so, whether there is an optimal τ. If an optimal τ exists, we also ask whether it is consistent across experiments or depends on the experimental setup.

Data and Visualization: For data, we use E^{revenue} and $E^{\text{depth}(\tau)}$, as well as Σ^{revenue} and $\Sigma^{\text{depth}(\tau)}$ for $\tau \in \{0.1, 1, 2, 3, 4\}$. The corresponding scatter plots and relationships are displayed in Figure 15.

As shown in Figure 15, a different τ is optimal in experiment 1 than in experiment 2. We also see that if we choose τ too small, there is a rightward curvature, whereas if we choose it too large, there is a leftward curvature.

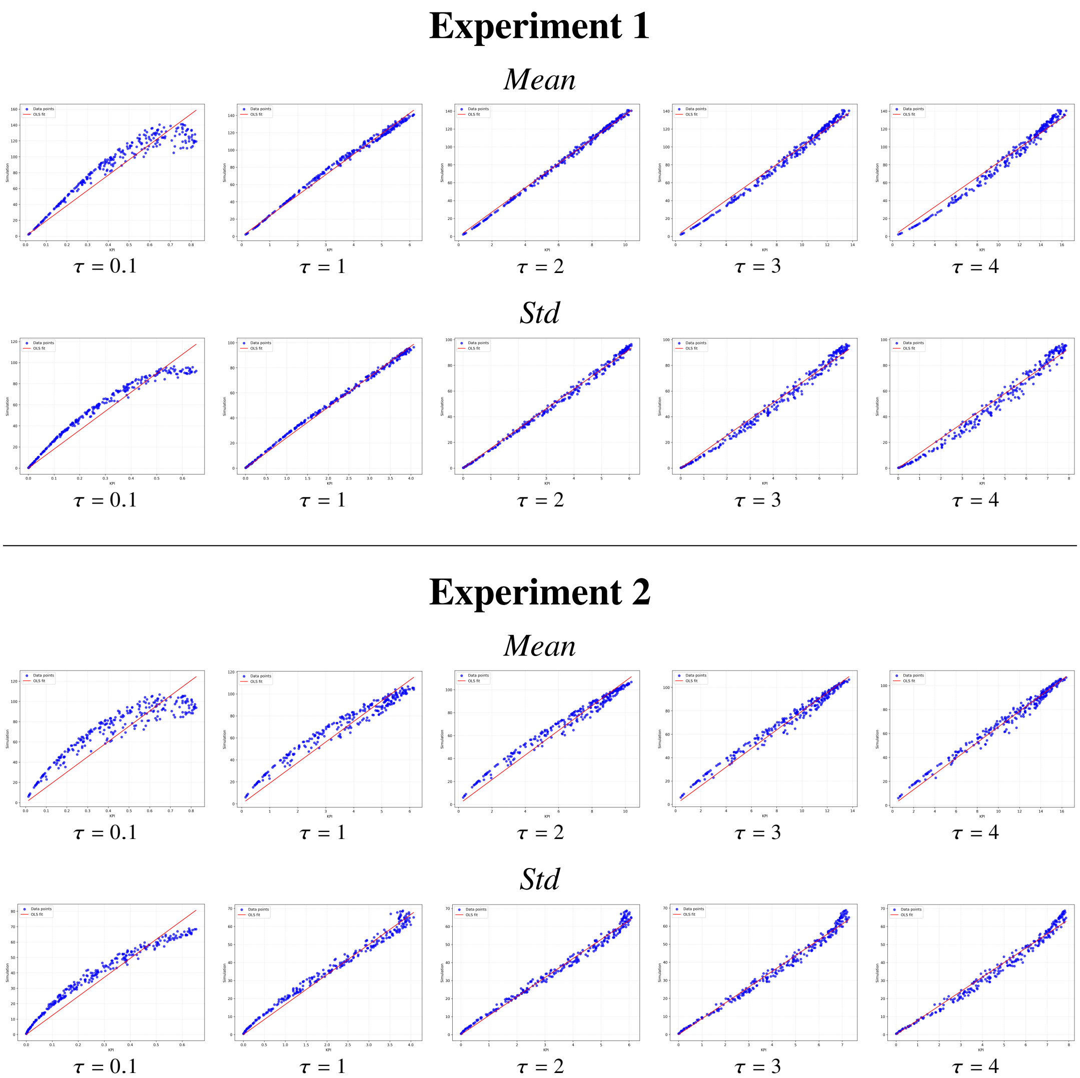

Figure 15: For each experiment different choices of τ to compare mean market depth and mean revenue, as well as standard deviation of market depth and standard deviation of revenue.

Statistical Analysis: The goal here is to determine the optimal τ for each setting. As seen above in Figure 15, we can choose τ too large or too small. Therefore, we compute the correlation between market depth and revenue for different choices of τ:

$$\rho_E(\tau) = \rho_{E^{\text{revenue}}, E^{\text{depth}(\tau)}} \quad \text{and} \quad \rho_\Sigma(\tau) = \rho_{\Sigma^{\text{revenue}}, \Sigma^{\text{depth}(\tau)}}$$

and maximize $\rho_E(\tau) + \rho_\Sigma(\tau)$.

Based on the results from Table 2, we choose $\tau = 2$ for experiment 1, and based on Table 3, we choose $\tau = 3$ for experiment 2.

τ	0.1	1	2	3	4
Mean	0.943	0.997	0.999	0.997	0.994
Std	0.972	0.999	0.998	0.995	0.991
Sum	1.915	1.996	1.997	1.992	1.985

Table 2: Correlation for Experiment 1 when choosing different values for τ.

τ	0.1	1	2	3	4
Mean	0.917	0.980	0.991	0.993	0.992
Std	0.978	0.993	0.994	0.994	0.992
Sum	1.895	1.974	1.985	1.986	1.984

Table 3: Correlation for Experiment 2 when choosing different values for τ.

6.2.2 Step 2: Approximating the Standard Deviation of the Profit

Description: Returning to our goal of approximating the risk-adjusted return using market depth and impermanent loss, so far we have

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})} = \frac{\mathbb{E}(X^{\text{revenue}}) - \mathbb{E}(X^{\text{cost}})}{\sigma(X^{\text{revenue}} - X^{\text{cost}})}.$$

In this step, we will see that there exist a constant m_σ, such that

$$\sigma(X^{\text{revenue}} - X^{\text{cost}}) \approx m_\sigma \cdot \big(\sigma(X^{\text{revenue}}) + \sigma(X^{\text{cost}})\big),$$

and we can approximate the risk-adjusted return by

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})} \approx \frac{\mathbb{E}(X^{\text{revenue}}) - \mathbb{E}(X^{\text{cost}})}{m_\sigma\big(\sigma(X^{\text{revenue}}) + \sigma(X^{\text{cost}})\big)}.$$

Therefore, the goal is to approximate Σ^{profit} using Σ^{revenue} and Σ^{cost}

$$\Sigma^{\text{profit}} = m_\sigma\big(\Sigma^{\text{revenue}} + \Sigma^{\text{cost}}\big).$$

In other words, we try to approximate Σ^{profit} using linear regression. To do that, we set

$$Z := \Sigma^{\text{revenue}} + \Sigma^{\text{cost}}.$$

For the linear regression, we also set the y-intercept to 0, because if the standard deviation of cost and of revenue is 0, then the standard deviation of profit is also 0. In general, the relation

$$\sigma(X+Y) = M\big(\sigma(X) + \sigma(Y)\big)$$

for some $M \in \mathbb{R}$ and random variables X, Y does not hold; however, in our specific case we will see that it provides a good approximation.

Data and Visualization: For data, we use Σ^{profit}, Σ^{cost}, and Σ^{revenue} as described in 6.1. In Figure 16, we illustrate the relationship between Σ^{profit} and Z.

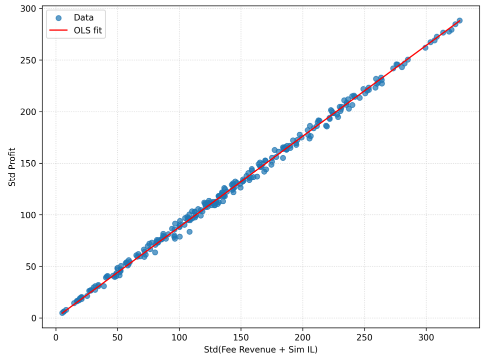

(a) Experiment 1

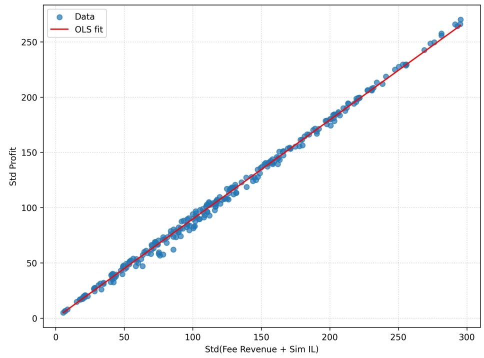

(b) Experiment 2

Figure 16: Relationship between $\sigma(X^{\text{profit}})$ and $\sigma(X^{\text{revenue}}) + \sigma(X^{\text{cost}})$.

		Correlation		Regression	
Experiment	τ	Corr.	p-value	Slope	95% Confidence Interval (CI)
Experiment 1	2	0.999	< 0.001	0.880	[0.878, 0.883]
Experiment 2	3	0.999	< 0.001	0.897	[0.894, 0.900]

Table 4: Correlation and regression results for Std(Fee Revenue) + Std(IL) vs Std(Profit).

Statistical Analysis: In Table 4 we see that in both experiments Z and Σ^{profit} are highly correlated and the p-value is smaller than 0.001, which makes the relationship statistically significant. When we examine it further using linear regression, we obtain a rather tight confidence interval for m_σ.

6.2.3 Step 3: Linear Relationship between Revenue and Market Depth

Description: Returning to our goal of approximating the risk-adjusted return using KPIs, so far we have

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})} \approx \frac{\mathbb{E}(X^{\text{revenue}}) - \mathbb{E}(X^{\text{cost}})}{m_\sigma\left(\sigma(X^{\text{revenue}}) + \sigma(X^{\text{cost}})\right)}.$$

In this step, we will see that

$$\mathbb{E}(X^{\text{revenue}}) \approx m_E\, \mathbb{E}(X^{\text{depth}(\tau)}), \quad \sigma(X^{\text{revenue}}) \approx m_\Sigma\, \sigma(X^{\text{depth}(\tau)})$$

and we can approximate the risk-adjusted return by

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})} \approx \frac{m_E \mathbb{E}(X^{\text{depth}(\tau)}) - \mathbb{E}(X^{\text{cost}})}{m_\sigma\left(m_\Sigma \sigma(X^{\text{depth}(\tau)}) + \sigma(X^{\text{cost}})\right)}.$$

In Step 1, we already saw that, when we choose τ correctly, there is a linear correlation between E^{revenue} and $E^{\text{depth}(\tau)}$, as well as between Σ^{revenue} and $\Sigma^{\text{depth}(\tau)}$.

This step concerns the statistical significance of that correlation. In Chapter 5.3.2, we also discussed a causal relationship between market depth and revenue; therefore, we quantify it by linear regression. We use the model

$$E^{\text{revenue}} = m_E \, E^{\text{depth}}, \qquad \Sigma^{\text{revenue}} = m_\Sigma \, \Sigma^{\text{depth}}.$$

In both cases, the regression lines are constrained to pass through the origin. For the first model, $E^{\text{revenue}} = m_E \, E^{\text{depth}}$, this reflects the fact that when market depth is zero, no trades occur, and hence the expected revenue must also be zero. Similarly, for the second model, $\Sigma^{\text{revenue}} = m_\Sigma \, \Sigma^{\text{depth}}$, the data indicate that Σ^{revenue} approaches zero as Σ^{depth} tends to zero. Therefore, in both regressions, we fix the intercept to zero.

Data and Visualization: For data, we use E^{revenue} and $E^{\text{depth}(\tau)}$, as well as Σ^{revenue} and Σ^{depth} as described in 6.1. In Figure 13, we see an almost linear relationship between E^{revenue} and E^{depth}, as well as between Σ^{revenue} and Σ^{depth}.

Experiment 1

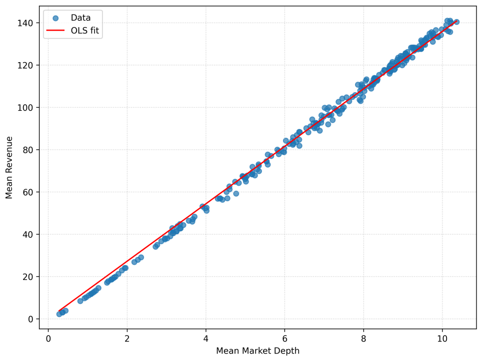

(a) Mean market depth vs mean revenue.

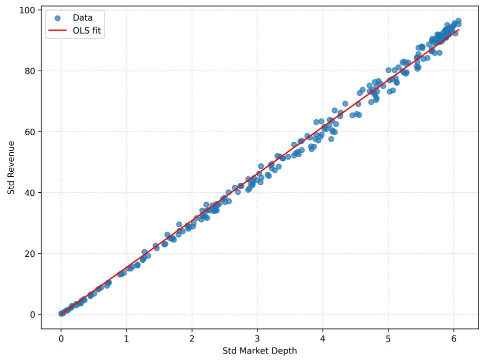

(b) Std market depth vs std revenue.

Experiment 2

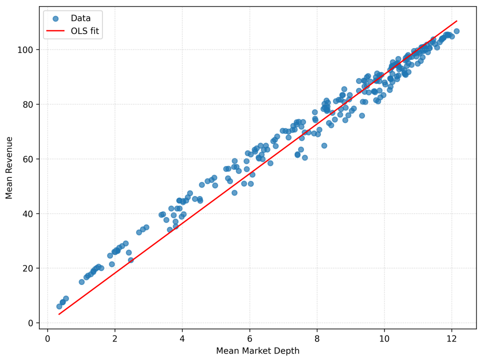

(c) Mean market depth vs mean revenue.

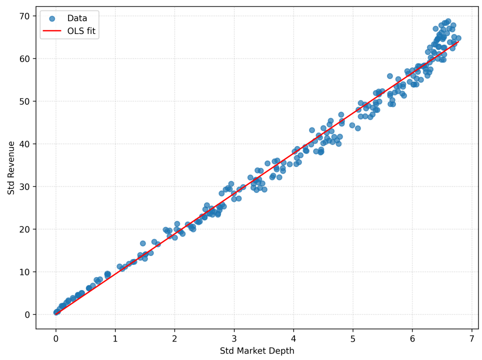

(d) Std market depth vs std revenue.

Figure 17: Relation between market depth and fee revenue for Experiment 1 and 2.

		Mean		Std	
Experiment	τ	Corr.	p-value	Corr.	p-value
Experiment 1	2	0.9991	< 0.001	0.9981	< 0.001
Experiment 2	3	0.9927	< 0.001	0.9935	< 0.001

Table 5: Correlation results between revenue and market depth.

		Mean		Std	
Experiment	τ	Slope	95% CI	Slope	95% CI
Experiment 1	2	13.586	[13.549, 13.623]	15.384	[15.324, 15.444]
Experiment 2	3	7.968	[7.914, 8.021]	8.733	[8.675, 8.791]

Table 6: Regression results between revenue and market depth.

Statistical Analysis In Table 5, we see that for both experiments, market depth and revenue are highly correlated. In all cases, the p-value is smaller than 0.01, which makes the relationship statistically significant. Therefore, we examine it further using linear regression as described above, again setting the y-intercept to 0.

In Table 6, we report the values m_E and m_Σ for experiments 1 and 2, and we observe that the 95% confidence interval is quite tight around those values in both experiments.

6.2.4 Step 4: Linear Relationship between Cost and Impermanent Loss

Description: Continuing our goal of approximating the risk-adjusted return using KPIs, we have

$$\frac{\mathbb{E}\left(X^{\text{profit}}\right)}{\sigma\left(X^{\text{profit}}\right)} \approx \frac{m_E \mathbb{E}(X^{\text{depth}(\tau)}) - \mathbb{E}(X^{\text{cost}})}{m_\sigma\left(m_\Sigma \sigma(X^{\text{depth}(\tau)}) + \sigma(X^{\text{cost}})\right)}.$$

In this step, we will see that

$$E^{\text{cost}} \approx E^{\text{IL}}, \quad \Sigma^{\text{cost}} \approx \Sigma^{\text{IL}},$$

and we can approximate the risk-adjusted return by

$$\frac{\mathbb{E}\left(X^{\text{profit}}\right)}{\sigma\left(X^{\text{profit}}\right)} \approx \frac{m_E \mathbb{E}(X^{\text{depth}(\tau)}) - \mathbb{E}(X^{\text{IL}})}{m_\sigma\left(m_\Sigma \sigma(X^{\text{depth}(\tau)}) + \sigma(X^{\text{IL}})\right)}.$$

Similar to Step 3, we now analyze the linear correlation between E^{cost} and E^{IL}, as well as between Σ^{cost} and Σ^{IL}. Based on the discussion in Section 5.2.1, we expect

$$E^{\text{cost}} = E^{\text{IL}}, \quad \Sigma^{\text{cost}} = \Sigma^{\text{IL}}.$$

Thus, impermanent loss is the only cost that occurs.

Data and Visualization: For data, we use E^{cost} and E^{IL}, as well as Σ^{cost} and Σ^{IL} as described in 6.1. In Figure 18, we see a nearly perfect linear relationship between E^{cost} and E^{IL}, as well as between Σ^{cost} and Σ^{IL}.

Experiment 1

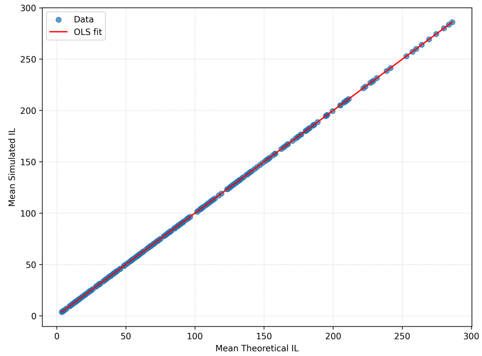

(a) Mean IL.

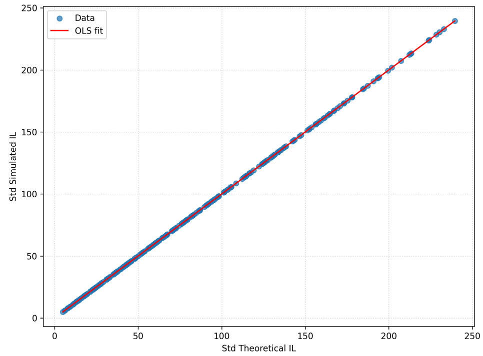

(b) Std IL.

Experiment 2

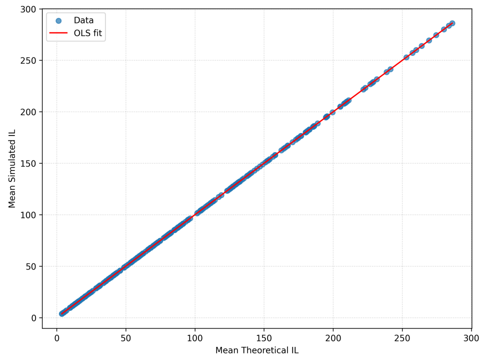

(c) Mean IL.

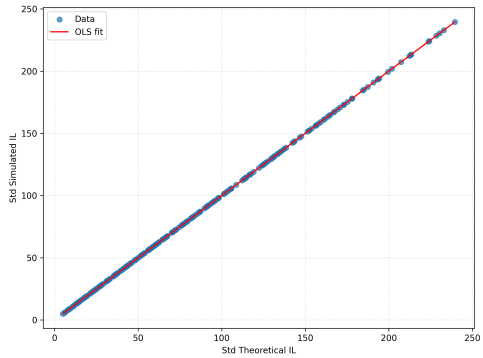

(d) Std IL.

Figure 18: Relation between cost and impermanent loss for Experiment 1 and 2.

Statistical Analysis: In Tables 7 and 8, we confirm our expectation and show that the cost can be reliably estimated using impermanent loss.

		Mean		Std	
Experiment	τ	Corr.	p-value	Corr.	p-value
Experiment 1	2	1.000	< 0.001	1.000	< 0.001
Experiment 2	3	1.000	< 0.001	1.000	< 0.001

Table 7: Correlation results between cost and impermanent loss.

		Mean		Std	
Experiment	τ	Slope	95% CI	Slope	95% CI
Experiment 1	2	1.000	[1.000, 1.000]	1.000	[1.000, 1.000]
Experiment 2	3	1.000	[1.000, 1.000]	1.000	[1.000, 1.000]

Table 8: Regression results between cost and impermanent loss.

6.2.5 Step 5: Sharpe Ratio Estimation and KPI-Based Estimators

Description: We are interested in finding a function g such that

$$\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})} \approx g\big(\mathbb{E}(X^{\text{IL}}), \sigma(X^{\text{IL}}), \mathbb{E}(X^{\text{depth}}), \sigma(X^{\text{depth}})\big).$$

Steps 1–4 have shown that the Sharpe ratio can be well approximated by

$$\frac{\mathbb{E}\big(X^{\text{profit}}\big)}{\sigma\big(X^{\text{profit}}\big)} \approx \frac{m_E\,\mathbb{E}\big(X^{\text{depth}(\tau)}\big) - \mathbb{E}\big(X^{\text{IL}}\big)}{m_\sigma\big(m_\Sigma\,\sigma\big(X^{\text{depth}(\tau)}\big) + \sigma\big(X^{\text{IL}}\big)\big)},$$

where m_E, m_Σ, and m_σ are obtained by linear regressions in Steps 2–4, and τ is the depth parameter chosen in Step 1. This provides a natural candidate for the functional form of g.

To evaluate how well this candidate predicts the simulated Sharpe ratio, we define

$$Z := \frac{m_E\,\mathbb{E}\big(X^{\text{depth}(\tau)}\big) - \mathbb{E}\big(X^{\text{IL}}\big)}{m_\sigma\big(m_\Sigma\,\sigma\big(X^{\text{depth}(\tau)}\big) + \sigma\big(X^{\text{IL}}\big)\big)},$$

and analyze how closely Z tracks the Sharpe ratio obtained from the full simulation. In particular, we first study the correlation between Z and the simulated Sharpe ratio and then fit a linear regression with zero or small intercept to quantify any systematic bias.

The parameters entering the construction of Z are summarized in Table 9.

	Experiment 1	**Experiment 2**
Depth parameter τ	2	3
m_E	13.586	7.968
m_Σ	15.384	8.733
m_σ	0.880	0.897

Table 9: Estimated parameters used in the KPI-based Sharpe ratio estimators.

Data and Visualization: For each AMM, we compute, on the one hand, the simulated Sharpe ratio $\frac{\mathbb{E}(X^{\text{profit}})}{\sigma(X^{\text{profit}})}$ and, on the other hand, the approximation Z defined above. To construct Z, we use

$$\sigma(X^{\text{IL}}),\ \sigma(X^{\text{depth}(\tau)}),\ \mathbb{E}(X^{\text{IL}}),\ \mathbb{E}(X^{\text{depth}(\tau)})$$

as described in Section 6.1, together with the parameters τ, m_E, m_Σ, and m_σ obtained in Steps 1–4. A comparison of Z and the simulated Sharpe ratio is shown in Figure 19. Visually, the points lie close to the 45-degree line, indicating that the KPIs provide a good approximation of the Sharpe ratio, with a few noticeable deviations for AMMs with very low market depth.

Experiment 1

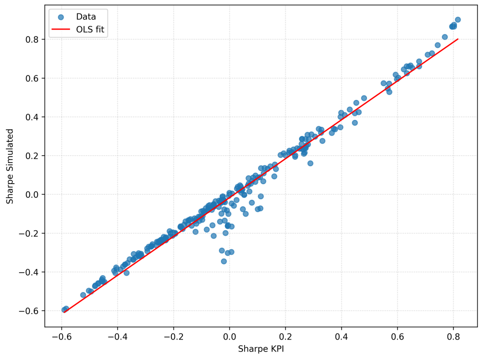

(a) Using all data points

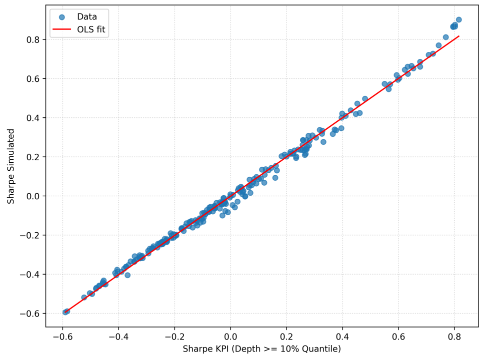

(b) Use a subset of AMM

Experiment 2

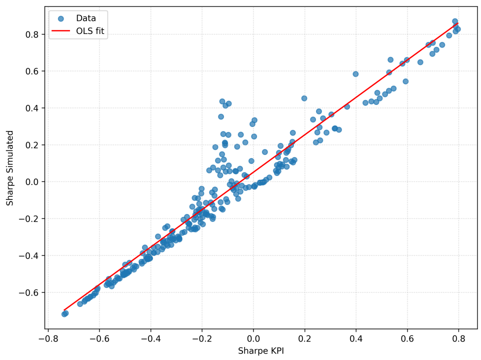

(c) Using all data points

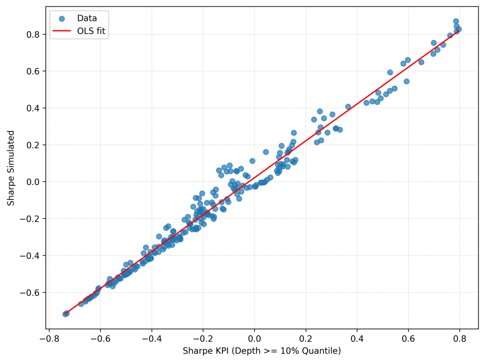

(d) Use a subset of AMM

Figure 19: Comparison of Sharpe ratios: Once we use all 256 AMMs and once we use only the top 90% percent tile with respect to the market depth.

Statistical Analysis: Table 10 shows that the approximation performs well when all AMMs are included. The correlation between Z and the simulated Sharpe ratio exceeds 0.95 in both experiments, with a p-value below 0.001. The regression slope is slightly above 1, suggesting that Z tends to underestimate the true Sharpe ratio on average.

When we exclude the 10% of AMMs with the lowest market depth, Table 11 shows that the correlation increases further and the 95% confidence interval of the regression slope becomes tighter. In Figure 19, the biggest outliers disappear once we focus on AMMs with sufficient market depth. This indicates that the KPI-based approximation works particularly well for AMMs with moderate to high market depth.

		Correlation		Regression			
Experiment	τ	Corr.	p-value	Slope	95% CI	y-Intercept	95% CI
Experiment 1	2	0.984	< 0.001	1.002	[0.980, 1.024]	-0.017	[-0.024, -0.010]
Experiment 2	3	0.957	< 0.001	1.014	[0.976, 1.052]	0.051	[0.037, 0.064]

Table 10: Correlation and regression between simulated Sharpe ratio and KPI-based Sharpe ratio using all AMMs.

		Correlation		Regression			
Experiment	τ	Corr.	p-value	Slope	95% CI	y-Intercept	95% CI
Experiment 1	2	0.997	< 0.001	1.003	[0.993, 1.013]	-0.002	[-0.005, 0.001]
Experiment 2	3	0.990	< 0.001	1.000	[0.981, 1.018]	0.021	[0.014, 0.028]

Table 11: Correlation and regression between simulated Sharpe ratio and KPI-based Sharpe ratio using top 90% AMMs regarding the market depth.

Resulting Sharpe Ratio Estimators g_1 and g_2: We now use our results to provide explicit estimators g_1 and g_2 for experiments 1 and experiment 2, respectively. These estimators are obtained by inserting the empirically estimated parameters from Table 9 into the expression for Z, together with the linear adjustment from the regression of the simulated Sharpe ratio on Z.

For experiment 1, we obtain

$$\begin{aligned} &g_1\big(\mathbb{E}(X^{\mathrm{IL}}), \sigma(X^{\mathrm{IL}}), \mathbb{E}(X^{\mathrm{depth}}), \sigma(X^{\mathrm{depth}})\big) \\ &= 1.003 \cdot \frac{13.586\, \mathbb{E}\big(X^{\mathrm{depth}(2)}\big) - \mathbb{E}\big(X^{\mathrm{IL}}\big)}{0.880\,\big(15.384\, \sigma\big(X^{\mathrm{depth}(2)}\big) + \sigma\big(X^{\mathrm{IL}}\big)\big)} - 0.002. \end{aligned}$$

For experiment 2, we analogously obtain

$$\begin{aligned} &g_2\big(\mathbb{E}(X^{\mathrm{IL}}), \sigma(X^{\mathrm{IL}}), \mathbb{E}(X^{\mathrm{depth}}), \sigma(X^{\mathrm{depth}})\big) \\ &= 1.000 \cdot \frac{7.968\, \mathbb{E}\big(X^{\mathrm{depth}(3)}\big) - \mathbb{E}\big(X^{\mathrm{IL}}\big)}{0.897\,\big(8.733\, \sigma\big(X^{\mathrm{depth}(3)}\big) + \sigma\big(X^{\mathrm{IL}}\big)\big)} + 0.021, \end{aligned}$$

for AMMs where market depth is not low, i.e., those in the upper 90% quantile of the depth distribution. For the lower 10% of AMMs, we know from the regression results that g_1 tends to overestimate and g_2 tends to underestimate the true Sharpe ratio. For the vast majority of AMMs in our sample, however, g_1 and g_2 provide accurate KPI-based approximations of the risk-adjusted profitability.

7 Discussion

Summary of the Findings and Their Relevance

The aim of this thesis was to explore how Automated Market Makers (AMMs) can be designed to allow negative prices and how their performance can be understood using key performance indices of the invariant. For this purpose, the Power Sum Invariant was introduced as a class of invariants that allows both positive and negative spot prices.

Additionally, a simulation environment was built that combines fair-value paths of the underlying (generated with an Offset-GBM process) with an agent-based trading setup consisting of arbitrage and noise traders. This allowed us to measure fee revenue, simulated impermanent loss, and ultimately the risk-adjusted profit of a liquidity provider for a set of 256 different invariants. In parallel, the two key performance indicators—market depth and theoretical impermanent loss—were computed for each invariant in the same market environment.

When comparing these KPIs with the simulation results, a clear pattern emerged: the Sharpe ratio of the liquidity provider appears to be largely determined by a combination of aggregated market depth and theoretical impermanent loss. The estimator g introduced in Section 6.2.5 captures this relation reasonably well and provides a simple way to approximate the risk-adjusted return directly from the invariant's parameters. The behavior was consistent across the 256 sampled invariants and suggests that these KPIs capture essential structural features of the AMM's performance in environments where prices can become negative. It is important to note, however, that this approximation worked reliably only for invariants whose market depth lay within roughly the top 90% of the sampled range. For invariants with very low market depth, the estimator showed large deviations.

Limitations

The results rely on several simplifying assumptions. Most importantly, the fair-value dynamics of the underlying were modeled using GBM. While this is mathematically convenient, it does not fully capture the empirical behavior of financial assets, such as jumps or time-varying volatility. The noise traders in the simulation follow a simple rule-based behavior, which does not reflect the diversity or adaptiveness of real market participants.

Another limitation is that the liquidity level of the AMM is assumed to be constant throughout the one-year simulation horizon, and we fixed it at 30 across all invariants. In reality, liquidity providers react to market conditions, expected returns, and collateral requirements, which is especially relevant in settings with potentially negative prices. The simulation therefore evaluates a static environment and does not account for changes in liquidity over time.

Furthermore, the invariant family studied in this thesis is intentionally narrow. The Power Sum Invariant provides a tractable class with clear mathematical structure, but it is only one possible approach to designing AMMs that allow negative prices. The thesis therefore cannot identify an optimal invariant and does not claim that such an invariant exists within the tested class.

Finally, the observed relationship between KPIs and the Sharpe ratio is derived within the specific simulation setup used here. It is unclear how robust this relationship is under alternative market models, trader behavior, or different types of stochastic processes. The results should therefore be seen as an indication of a structural link rather than a general statement.

Further Work

There are several natural directions for future research. One possibility is to replace the GBM-based fair-value process with more realistic stochastic models and to study whether the KPI-based approximation remains valid. More sophisticated noise-trader models could also be introduced, potentially allowing calibration to empirical trading patterns.

Another direction is to explore broader families of invariants, for example invariants with different curvature properties or alternative boundary behavior. If the KPI-based relationship continues to hold, this could help identify classes of invariants that tend to produce better risk-adjusted outcomes.

Finally, future work could focus on modeling dynamic liquidity provision. In practice, liquidity providers adjust their positions in response to market conditions, expected returns, and the risk of holding offset tokens. Incorporating such entry and exit decisions into the simulation would make it possible to analyze how changing liquidity levels interact with the invariant and whether the relationships identified in this thesis persist once liquidity is no longer fixed.

References

[1] *Annual Crypto Industry Report*. URL: https://assets.coingecko.com/reports/2024/CoinGecko-2024-Annual-Crypto-Industry-Report.pdf (visited on Aug. 4, 2025).

[2] Siemens. *Siemens Issues First Digital Bond on Blockchain*. URL: https://press.siemens.com/global/en/pressrelease/siemens-issues-first-digital-bond-blockchain (visited on Aug. 4, 2025).

[3] *Price: The Language of Exchange*. URL: https://www.imf.org/external/pubs/ft/fandd/basics/28-price.htm?utm_source=chatgpt.com (visited on Dec. 18, 2025).

[4] J. Hull. *Options, Futures, and Other Derivatives, EBook, Global Edition*. Harlow, UNITED KINGDOM: Pearson Education, Limited, 2017. ISBN: 978-1-292-21292-0. URL: http://ebookcentral.proquest.com/lib/aalto-ebooks/detail.action?docID=5186416 (visited on Sept. 4, 2025).

[5] *Ethereumbook*. original-date: 2016-08-10T15:07:54Z. Nov. 19, 2025. URL: https://github.com/ethereumbook/ethereumbook (visited on Nov. 20, 2025).

[6] E. I. Proposals. *ERC-20: Token Standard*. Ethereum Improvement Proposals. URL: https://eips.ethereum.org/EIPS/eip-20 (visited on Dec. 9, 2025).

[7] J. V. Burke. *Convexity (Lecture Notes) (University of Washington)*. URL: https://sites.math.washington.edu//~burke/crs/cvx15/notes/RoW_convexity.pdf.

[8] C. Y. Yan et al. *Better market Maker Algorithm to Save Impermanent Loss with High Liquidity Retention*. version: 1. Feb. 27, 2025. DOI: 10.48550/arXiv.2502.20001. arXiv: 2502.20001[q-fin]. URL: http://arxiv.org/abs/2502.20001 (visited on Aug. 12, 2025).

[9] V. Tolstikov. *Concentrated Superelliptical Market Maker*. Nov. 26, 2024. DOI: 10.48550/arXiv.2410.13265. arXiv: 2410.13265[q-fin]. URL: http://arxiv.org/abs/2410.13265 (visited on Aug. 3, 2025).

[10] V. Tolstikov et al. *Concentrated N-dimensional AMM with Polar Coordinates in Rust*. Oct. 6, 2025. DOI: 10.48550/arXiv.2510.05428. arXiv: 2510.05428[q-fin]. URL: http://arxiv.org/abs/2510.05428 (visited on Nov. 1, 2025).

[11] *Negative Price AMM - Decentralized Exchanges*. Ethereum Research. Section: Decentralized exchanges. Dec. 15, 2023. URL: https://ethresear.ch/t/negative-price-amm/17833 (visited on Nov. 1, 2025).

[12] D. Toth and B. Jones. *Against the Norm: Modeling Daily Stock Returns with the Laplace Distribution*. June 25, 2019. DOI: 10.48550/arXiv.1906.10325. arXiv: 1906.10325[q-fin]. URL: http://arxiv.org/abs/1906.10325 (visited on Nov. 4, 2025).

[13] G. Angeris et al. "An Analysis of Uniswap markets". In: *Cryptoeconomic Systems* 0.1 (Apr. 7, 2021). Publisher: Metagov. ISSN: 2767-4207, DOI: 10.21428/58320208.c9738e64. URL: https://cryptoeconomicsystems.pubpub.org/pub/angeris-uniswap-analysis/release/15 (visited on Aug. 3, 2025).

[14] W. F. Sharpe. "Mutual Fund Performance". In: *The Journal of Business* 39.1 (1966). Publisher: University of Chicago Press, pp. 119–138. ISSN: 0021-9398. URL: https://www.jstor.org/stable/2351741 (visited on Nov. 7, 2025).

[15] D. Engel and M. Herlihy. "Presentation and Publication: Loss and Slippage in Networks of Automated Market Makers". In: *OASIcs, Volume 97, Tokenomics 2021* 97 (2022). Ed. by V. Gramoli, H. Halaburda, and R. Pass. Artwork Size: 23 pages, 847283 bytes ISBN: 9783959772204 Medium: application/pdf Publisher: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 13:1–13:23. ISSN: 2190-6807. DOI: 10.4230/OASICS.TOKENOMICS.2021.13. URL: https://drops.dagstuhl.de/entities/document/10.4230/OASIcs.Tokenomics.2021.13 (visited on Aug. 3, 2025).

[16] S. Loesch et al. *Impermanent Loss in Uniswap v3*. Nov. 17, 2021. DOI: 10.48550/arXiv.2111.09192. arXiv: 2111.09192[q-fin]. URL: http://arxiv.org/abs/2111.09192 (visited on Aug. 19, 2025).

[17] G. Liao and D. Robinson. "The Dominance of Uniswap v3 Liquidity". In: *Uniswap* (). URL: https://app.uniswap.org/TheDominanceofUniswapv3Liquidity.pdf.

[18] A. Kempf and O. Korn. "Market Depth and Order Size: An Analysis of Permanent Price Effects of DAX Futures' Trades". In: Zentrum für Europäische Wirtschaftsforschung (ZEW) (). URL: https://www.econstor.eu/bitstream/10419/24262/1/dp1098.pdf.

A Additional Figures: Pseudo Code

```
contract MinimalAMM {
    ERC20 public immutable alpha;
    ERC20 public immutable beta;
    uint256 private kappa; // fee in per mille
    address liquidity_provider;
    function f(uint256 rAlpha) public view returns (uint256) {
        // Implementation of invariant
        return rBeta;
    }
    function buy_sell_alpha(int256 dalpha) external returns bool
         {
        // 1) Read pool reserves
        uint256 rAlpha = alpha.balanceOf(address(this));
        uint256 rBeta = beta.balanceOf(address(this));
        // 2) Compute new pool state and trader delta
        int256 dBetaPool = int256(f(rAlpha + uint256(dalpha))) -
             int256(rBeta);
        uint256 fee = (abs(dalpha) * kappa) / 1000;
        int256 dBetaTrader = dBetaPool + int256(fee);
        // 3) Apply token flows
        collect_or_payout(alpha, dalpha);
        collect_or_payout(beta, dBetaTrader);
        beta.transfer(liquidity_provider, fee);
        return true;
    }
    function collect_or_payout(ERC20 token, int256 amount)
        internal {
        if (amount > 0) {
            // Trader sends tokens to the pool
            token.transferFrom(msg.sender, address(this), uint256
                (amount));
        } else if (amount < 0) {
            // Pool sends tokens to the trader
            token.transfer(msg.sender, uint256(-amount));
        }
    }
}
```

Figure 20: Minimalistic AMM as a smart contract

```
contract ERC20 {
    uint8 public decimals = 18;
    uint256 public totalSupply;
    mapping(address => uint256) public balanceOf;
    mapping(address => mapping(address => uint256)) public
        allowance;

    constructor(uint256 initialSupply) {
        totalSupply = initialSupply * (10 ** uint256(decimals));
        balanceOf[msg.sender] = totalSupply;
        emit Transfer(address(0), msg.sender, totalSupply);
    }
    function _transfer(address _from, address _to, uint256
        _value) internal {
        require(_to != address(0));
        require(balanceOf[_from] >= _value);
        unchecked {
            balanceOf[_from] -= _value;
            balanceOf[_to] += _value;
        }
    }
    function transfer(address _to, uint256 _value) external
        returns (bool) {
        _transfer(msg.sender, _to, _value);
        return true;
    }
    function approve(address _spender, uint256 _value) external
        returns (bool) {
        allowance[msg.sender][_spender] = _value;
        return true;
    }
    function transferFrom(address _from, address _to, uint256
        _value) external returns (bool) {
        uint256 allowed = allowance[_from][msg.sender];
        require(allowed >= _value);
        unchecked {
            allowance[_from][msg.sender] = allowed - _value;
        }
        _transfer(_from, _to, _value);
        return true;
    }
}
```

Figure 21: Minimal ERC20 Token as a smart contract