

The Potential of Self-Regulation for Front-Running Prevention on Decentralized Exchanges

Lioba Heimbach
ETH Zurich
hlioba@ethz.ch

Eric Schertenleib
unaffiliated
eric.schertenleib@gmail.com

Roger Wattenhofer
ETH Zurich
wattenhofer@ethz.ch

Abstract

The transaction ordering dependency of the smart contracts building *decentralized exchanges (DEXes)* allow for predatory trading strategies. In particular, front-running attacks present a constant risk for traders on DEXes. Whereas legal regulation outlaws most front-running practices in traditional finance, such measures are ineffective in preventing front-running on DEXes. While novel market designs hindering front-running may emerge, it remains unclear whether the market’s participants, in particular, liquidity providers, would be willing to adopt these new designs. A misalignment of the participant’s private incentives and the market’s social incentives can hinder the market from adopting an effective prevention mechanism.

We present a game-theoretic model to study the behavior of sophisticated traders, retail traders, and liquidity providers in DEXes. Sophisticated traders adjust for front-running attacks, while retail traders do not, likely due to lack of knowledge or irrationality. Our findings show that with less than 1% of order flow from retail traders, traders’ and liquidity providers’ interests align with the market’s social incentives – eliminating front-running attacks. However, the benefit from embracing this novel market is often small and may not suffice to entice them. With retail traders making up a larger proportion (around 10%) of the order flow, liquidity providers tend to stay in pools that do not protect against front-running. This suggests both educating traders and providing additional incentives for liquidity providers are necessary for market self-regulation.

1 Introduction

The emergence of *decentralized finance (DeFi)* on Ethereum [47] greatly enhanced the interest in cryptocurrency applications. DeFi is a blockchain-based form of finance that utilizes *smart contracts* to offer many traditional financial instruments, but without relying on financial intermediaries. A prime example thereof is *decentralized exchanges (DEXes)*.

While traditional exchanges match individual buyers and sellers with the limit order book mechanism, a DEX algorithmically sets the exchange rate for a trade. To this end, DEXes store liquidity for exchanges between individual cryptocurrency pairs in smart contracts, referred to as *liquidity pools*. The trade size and the respective cryptocurrency pair’s amount and ratio of reserves control the price. The pool charges a fee for every trade which is proportional to the trade’s input amount and distributed pro-rata amongst the pool’s *liquidity providers*.

DEXes are becoming increasingly popular. Yet, the rise of DEXes does not come without caveats, leading to the characterization of the Ethereum peer-to-peer network as a dark forest. Predatory trading bots prey on user transactions in Ethereum’s *mempool*, the public waiting area for transactions. Predatory trading schemes exploit the lack of privacy given to transactions prior to their execution. Moreover, the smart contracts that build DEXes are dependent on the transaction order. Generally, these attacks involve the attacker *front-running* a victim’s transaction. One of the most frequently observed strategies exploiting this dependency on transaction ordering is the *sandwich attack* [42] which we describe in Section 3. We focus on sandwich attacks in this work, as this is the only front-running attack directly impacting the welfare of liquidity providers.¹ Such an attack occurs when a trading bot front- and back-runs a victim’s transaction, forcing the trade to execute at an unfavorable price. Between 1 August and 31 August 2024, more than 130,000 transactions were sandwiched on Ethereum blockchain’s DEXes [11].

Given the severity of front-running attacks on DEXes, the market is seeking mechanisms that can prevent such attacks. While front-running attacks are outlawed in traditional finance, the anonymity of market participants and the absence of a central authority do not allow for an effective regulatory approach for DEXes. Therefore, they require new innova-

¹Apart from sandwich attacks, there are destructive front-running attacks [28]. Thereby, the attacker searches for trades that exploit arbitrage opportunities and front-runs these and is indifferent to whether the victim’s transaction executes. The DEXes volume remains unchanged.

tive solutions to prevent front-running attacks. In response, multiple approaches to prevent front-running and, more generally, transaction reordering manipulation are currently under development [30]. These approaches generally ensure that transaction contents are private from the public until an order is agreed upon. Further, some designs are already being adopted [3, 6, 7].

Even though successful market designs preventing front-running attacks increase market efficiency as a whole [35], there may still be insufficient incentives to adopt these prevention schemes. Liquidity providers, who potentially benefit from the additional trading volume from sandwich attacks, might be reluctant to embrace such a market and, similarly, some traders might be slow to adapt to such new markets. Examples, where the divergence of private and social incentives has led to adoption failure exist in traditional limit order book exchanges [21].

In this work, we develop a game-theoretical model to study whether traders and liquidity providers would embrace DEXes that incorporate a front-running prevention mechanism. Our repeated game consists of two hypothetical liquidity pools, one allowing sandwich attacks and one implementing a sandwich attack prevention scheme. Traders and liquidity providers distribute across the two pools, thereby maximizing their private incentives. The objective of liquidity providers is to maximize their fees earned while the traders seek to execute their orders at the best possible price. Thus, while the presence of sandwich attacks leads to additional trades from the attack, it reduces the volume of ordinary traders as they receive a poor price.

To capture different behavior traders, we distinguish between two types of traders. A portion of the trade orders originates from *sophisticated* traders who are aware of sandwich attacks and adjust their behavior accordingly (fully rational), whereas, another fraction of trades stem from *retail* traders who are either oblivious or indifferent to these attacks. This distinction is made to account for information asymmetry and a degree of irrationality present in the behavior of users.

Contributions. We summarize our contributions below:

- *Small Retail Volume.* We find that a parameter regime exists for which the players utilize the liquidity pool with front-running prevention – indicating that the market can fix itself. This regime dominates when we have a very small proportion of order flow stemming from retail traders. However, even when the private incentives of liquidity providers align with the market’s social incentives, the benefit of embracing the new market can be small.
- *Significant Retail Volume.* When the proportion of retail traders increases, i.e., we observe more irrational behavior from traders, market conditions where the private incentives of liquidity providers and the system’s social incentives are misaligned become more common.
- *Moving Towards Self-Regulation.* Finally, we highlight the

importance of educating traders and point towards additional incentives the market could provide to entice liquidity providers to adjust their strategy.

2 Related Work

Front-running has long been prevalent in traditional finance [19, 23, 26] where the regulator is tasked with banning such practices [36, 39]. The lack of a central authority in DeFi, however, means that the market must regulate itself. Thus, we study whether the private incentives of market participants obstruct the adoption of innovative market designs preventing front-running.

Eskandir et al. [28] are the first to systematize work surrounding front-running on DeFi. In a similar line of research, Daian et al. [25] study the risks of front-running on DEXes. They observe traditional forms of predatory trading behaviors adapting to the blockchain ecosystem. Park [41] further shows that the pricing rule of most DEXes gives rise to intrinsically profitable front-running opportunities. By analyzing the market participants’ private incentives to prevent front-running, we build on these earlier works and, in particular, study sandwich attacks, as they influence the welfare of traders and liquidity providers.

The prevalence of front-running attacks on DEXes is first quantified by Qin et al. [42]. Zhou et al. [51] study sandwich attacks both analytically and empirically. Both demonstrate the risk stemming from front-running attacks on DEXes. Our work, on the other hand, focuses on whether the market participant’s private incentives are disruptive to the adoption of market designs preventing front-running attacks and the associated risks.

Recently, many suggestions for DEX front-running prevention schemes have emerged. For a comprehensive overview, we refer the reader to Heimbach and Wattenhofer [30]. Their work compares state-of-the-art prevention mechanisms and finds that current schemes do not meet the blockchain ecosystem’s requirements.

We summarise the core ideas behind the most relevant suggestions in the following. The simplest schemes, tune the transaction parameters to prevent specific attacks on specific protocols [29, 50]. Further, several suggestions propose that transactions are sent to a trusted third party that is put in charge of ordering the transactions fairly [2, 4, 6–9, 12, 18, 44]. A parallel line of work, instead of relying on a single entity, trust a generally permissioned committee to order the transactions in a fair manner [16, 17, 22, 24, 31–34, 37, 38, 40, 43, 48, 49] — preventing front-running. Finally, several schemes set the order of transactions by first having users commit to their transactions on-chain and then only revealing the transaction contents later in a second phase [20, 27, 45]. All of these schemes, thus, aim to preserve the privacy of the transaction contents until an execution order is agreed upon. In this work, instead of assessing or designing prevention mechanisms, we

study a market with an ideal prevention mechanism to analyze whether the private incentives would steer market participants to accept such a market design.

Budisch et al. [21] examine the incentives of exchanges to embrace market design innovations that eliminate latency arbitrage and HFT trading. Their work finds that adoption failures arise in traditional limit order book exchanges. In particular, the divergence of private and social incentives hinders the market from accepting new market designs. We study the incentives of liquidity providers in DEXes and show that their private incentives generally align with the system’s social interests: demonstrating that liquidity providers can be incentivized to adopt new market designs preventing front-running attacks.

3 Preliminaries

In the following, we detail the trading mechanism of the biggest DEXes and introduce the sandwich attack specifics.

3.1 Automated Market Maker

As its name suggests, trade execution on *automated market makers (AMMs)* is automatic, and the price is controlled by an algorithm with liquidity being supplied by individual liquidity providers rather than brokers or market makers. A host of AMM variants exist, each with its specific pricing mechanism. We focus on the most widely adopted subclass of AMMs: *constant product market makers (CPMMs)* [5]. For each tradeable cryptocurrency pair, the CPMM stores assets of both cryptocurrencies in a liquidity pool. The CPMM then guarantees that the product between the amounts of the two reserved pool currencies stays constant. This property ensures that the price for swapping between these pairs mimics the behavior of a demand curve of a normal good. Both Uniswap and Sushiswap, two of the biggest DEXes, employ the CPMM for pricing. The original CPMM design, as deployed by Uniswap V2 [14] and Sushiswap [13], utilizes the same liquidity for the pool’s entire price range. Consider a pool $X \rightleftharpoons Y$ between X -tokens and Y -tokens with respective reserves x and y . Then the pool’s marginal price indicating the pool’s current price for X -token in terms of Y -token is $P = y/x$ [14]. Further, the pool’s liquidity is defined as $L = \sqrt{x \cdot y}$. This liquidity needs to support trading along the entire price range $(0, \infty)$ in the original CPMM design.

In the newest Uniswap design (V3) [15], liquidity providers choose the price range $[P_a, P_b]$ for which they provide liquidity. This concentration of liquidity is intended to increase capital efficiency, as the liquidity only needs to support trade execution in the corresponding price range. Liquidity providers can only choose from a discrete set of price range boundaries that are defined by the pool’s initialized *ticks*. Between each pair of initialized ticks, the CPMM only needs to maintain

enough reserves to support trading between the price boundaries. One can simulate a constant product pool with adjusted larger reserves, referred to as *virtual reserves*, between any pair of neighboring initialized ticks.

For a price range $[P_a, P_b]$ between two neighboring initialized ticks, the *liquidity* inside the tick is given by L and the *marginal price* is P . The CPMM ensures that the constant product of the virtual reserves x and y stays constant, i.e., $x \cdot y = k = L^2$, where k is the constant product of the reserves in the considered price interval. As on Uniswap V2, the marginal price is $P = y/x$ and the liquidity is $L = \sqrt{x \cdot y}$ [15]. Thus, the virtual reserves are then given by $x = L/\sqrt{P}$ and $y = L \cdot \sqrt{P}$. For the sake of simplicity, we focus on trading between two neighboring initialized ticks and refer to virtual reserves simply as reserves in this work.²

The exchange rate received by traders is dependent on their trade size and the number of tokens reserved in the liquidity pool. Consider, again, a liquidity pool between X -token and Y -token, $X \rightleftharpoons Y$. We denote the respective initial (virtual) reserves prior to any trading as x and y , respectively, and use ξ and υ for the reserves of X and Y at any time during the trading process. Thus, if a trader adds an infinitesimal amount $d\xi$ to the pool, the following amount $d\upsilon$ of Y is extracted where

$$d\upsilon = -\frac{xy}{\xi^2} d\xi.$$

This expression follows directly from the constant product property. Note that the sign convention we choose is relative to the pool, i.e., $\Delta\upsilon < 0$ for a trader buying Y -tokens. Further, observe that the price per Y -token increases with the input amount. Thus, traders have to pay more per desired token the larger their trade is, resulting in *expected slippage* which is the difference between the pool’s marginal price and the actual price received by the trader. Note that the expected slippage is lower in more liquid pools, i.e., those with larger stored reserves.

From the infinitesimal price, it follows that a trader wishing to sell δ_x X -tokens will receive δ_y Y -tokens, where

$$\begin{aligned} \delta_y &= -\int_x^{x+(1-f)\delta_x} \frac{-x \cdot y}{\xi^2} d\xi \\ &= y - \frac{x \cdot y}{x + (1-f)\delta_x} = \frac{y(1-f)\delta_x}{x + (1-f)\delta_x}, \end{aligned}$$

and f is the transaction fee which is charged relative to the input amount δ_x and is distributed pro-rata to the tick’s liquidity providers. The sign in front of the integral is negative as the trader receives the Y -tokens extracted from the pool. Note that the $(1-f)\delta_x$ in the upper integral bound corresponds

²As our analysis focuses on a time frame where the fair market price between X - and Y -tokens remains constant (cf. Section 4), trading is also likely to occur within a small price range and thus will likely remain within one tick. To cover trading across ticks, one can reapply our analysis. Note that trading within a tick on Uniswap V3 is mostly the same as on Uniswap V2.

to the amount of X added to the pool after deduction of the transaction fee. Thus, post-execution the reserves of X and Y will be $x + (1 - f)\delta_x$ X -token and $y - \delta_y$ Y -tokens.

The time at which a trade executes is unclear to the trader, as their transactions will only be confirmed upon block inclusion. In the meantime, other transactions changing the pool’s state might occur. The change in the pool’s state introduces a difference between the trader’s expected price at the time of submission and the actual price at the time of execution. This price change is known as *unexpected slippage*. To ensure the price of the transaction does not deviate significantly from the expectation, traders specify a *slippage tolerance*, indicating the maximum unexpected price movement they are willing to accept. A trade expecting δ_y Y -tokens at the time of transaction submission will only execute if it receives no less than $(1 - s)\delta_y$ Y -tokens, where s is the slippage tolerance. Typical slippage tolerances are $s < 0.03$ [46].

3.2 Sandwich Attacks

A too small slippage tolerance results in frequent transaction failure. However, the slippage tolerance also gives an opening for sandwich attacks. On Ethereum, users broadcast their transactions to the network. The transaction waits in the mempool until it is included in a block by a validator. During this time, the transaction is visible to predatory trading bots and runs the risk of being front- and back-run as part of a sandwich attack. Predatory trading bots scan the mempool’s inflowing transaction stream searching for profitable sandwich attack opportunities.

As validators³ control the ordering in a block, sandwich attackers can provide validators with the necessary (financial) incentives to achieve their desired transaction ordering. In fact, *front-running-as-a-service* schemes, such as Flashbots [7] and Eden network [6], facilitate this interaction between sandwich attackers and validators. On the other hand, these services can also be used for front-running prevention, but users must deliberately seek them out rather than being truly incorporated into the market design.

A sandwich attack involves the attacker front-running the victim’s transaction, exchanging X -token for Y -token in transaction A_F . The attacker’s front-running transaction purchases the same asset as the victim: Y -token. Thereby, the attacker drives up the asset Y ’s price. The following victim transaction T then buys Y -token at a higher price and further inflates Y ’s price. To conclude the attack, the attacker back-runs the victim’s transaction, selling Y -assets at the inflated price with transaction A_B .

To provide a conceptual understanding of sandwich attacks, we visualize a victim’s trade T without a sandwich attack

³Note that currently most blocks are built through the proposer-builder separation scheme, where block building is outsourced to the specialized builders [10]. The same reasoning we detail for the validator also applies to these builders.

in Figure 1a. Figure 1b then shows how a sandwich attack alters the transaction T . We observe that without the sandwich attack, the victim expects a greater output δ_y (cf. Figure 1a) than the output $\tilde{\delta}_y$ it receives in the presence of a sandwich attack (cf. Figure 1b). The attacker’s front-running inflates Y -asset’s price. Further, we observe that the attacker’s output a_x^{out} of the back-running transaction A_B exceeds the attacker’s input a_x^{in} (cf. Figure 1b). The difference $a_x^{\text{out}} - a_x^{\text{in}}$ presents the attacker’s profit, as the attacker’s input a_y to transaction A_B is the output of transaction A_F .

Lastly, we note that at first glance liquidity providers appear to benefit from sandwich attacks as they lead to increased trading volume, and therefore, collected fees. However, traders aware of this threat could reduce their trading activity, as they receive a worse price than the market price if they fall victim to the attack. We will study this interplay by analyzing the effects of sandwich attacks on the utility of both traders and liquidity providers.

4 Model

We model a system with two liquidity pools Pool_N and Pool_W . Both pools facilitate exchanges for the same cryptocurrency pair: $X \rightleftharpoons Y$. While a scheme to prevent sandwich attacks is implemented in Pool_N , sandwich attacks are common practice in Pool_W . With our model, we will study whether DEX participants are able to self-regulate and adopt a DEX with front-running prevention in place.


Our model has four types of players: (sophisticated and retail) traders, liquidity providers, sandwich attackers, and price arbitrageurs. Traders and liquidity providers strive to maximize their personal utility (cf. Section 4.3 and Section 4.4) across two liquidity pools Pool_N and Pool_W . In maximizing their utilities, sophisticated traders and liquidity providers account for the effects of sandwich attacks and price arbitrageurs. Our model also captures the effects of less sophisticated traders who are oblivious to sandwich attacks. We will call this group retail traders. Further, for liquidity providers, we will also consider the consequences of them being inert.

Without sandwich attacks, trades in Pool_N execute at the expected price.⁴ In Pool_W , on the other hand, sandwich attack bots make an attack whenever it is profitable (cf. Section 4.1). We denote the fraction of the total liquidity placed in Pool_N by p . Thus, the reserves in Pool_N are $x_N = p \cdot x$ X -tokens and $y_N = p \cdot y$ Y -tokens, and $x_W = (1 - p)x$ X -tokens and $y_W = (1 - p)y$ Y -tokens in Pool_W . Given the price $P_{X \rightarrow Y}$ of X -token, we have


$$P_{X \rightarrow Y} = \frac{y}{x} = \frac{y_N}{x_N} = \frac{p \cdot y}{p \cdot x} = \frac{y_W}{x_W} = \frac{(1 - p)y}{(1 - p)x}.$$

We emphasize that the social incentives of our system are

⁴While it is possible for there to be several trades in a single block, we can assume them to only amount to natural price fluctuations. In the time frame of a block, they can be assumed to be negligible [29].



(a) The execution of transaction T without a sandwich attack. The transaction T receives the Y -assets at the expected price.



(b) The execution of transaction T with a sandwich attack. Transaction T is first front-run by transaction A_F and then back-run by transaction A_B .

Figure 1: Execution of victim transaction T in pool $X = Y$. without (cf. Figure 1a) and with (cf. Figure 1b) sandwich attack. In the presence of an attack the trader receives fewer Y -tokens $\tilde{\delta}_y < \delta_y$ while the attacker makes a profit, i.e., $a_x^{\text{in}} < a_x^{\text{out}}$.

to completely adopt Pool_N , the pool without front-running. In the presence of sandwich attacks in Pool_W , the trades of ordinary traders do not execute at the effective market price but rather at an unfavorable rate. Further, we purposefully exclude incentives of sandwich attackers and price arbitrageurs when discussing the system's incentives. Including their incentives would turn the game into a zero-sum game. Thus, in the presence of profitable sandwich attacks and price arbitrages, the remaining market participants (traders and liquidity providers) collectively lose money.

Further note that throughout, we assume that the transaction fee f ($0 < f < 1$) is identical in both pools. Additionally, we disregard the *gas fee*, the fee paid to validators for block inclusion on the Ethereum blockchain, for all players in our analysis. The gas fee would add a fixed cost to every trade and liquidity movement. However, for the sake of simplicity and as the gas fee is not part of the CPMM market mechanism itself, we assume it to be zero.

4.1 Sandwich Attackers

Sandwich attackers observe the inflowing transactions in Pool_W . Upon noticing a trade entering the mempool of Pool_W , they compute the maximal input for the sandwich attack and assess the attack's profitability. The maximal input infers the maximal acceptable price movement on the trader, such that the trade still executes. Attackers conduct any such profitable attack. We find the maximal input of a sandwich attack and study their profitability in Section 5.1.

The victim submits an order T_W wishing to exchange $\delta_{x,W} > 0$ X -tokens in Pool_W for Y -tokens and sets a slippage tolerance s . When submitting the trade T_W , the victim is estimated to receive $\delta_{y,W}$ Y -tokens, i.e., the number of tokens the victim would receive if no other trade is executed beforehand. On the other hand, when a sandwich attack occurs, the attacker front-runs the victim with transaction A_F exchanging $a_x^{\text{in}} > 0$ X -tokens for a_y Y -tokens. Now the victim's transaction executes at a worse price. To finish the attack, the attacker exchanges a_y Y -tokens for a_x^{out} X -tokens in the back-running attack transaction A_B .

We define the attacker's utility as their profit:

Definition 1. The attacker's utility $U^A(\delta_{x,W}, f, s, p, x, y)$ is given by

$$a_x^{\text{out}}(\delta_{x,W}, f, s, p, x, y) - a_x^{\text{in}}(\delta_{x,W}, f, s, p, x, y).$$

Here, $a_x^{\text{in}}(\delta_{x,W}, f, s, p, x, y)$ is the input of the front-running transaction and $a_x^{\text{out}}(\delta_{x,W}, f, s, p, x, y)$ is output of the back-running transaction.

We will assume that if a profitable sandwich attack exists, it executes successfully. A bot must have access to the necessary funds and achieve its desired transaction ordering which can be accomplished through *front-running-as-a-service* platforms such as Flashbots [7]. These services further guarantee their users that a transaction will only be included in a block if it executed successfully. Therefore, it is reasonable to assume that profitable sandwich attacks execute successfully.

4.2 Price Arbitrageurs

We consider a time window, during which the external market price between the pools' two cryptocurrencies is constant. Then, price arbitrageurs ensure that the pool's price returns to $P_{X \rightarrow Y}$ after every trade sequence (either an individual victim transaction in Pool_N or a victim transaction together with a sandwich attack in Pool_W). Thus, price arbitrageurs balance the market after any set of trades to reflect the fair market price. Letting the pool return to its initial state allows us to study the system analytically in the presence of an infinitely long trade flow as opposed to a fixed set of trades.

4.3 Traders

Our game captures a continuous stream of trade orders from two types of traders: sophisticated and retail. Sophisticated traders are aware and adjust to the presence of sandwich attacks, while retail traders are oblivious to the presence of these attacks, i.e., they trade in the pool as if there were no sandwich attacks. In the continuous stream of trade orders, a proportion $(1 - \omega)$ of orders originate from sophisticated traders and a proportion ω of orders from retail traders.

In Pool_N , where there are no sandwich attacks, the tokens received by traders equal the expected trade output. On the other hand, in Pool_W traders experience sandwich attacks which reduce the expected output. Sophisticated traders account for these attacks, while retail traders do not.⁵

All traders wish to sell X -tokens for Y -tokens, as they have a personal use for Y -tokens. Thus, the sophisticated trader's strategy space is

$$S^S = \{(\delta_{x,N}, \delta_{x,W}) \mid \delta_{x,N}, \delta_{x,W} \in \mathbb{R}^{\geq 0}\},$$

while the retail trader's strategy space is

$$S^T = \{(\Delta_{x,N}, \Delta_{x,W}) \mid \Delta_{x,N}, \Delta_{x,W} \in \mathbb{R}^{\geq 0}\},$$

where $\delta_{x,N}$ and $\Delta_{x,N}$ are the respective trade input sizes in Pool_N , whereas $\delta_{x,W}$ and $\Delta_{x,W}$ are the corresponding trade inputs in Pool_W .

Traders set their trade sizes across both pools to maximize their personal benefit. As the traders have a personal use for Y -tokens, they associate a relative benefit $\alpha > 0$ with Y -tokens. The private benefit associated with the number of Y -tokens a trader buys, $\delta_{y,\bullet}$, is thus given by $(1 + \alpha)\delta_{y,\bullet}$ in Pool_\bullet . In addition to the benefits traders obtain from the received Y -tokens, they associate a cost with the trade's input, which is given by $P_{X \rightarrow Y}\delta_{x,\bullet}$. Here, $\delta_{x,\bullet}$ is the trade input in X -tokens, and $P_{X \rightarrow Y}$ is the fair exchange rate from X -tokens to Y -tokens. Combining the trader benefit and cost in both

⁵Sophisticated traders assume there to be a sandwich attack for every transaction in Pool_W . As sandwich attacks only execute when they are profitable, there is not always a sandwich attack. However, this is only the case for small transactions (cf. Section 5.1) and unrealistic parameter configurations (cf. Section 6), and it is, therefore, negligible.

pools, we obtain their utility for the sophisticated traders in Definition 2 and for the retail trader in Definition 3. Notice that the important difference between the two utilities below is that the sophisticated trader takes the change in the output amount in Pool_W as a consequence of sandwich attacks into account, while the retail trader does not. We indicate this with the lack of the argument s in Definition 3. Further, note that the retail trader's utility U^R can be seen as the *expected* utility, i.e., what the retail trader expects and thus behaves according to. The *realized* utility of the retail trader is lower in the presence of attacks and equivalent to that of the sophisticated trader.

Definition 2. *The sophisticated trader's utility $U^S(\delta_{x,N}, \delta_{x,W}, \alpha, f, s, p, x, y)$ for a trade with input $\delta_{x,N} \geq 0$ in Pool_N and input $\delta_{x,W} \geq 0$ in Pool_W is given by*

$$(1 + \alpha)\delta_{y,N}(f, p, x, y) - \frac{y}{x}\delta_{x,N} \\ + (1 + \alpha)\delta_{y,W}(f, p, x, y, s) - \frac{y}{x}\delta_{x,W},$$

Here, $\delta_{y,N}(f, p, x, y)$ and $\delta_{y,W}(f, p, x, y, s)$ are the outputs of the trade in each pool.

Definition 3. *The retail trader's utility $U^R(\Delta_{x,N}, \Delta_{x,W}, \alpha, f, p, x, y)$ for a trade with input $\Delta_{x,N} \geq 0$ in Pool_N and input $\Delta_{x,W} \geq 0$ in Pool_W is given by*

$$(1 + \alpha)\Delta_{y,N}(f, p, x, y) - \frac{y}{x}\Delta_{x,N} \\ + (1 + \alpha)\Delta_{y,W}(f, p, x, y) - \frac{y}{x}\Delta_{x,W},$$

Here, $\Delta_{y,N}(f, p, x, y)$ and $\Delta_{y,W}(f, p, x, y)$ are the outputs of the trade in each pool.

Observe that in this framework, trades execute across both pools to maximize the respective trader's utility. Given a distribution on the relative benefit α and slippage tolerance s , the trading volume in either pool depends on the pool's reserve, transaction fee, and slippage tolerance. Throughout we assume all traders have the same relative benefit α and slippage tolerance s . Later we will also consider distribution on α to capture non-uniformity among traders. Further note that while we focus on trades from X to Y , the analysis applies directly in the opposite direction by symmetry.

4.4 Liquidity Providers

Liquidity providers supply reserves to the two pools. Knowledge of the sophisticated and retail trader's utility is assumed for liquidity providers. Further, liquidity providers are aware of the behavior of sandwich attackers and price arbitrageurs. We consider the liquidity providers to be rational, i.e., they optimally place their liquidity across the pools such that they maximize their received fees. The system has $n \in \mathbb{N}$ liquidity providers. Both the number of liquidity providers and the system's total reserves are fixed during the time of this analysis.

A liquidity provider LP_i for $i \in [0, \dots, n-1]$ holds a proportion l_i ($0 < l_i \leq 1$) of the total liquidity $L = \sqrt{x \cdot y}$, where x and y are the system's total reserves. We note that $\sum_{i=0}^{n-1} l_i = 1$.

Liquidity provider LP_i 's *strategy space* is given by all possible distributions of their liquidity across both pools:

$$S_i^{LP} = \{(p_i l_i L, (1-p_i) l_i L) | 0 \leq p_i \leq 1\}.$$

More precisely, a liquidity provider LP_i can choose the proportion p_i of their liquidity in Pool_N . They automatically place the remaining proportion $1-p_i$ of their liquidity in Pool_W . Knowing the distribution of the remaining liquidity $(1-l_i)L$ across the pools and the behavior of the other market participants, the liquidity provider chooses the strategy that maximizes the received fees. We define the liquidity provider's utility as the earned fees:

Definition 4. *The utility $U^{LP}(f, \alpha, s, x, y, p_i, l_i)$ of liquidity provider LP_i that places $p_i l_i L$ liquidity in Pool_N and $(1-p_i) l_i L$ in Pool_W represents the fees collected in both pools.*

Our game starts with an arbitrary initial liquidity distribution. One after the other, liquidity providers can change their personal liquidity distribution. The system is in a *Nash equilibrium* if no liquidity provider can improve their utility by unilaterally changing their liquidity distribution (strategy). We will loosen the restriction on equilibria and also consider ϵ -*equilibria*, where a liquidity provider only changes strategy if it increases their utility by a factor greater than $1 + \epsilon$ ($\epsilon \geq 0$). We analyze the system with this relaxation on equilibria, as inert liquidity providers are unlikely to change strategies for infinitesimal utility increases due to the effort involved. This adjustment allows us to analyze whether the potential private benefits of liquidity providers suffice.

5 Strategies

The optimal strategies of sandwich attackers and price arbitrageurs are straightforward. Sandwich attackers always execute the largest possible profitable attack, i.e., the attack inferring the maximal acceptable price movement on the trader (cf. Section 5.1), and price arbitrageurs re-balance the market after every trade sequence.

Sophisticated and retail traders set their trade sizes across both pools optimally to maximize their utility, knowing the pools' liquidity, transaction fee, and, in the prior case, the potential presence of sandwich attacks (cf. Section 5.2). Finally, liquidity providers distribute their liquidity to maximize the received fees. Liquidity providers account for the effects their decision to alter the liquidity distribution would have on the trading volume of the other market participants (cf. Section 5.3). Note that the section's omitted proofs can be found in Appendix B.

5.1 Sandwich Attack Profitability

A sandwich attacker only executes an attack whenever it is profitable, i.e., when U^A is positive (cf. Definition 1). We find that the sandwich attacker's profit for a front-running transaction of size a_x^{in} can be calculated analytically in Lemma 1. The expression is given in Appendix B.

Lemma 1. *The sandwich attacker's profit from an attack of size a_x^{in} to the front-running transaction on a victim's transaction $\delta_{x,W}$ can be given analytically.*

We will analyze the conditions under which profitable sandwich attacks exist. First, we determine a bound for the victim's trade size $\delta_{x,W}$ such that a profitable sandwich attack exists (cf. Lemma 2). From Lemma 2 we can follow that a profitable attack only exists, if the victim's trade size $\delta_{x,W}$ exceeds a fee dependent threshold

$$\delta_x^{\text{min}} = \frac{f(1-p)x}{(1-f)^2}.$$

Hence, only relatively large trades are prone to sandwich attacks.


Lemma 2. *A sandwich attack of size a_x^{in} is only profitable if the trader's transaction size exceeds*

$$\frac{f((1-p)x + a_x^{\text{in}}(1-f))}{(1-f)^2}.$$


Next, we explore what limits the attacker's maximum profit to show that it is optimal for sandwich attackers to execute the attack with maximal input size, i.e., the attack that infers the maximal acceptable price movement, as dictated by the slippage tolerance of the trader. In Lemma 3 we show that the attacker's maximal input a_x^s for which a victim's transaction still executes can be calculated analytically.

Lemma 3. *The sandwich attacker's maximal input, a_x^s , for a transaction exchanging $\delta_{x,W}$ X-tokens with slippage tolerance s such that the victim's trade still executes can be given analytically.*

However, for very large slippage tolerances the size of the sandwich attack is limited. To see this we can consider the asymptotic behavior of Lemma 1 in the limit of very large attack sizes a_x^{in} : $\lim_{a_x^{\text{in}} \rightarrow \infty} U^A = \lim_{a_x^{\text{in}} \rightarrow \infty} -f a_x^{\text{in}} \rightarrow -\infty$. We, thus, analyze whether the slippage tolerance or profitability limits the sandwich attack size and plot the sandwich attack size that achieves the maximum profit U^A as a function of the victim's transaction size $\delta_{x,W}$ in Figure 2a. Figure 2b shows the sandwich attacker's maximal input a_x^s as a function of the victim's transaction size $\delta_{x,W}$ for different slippage tolerances. Note the vast difference in scale, demonstrating that the sandwich attack size is clearly limited by the slippage tolerance. Thus, in realistic market configurations, the sandwich attackers always execute the attack with maximal possible input size a_x^s .



(a) Attack size a_x^{in} achieving max profit U^A vs. the victim's trade size. a_x^{in} is found by maximizing the attacker's profit w.r.t. a_x^{in} .



(b) Maximum sandwich attack size dependent on the victim's transaction size and slippage tolerance.

Figure 2: Limits on the sandwich attack size in terms of profitability (left) and slippage tolerance (right) for $f = 0.3\%$. Note the vast difference in scale of the vertical axis, demonstrating that the attack is limited by the slippage tolerance.

5.2 Trade Sizes

Sophisticated traders wish to maximize their utility $U^S(\delta_{x,N}, \delta_{x,W}, \alpha, f, s, p, x, y)$ (cf. Definition 2), i.e., the difference between the benefit from receiving Y -tokens and the trade's costs. The utility function accounts for sandwich attacks in Pool_W and assumes that the transaction output is reduced by the slippage tolerance. Retail traders, on the other hand, wish to maximize their utility $U^R(\Delta_{x,N}, \Delta_{x,W}, \alpha, f, p, x, y)$ (cf. Definition 3) that does not account for the presence of sandwich attacks.

Sophisticated Trader. We start with sophisticated traders and show in Lemma 4 that the optimal transaction size, maximizing utility U^R , in Pool_N ($\delta_{x,N}^{\text{opt}}$) and Pool_W ($\delta_{x,W}^{\text{opt}}$) can be expressed analytically. Observe that the transaction size is proportional to the pool's reserves of X -token. Further, we can see that the effects of slippage tolerance on the trade size are identical to those of the transaction fee. Thus, the combination of transaction fee f in Pool_W and the trader's slippage tolerance s is from the trader's perspective equivalent to a larger transaction fee equaling $f + s - f \cdot s$ in Pool_N . Therefore, the

transaction size in Pool_N is always larger than in Pool_W at the same liquidity level, and we follow that the trader's utility is maximized for $p = 1$. We also note that the optimal transaction size increases with α and decreases with the transaction fee f , as well as, where applicable, the slippage tolerance s .

Lemma 4. *A trade of size $\delta_{x,N}^{\text{opt}} = \max(0, p \cdot x(\sqrt{(1+\alpha)(1-f)} - 1)/(1-f))$ maximizes a sophisticated trader's utility U^R in Pool_N and in Pool_W the optimum is at $\delta_{x,W}^{\text{opt}} = \max(0, (1-p)x(\sqrt{(1+\alpha)(1-s)(1-f)} - 1)/(1-f))$.*

With the help of Lemma 4, we can obtain bounds for relative benefit α as a function of the slippage tolerance s , such that sophisticated traders benefit from trading in Pool_N and Pool_W . A trader executes a trade in Pool_N , as long as their α exceeds

$$\alpha > \alpha_N^{\text{min}} = \frac{f}{(1-f)}.$$

Notice that this bound only depends on the transaction fee f . In Pool_W , a sophisticated trader will only execute a trade if

$$\alpha > \alpha_W^{\text{min}} = \frac{f + s - s \cdot f}{((1-f)(1-s))}.$$

Retail Trader. In Lemma 5, we show the optimal trade sizes for the retail trader (i.e., those that maximize utility U^R) in Pool_N ($\Delta_{x,N}^{\text{opt}}$) and Pool_W ($\Delta_{x,W}^{\text{opt}}$) can also be expressed analytically. Note that the proof for Lemma 5 is analogous to that of Lemma 4. In particular, in Pool_N the behavior of the retail trader mirrors that of the sophisticated trader (i.e., $\Delta_{x,N}^{\text{opt}} = \delta_{x,N}^{\text{opt}}$). In Pool_W , on the other hand, the retail traders, who is oblivious to sandwich attacks, only adjust their trade size in response to differing levels of liquidity relative to Pool_N . Thus, the retail trader's trade size in Pool_N is equivalent to that in Pool_W at the same liquidity level, and we follow that the trader's utility is independent of p . Importantly, this is only due to the retail trader optimizing their *expected* utility, which ignores the attacks. The *realized* utility, which matches that of the sophisticated trader, is maximized when $p = 1$.

Lemma 5. *A trade of size $\Delta_{x,N}^{\text{opt}} = \max(0, p \cdot x(\sqrt{(1+\alpha)(1-f)} - 1)/(1-f))$ maximizes a retail trader's utility U^R in Pool_N and in Pool_W the optimum is at $\Delta_{x,W}^{\text{opt}} = \max(0, (1-p)x(\sqrt{(1+\alpha)(1-f)} - 1)/(1-f))$.*

Further note that a retail trader executes a trade in Pool_N and Pool_W , as long as their α exceeds

$$\alpha > \alpha_N^{\text{min}} = \frac{f}{(1-f)},$$

which is the same bound for the sophisticated trader in Pool_N .

5.3 Liquidity Distribution

A liquidity provider's utility directly corresponds to the received fees (cf. Definition 4). We will investigate how liquidity providers distribute their liquidity across the two pools', knowing that sophisticated and retail traders execute their respective optimal transactions. Recall, that in the continuous stream of trade orders that we study, a proportion $(1 - \omega)$ stems from sophisticated traders, while a proportion ω stems from retail traders. Each of these trades is accompanied by a sandwich attack (whenever applicable), and price arbitrage.

We quantify the system's total fees in Lemma 6 and that the total fees are proportional to p . If the fee gradient with respect to p is zero, all liquidity distributions maximize the game's fees. Otherwise, the game's fees are maximized, either when all liquidity is in Pool_N ($p = 1$) or when all liquidity is in Pool_W ($p = 0$).

Lemma 6. *The total transaction fees $F(f, \alpha, s, y, p, \omega)$ collected across both pools for retail and sophisticated traders with relative benefit α are proportional to p .*

Proof. We consider the fee revenue stemming from the order flow related to sophisticated and retail traders separately.

We start with sophisticated traders. There, we consider the following four intervals:

$$0 < \alpha \leq \alpha_N^{\min}, \alpha_N^{\min} < \alpha < \alpha_W^{\min},$$

$$\alpha > \alpha_W^{\min} \text{ and } U_S^A \leq 0, \alpha > \alpha_W^{\min} \text{ and } U_S^A \geq 0,$$

where $U_S^A(\delta_{x,W}, f, s, p, x, y)$ is the sandwich attacks' profitability for the trades from sophisticated traders (cf. Definition 1).

Following from Lemma 4, we conclude that no trades from sophisticated traders execute in either pool and, thereby, no fees collected for $\alpha \leq \alpha_N^{\min}$.

We continue with the second interval, i.e., $\alpha_N^{\min} < \alpha \leq \alpha_W^{\min}$. Following from Lemma 4 sophisticated traders exclusively execute trades in Pool_N on this interval. The fees collected in Pool_N for a transaction by a sophisticated trader with relative benefit α are

$$\begin{aligned} F_{N,S}(f, \alpha, s, y, p) &= f \left(\delta_{x,N}^{\text{opt}} \cdot \frac{y}{x} + \frac{p \cdot y(1-f)\delta_{x,N}^{\text{opt}}}{p \cdot x + (1-f)\delta_{x,N}^{\text{opt}}} \right) \\ &= p \cdot y \cdot f \left(\frac{\alpha}{\sqrt{(1+\alpha)(1-f)}} - \frac{f}{1-f} \right) \end{aligned}$$

Y -tokens. In the previous, $\delta_{x,N}^{\text{opt}} \cdot \frac{y}{x}$ is the sophisticated trader's transaction size in Y -tokens in Pool_N and

$$\frac{p \cdot y \delta_{x,N}^{\text{opt}}}{p \cdot x + (1-f)\delta_{x,N}^{\text{opt}}}$$

is the size of the price arbitrageur's transaction.

In the third interval, i.e., $\alpha > \alpha_W^{\min}$ and $U_S^A \leq 0$, sophisticated traders execute trades in both pools (cf. Lemma 4). However, there is no profitable sandwich attack, due to the

small trade size in Pool_W (cf. Lemma 2). The fees collected from the sophisticated trader order flow in Pool_N are again given by $F_{N,S}(f, \alpha, s, y, p)$ but liquidity providers collect additional fees in Pool_W. The fees collected in Pool_W for a transaction by a sophisticated trader with relative benefit α are given by

$$\begin{aligned} F_{W,S}^{U_S^A \leq 0}(f, \alpha, s, y, p) &= f \left(\delta_{x,W}^{\text{opt}} \cdot \frac{y}{x} + \frac{(1-p)y(1-f)\delta_{x,W}^{\text{opt}}}{(1-p)x + (1-f)\delta_{x,W}^{\text{opt}}} \right) \\ &= (1-p)y \cdot f \left(\frac{(n_1(f, \alpha, s) - 1)(1-f + n_1(f, \alpha, s))}{(1-f)n_1(f, \alpha, s)} \right) \end{aligned}$$

where

$$\delta_{x,W}^{\text{opt}} \cdot \frac{y}{x}$$

is the trader's transaction size in Y -tokens in Pool_W and

$$\frac{p \cdot y \delta_{x,W}^{\text{opt}}}{p \cdot x + (1-f)\delta_{x,W}^{\text{opt}}}$$

is the size of the price arbitrageur's transaction that returns the pools to its initial state. Further

$$n_1(f, \alpha, s) = \sqrt{(1+\alpha)(1-s)(1-f)}.$$

Finally, we analyze the fourth interval, i.e., $\alpha > \alpha_W^{\min}$ and $U_S^A > 0$. In this interval, sophisticated trades execute in both pools and sandwich attacks execute in Pool_W. Thus, in addition to the fees $F_{N,S}(f, \alpha, s, y, p)$ collected in Pool_N, we also consider the fees collected in Pool_W from traders, price arbitrageurs, and sandwich attackers for the liquidity provider utility. In the presence of sandwich attacks, the fees from sophisticated flow in Pool_W are given by

$$\begin{aligned} F_{W,S}^{U_S^A > 0}(f, \alpha, s, y, p) &= \left(\left(\delta_{x,W}^{\text{opt}} + \alpha_x^s \right) \frac{y}{x} + \frac{(1-p)y(1-f)(\delta_{x,W}^{\text{opt}} + \alpha_x^s)}{(1-p)x + (1-f)(\delta_{x,W}^{\text{opt}} + \alpha_x^s)} \right) \\ &= \frac{(1-p)y \cdot f((n_1(f, \alpha, s) - 3) + n_2(f, \alpha, s))(n_1(f, \alpha, s) + 1 - 2f) + n_2(f, \alpha, s)}{1 - (1-f)((n_1(f, \alpha, s) - 1) + n_2(f, \alpha, s))} \end{aligned}$$

where $\left(\delta_{x,W}^{\text{opt}} + \alpha_x^s \right) \frac{y}{x}$ combines the trader's transaction size in Pool_W and the bot's front-running transaction size in Y -tokens (cf. Lemma 3). In the previous,

$$n_2(f, \alpha, s) = \sqrt{\frac{2n_1(f, \alpha, s)(1+s) + (1-s)(2+\alpha(1-s)-s) - (1+\alpha)f(1-s)^2}{1-s}}.$$

Further,

$$\frac{(1-p)y(1-f)(\delta_{x,W}^{\text{opt}} + \alpha_x^s)}{(1-p)x + (1-f)(\delta_{x,W}^{\text{opt}} + \alpha_x^s)}$$

is the combined size of the attacker's back-running transaction and the transaction that returns the pool to its initial state.

Next, we consider the fee revenue from the order flow associated with retail traders. Here, we consider three intervals

$$0 < \alpha \leq \alpha_N^{\min}, \alpha > \alpha_N^{\min} \text{ and } U_R^A \leq 0, \alpha > \alpha_N^{\min} \text{ and } U_R^A \geq 0,$$

where $U_R^A(\Delta_{x,W}, f, s, p, x, y)$ is the sandwich attacks' profitability for the trades from retail traders (cf. Definition 1).

Following from Lemma 5, we conclude that no trades from retail traders execute in either pool and, thereby, no fees collected for $\alpha \leq \alpha_N^{\min}$.

We continue with the second interval, i.e., $\alpha \geq \alpha_N^{\min}$ and $U_R^A \leq 0$. In contrast, to sophisticated traders, for $\alpha \geq \alpha_N^{\min}$, retail traders trade in both pools (cf. Lemma 5). Further, in this interval, there is no profitable sandwich attack on the retail flow, due to the small trade size in Pool_W (cf. Lemma 2). The fees collected from the retail trader order flow in Pool_N are given by $F_{N,R}(f, \alpha, s, y, p)$ which is equal to the fees collected in Pool_N by sophisticated traders (cf. Lemma 5), i.e.,

$$F_{N,R}(f, \alpha, s, y, p) = p \cdot y \cdot f \left(\frac{\alpha}{\sqrt{(1+\alpha)(1-f)}} - \frac{f}{1-f} \right) =: F_N.$$

Further, liquidity providers collect additional fees in Pool_W. The fees collected in Pool_W for a transaction by a retail trader with relative benefit α are given by

$$\begin{aligned} & F_{W,R}^{U_R^A \leq 0}(f, \alpha, s, y, p) \\ &= f \left(\Delta_{x,W}^{\text{opt}} \cdot \frac{y}{x} + \frac{(1-p)y(1-f)\Delta_{x,W}^{\text{opt}}}{(1-p)x + (1-f)\Delta_{x,W}^{\text{opt}}} \right) \\ &= (1-p)y \cdot f \left(\frac{f + f\alpha - \alpha\sqrt{(1-f)(1-\alpha)}}{(1-f)(1-\alpha)} \right). \end{aligned}$$

Finally, we analyze the third interval, i.e., $\alpha > \alpha_N^{\min}$ and $U_r^A > 0$, where retail traders trade in both pools and sandwich attacks are profitable in Pool_W. Here, in addition to the fees $F_{N,R}(f, \alpha, s, y, p)$ collected in Pool_N, we also consider the fees collected in Pool_W from traders, price arbitrageurs, and sandwich attackers for the liquidity provider utility. In the presence of sandwich attacks, the fees from retail flow in Pool_W are given by

$$\begin{aligned} & F_{W,R}^{U_R^A > 0}(f, \alpha, s, y, p) \\ &= \left(\left(\Delta_{x,W}^{\text{opt}} + a_x^s \right) \frac{y}{x} + \frac{(1-p)y(1-f)(\Delta_{x,W}^{\text{opt}} + a_x^s)}{(1-p)x + (1-f)(\Delta_{x,W}^{\text{opt}} + a_x^s)} \right) \\ &= (1-p)y \cdot f \left(\frac{n_3(f, \alpha, s)}{1-f} + \frac{n_3(f, \alpha, s)}{1+n_3(f, \alpha, s)} \right) \end{aligned}$$

where

$$n_3(f, \alpha, s) = \frac{\sqrt{(1+\alpha)(1-f)} - 3 + \sqrt{\left(\sqrt{(1+\alpha)(1-f)} - 1 \right)^2 + \frac{4\sqrt{(1+\alpha)(1-f)}}{1-s}}}{2}.$$

Through a combination, we obtain that the fees collected

across both pools are given by

$$F(f, \alpha, s, y, p, \omega) = \begin{cases} 0 & 0 < \alpha \leq \alpha_N^{\min} \\ F_N + \omega \cdot F_{W,R}^{U_R^A \leq 0} & \alpha_N^{\min} < \alpha \leq \alpha_W^{\min} \text{ \& } U_R^A \leq 0 \\ F_N + \omega \cdot F_{W,R}^{U_R^A > 0} & \alpha_N^{\min} < \alpha \leq \alpha_W^{\min} \text{ \& } U_R^A > 0 \\ F_N + (1-\omega)F_W^{U^A \leq 0} & \alpha \geq \alpha_W^{\min} \text{ \& } U_R^A, U_S^A \leq 0 \\ + \omega \cdot F_{W,R}^{U_R^A \leq 0} & \\ F_N + (1-\omega)F_W^{U^A \leq 0} & \alpha \geq \alpha_W^{\min}, U_R^A > 0 \text{ \& } U_S^A \leq 0 \\ + \omega \cdot F_{W,R}^{U_R^A > 0} & \\ F_N + (1-\omega)F_W^{U^A > 0} & \alpha \geq \alpha_W^{\min} \text{ \& } U_R^A, U_S^A > 0 \\ + \omega \cdot F_{W,R}^{U_R^A > 0} & \end{cases}$$


$F(f, \alpha, s, y, p, \omega)$ linearly combines the fees from the two streams (sophisticated and retail) according to their relative proportions (i.e., sophisticated trades make up a proportion $1 - \omega$ and retail traders the rest). We conclude that F is proportional to p for every $\alpha > 0$. \square

Importantly, Lemma 6 holds an infinite sequence of trade orders from sophisticated and retail traders with the same (α, s) along with the associated orders from sandwich attackers and arbitrageurs. The previous follows from price arbitrageurs returning the pool to its initial price $P_{X \rightarrow Y}$ after every trade sequence. We conclude that the total fees collected for a continuous stream of trade orders originating from a homogeneous set of traders with the same relative benefit α and slippage tolerance s is proportional to p , i.e., the fraction of the total liquidity placed in Pool_N.


Lemma 6 gives the system's total fees $F(f, \alpha, s, y, p, \omega)$. By virtue of the proportionality of the total fees F in both p and y , the fees received by an individual liquidity provider LP_i with liquidity $(p_i l_i L, (1-p_i) l_i L)$ are given by $F_i(f, \alpha, s, y, p_i, l_i, \omega) = l_i \cdot F(f, \alpha, s, y, p_i, \omega)$. Therefore, the optimal liquidity distribution for an individual liquidity provider also has all liquidity in Pool_N ($p_i = 1$) or all liquidity in Pool_W ($p_i = 0$), whenever the gradient of the fee with respect to p_i is non-zero. A liquidity provider will redistribute their liquidity optimally, i.e., such that their utility is maximized, whenever they can increase their received fees by more than a factor of $1 + \varepsilon$.

6 Game Equilibria

Before discussing the quantitative model, we will give an intuitive explanation. First, we note that liquidity providers profit from large trading volumes, irrespective of their origin. As the sandwich attackers extract their profits from the traders, sophisticated traders will reduce their trading volume if the attacks become too lucrative but retail trades will not. Therefore, whether a pool with sandwich attacks is the equilibrium



(a) We set $\omega = 0.01$.



(b) We set $\omega = 0.1$.

Figure 3: The Nash equilibrium (color shading), is dependent on the slippage tolerance and the relative benefit for $\omega = 0.01$ (cf. Figure 3a) and $\omega = 0.1$ (cf. Figure 3b). We set $x = 5,000,000 X$, $y = 5,000,000 Y$ and $f = 0.003$.

boils down to whether the increased trading volume the attackers generate can offset the diminished trading activity of sophisticated traders.

We will now substantiate this qualitative picture by locating the game's ε -equilibria to identify which pool is favored by the market actors. A liquidity distribution is an ε -equilibria if no liquidity provider can increase their utility by more than a factor of $1 + \varepsilon$ by adjusting the liquidity distribution. For $\varepsilon = 0$, any ε -equilibrium is considered a Nash equilibrium. We analyze the game's equilibria assuming a fixed (mean) benefit for the traders, α , in Section 6.1 and discuss the heterogeneous case in Section 6.2. Further note that the section's omitted proofs can be found in Appendix C.

6.1 Homogeneous Traders

We start by analyzing the game's equilibria given a homogeneous trader set, i.e., all traders have the same relative benefit α . In the simplest case, $\partial_p F = 0$, all liquidity distributions are both ε -equilibria and Nash equilibria (cf. Lemma 7).

Lemma 7. *The only Nash equilibria if $\partial_p F \neq 0$ are $p \in \{0, 1\}$. If $\partial_p F = 0$, all liquidity distributions are ε -equilibria in a homogeneous traders game.*

In Lemma 7 we further show that for $\partial_p F \neq 0$ the only possible Nash equilibria are the two corner cases: all liquidity in Pool_N ($p = 1$) or in Pool_W ($p = 0$). This follows from the proportionality of the fees to p which means that the sign of $\partial_p F$ dictates the location of the Nash equilibrium. In Figure 3, we plot the dependence of this equilibrium on the slippage tolerance and relative benefit for $\omega = 0.01$ (i.e., 1% of the trade flow originates from retail traders) and $\omega = 0.1$ (i.e., 10% of the trade flow originates from retail traders).

We first consider the setting when $\omega = 0.01$ (cf. Figure 3a). Notice that in areas where either the trader's relative benefit α or the slippage tolerance s is high, Pool_N is the Nash equilibrium. When α is comparatively large, so is the traders' transaction size. Liquidity providers, therefore, receive a substantial amount of fees from sophisticated traders and sandwich attacks would decrease the pool's trading volume more than the volume created by the attacker. Thus, all liquidity is in Pool_N . The same holds when the slippage tolerance is high compared to the trader's benefit. Sophisticated traders no longer trade in Pool_W (cf. Lemma 4) or their size in Pool_W is too small for there to be a profitable sandwich attack (cf. Lemma 2). Thus, the only volume in Pool_W stems from retail traders (cf. Lemma 5).

There is a small area in between where Pool_W is the equilibrium. Here, the slippage tolerance is just small enough not to exceed the bound given in Lemma 4 and the sophisticated trader's transaction size in Pool_W is just large enough to allow for a profitable attack (cf. Lemma 2). Further, the trades from retail traders also allow for profitable attacks as their trade size exceeds that of sophisticated traders. Thus, liquidity providers' private incentives are maximized in the presence of sandwich attackers.

When $\omega = 0.1$ the picture shifts and Pool_W is the Nash equilibrium in a larger proportion for the parameter space due to the increase in volume from retail traders (cf. Figure 3b). For relatively large α , Pool_W dominates as the Nash equilibrium. In this part of our parameter space, the trades from retail and sophisticated traders suffer from sandwich attacks. Thus, for relatively large α the extra volume associated with retail traders, in comparison to the setting in Figure 3a, prevents Pool_N from being the Nash equilibrium.

In the part of the parameter space where the slippage tolerance is large in comparison to the relative benefit, we observe that in a sliver of the space Pool_W is the Nash equilibrium. Here, the sandwich attack on the trades from retail traders are profitable. Further, the sandwich attack volume is relatively large in comparison to the volume from sophisticated and retail traders in Pool_N as the slippage tolerance is large in comparison to the relative benefit. However, when this difference decreases, Pool_N becomes the Nash equilibrium as the sandwich attack volume can no longer compensate the loss in volume from sophisticated traders.

As ω grows the parts of the parameter space where Pool_W is the Nash equilibrium move in on each other. In particular,

for $\omega = 1$ (i.e., there are only retail traders) Pool_W is the Nash equilibrium on the entire parameter space and the opposite holds for $\omega = 0$ (i.e., there are no retail traders). Thus, the composition of order flow from the two types of traders is a major factor determining where the Nash equilibrium lies.

While the sign of the gradient $\partial_p F$ dictates the location of the Nash equilibrium, it is not sufficient to determine if it is an ε -equilibrium. As we show in Theorem 1, it is the relative difference between the fees the liquidity provider earns with their current distribution and the maximum fees they can collect that dictate whether the liquidity provider will change strategy.

Theorem 1. *A liquidity distribution is an ε -equilibrium if there is no liquidity provider with initial liquidity distribution $(p_i l_i L, (1 - p_i) l_i L)$, such that*

$$\frac{\max\{F(f, \alpha, s, y, 0, \omega), F(f, \alpha, s, y, 1, \omega)\}}{F(f, \alpha, s, y, p_i, \omega)} - 1 < \varepsilon.$$


Currently, all liquidity is in markets that allow for sandwich attacks. Therefore, even if the Nash equilibrium is in Pool_N , liquidity providers would have to move their liquidity.

Thus, we also identify market configurations that are ε -equilibria independent of the current liquidity distribution. In this situation, a new market with a front-running protection mechanism would not attract any liquidity even if it were to maximize the liquidity provider's private incentives. The maximum relative change in fees for a given market configuration is given as $|\Delta_F|$, where $\Delta_F = \partial_p F / F_{\min}$ and $F_{\min} = \min(F(p = 0), F(p = 1))$. Independent of the liquidity provider's initial distribution, the relative benefit of switching strategy cannot exceed ε in case $|\Delta_F| < \varepsilon$. The sign of Δ_F corresponds to the sign of the fee's gradient $\partial_p F$ and therefore indicates the position of the Nash equilibrium.


We simulate the dependence of Δ_F on the slippage tolerance and relative benefit in Figure 4 for $\omega = 0.01$ and $\omega = 0.1$. Starting with the setting where $\omega = 0.01$ (cf. Figure 4a). For comparatively large slippage tolerances, the magnitude of Δ_F is large. Independent of the liquidity in Pool_W , the trading volume from sophisticated traders is either zero when $\alpha < \alpha_W^{\min}$ or relatively small when sandwich attacks are not profitable. Therefore, switching strategies by moving liquidity from Pool_W to Pool_N leads to a sizable increase in fees in this part of the parameter space, and we do not expect any ε -equilibria in Pool_W for this parameter range.

Turning to more realistic areas of the parameter space where the relative benefit is larger than the slippage tolerance, we notice that Δ_F 's magnitude is small (i.e., $\Delta_F < 0.02$). Thus, all liquidity providers who only change strategies for a relative benefit larger than 2% would not be inclined to move their liquidity. We follow that any liquidity distribution is an ε -equilibrium for a significant proportion of the parameter space even for small ε .

For $\omega = 0.1$ we observe a different picture (cf. Figure 4b).



(a) We set $\omega = 0.01$.



(b) We set $\omega = 0.1$.

Figure 4: Simulation of Δ_F across both pools depending on the trader's relative benefit and the slippage tolerance for $\omega = 0.01$ (cf. Figure 4a) and $\omega = 0.1$ (cf. Figure 4b). In blue areas, the Nash equilibrium is Pool_N , in red areas, it is Pool_W . Δ_F is cut off for better visibility and notice that the cutoff is different in the two plots. We set $x = 5,000,000 X$, $y = 5,000,000 Y$ and $f = 0.003$.

Recall, that in general for a higher ω a larger proportion of the parameter space has Pool_W (red areas) as the Nash equilibrium. In addition to that, we observe that in general Δ_F 's magnitude is larger. Only when the relative benefit is very large in comparison to the slippage tolerance is the relative benefit small, i.e., below 2%.

Therefore, when $\omega = 0.01$ liquidity providers are largely indifferent to whether the market utilizes a front-running protection mechanism, and might require additional financial incentives to migrate their liquidity to pools with front-running protection mechanisms. However, when $\omega = 0.1$ the difference between the pools are more pronounced and liquidity providers are more likely to move their liquidity to the pool with higher fee revenue. This, however, is Pool_W in the most realistic areas of the parameter space.

6.2 Heterogeneous Traders

We continue with analyzing where the ε -equilibria fall in a game with heterogeneous traders. We model a trader's rela-

tive benefit α as a random variable A with probability mass function $\psi_A(\alpha)$. The game is in an ε -equilibrium for any probability mass function $\psi_A(\alpha)$ that fulfills the condition provided in Theorem 2. Theorem 2 assumes that the random variable A is discrete. Note, however, that it could be adapted to the continuous case.

Theorem 2. *A liquidity distribution in a system with heterogeneous traders with distribution $\psi_A(\alpha)$ is an ε -equilibrium if there is no liquidity provider with initial liquidity distribution $(p_i)_i L, (1 - p_i)_i L$, such that*

$$\frac{\max \{ \sum_{\alpha} \psi_A(\alpha) F(f, \alpha, s, y, 0, \omega), \sum_{\alpha} \psi_A(\alpha) F(f, \alpha, s, y, 1, \omega) \}}{\sum_{\alpha} \psi_A(\alpha) F(f, \alpha, s, y, p_i, \omega)} - 1 < \varepsilon.$$

Further, we predict that for most probability distributions, extreme values of the system’s fee gradient $\partial_p F$ will be averaged out. To back up this assumption, we simulate the Δ_F for a two-point distribution in Appendix A.

7 Social Incentives and Self-Regulation

In our model, whenever the derivative of the liquidity provider’s utility is positive, i.e., $\partial_p U^{LP} = \partial_p F > 0$, the system’s Nash equilibrium maximizes social welfare.

Our analysis demonstrates that without or with very few retail traders, i.e., traders that act irrational, Pool_N is the Nash equilibrium for the vast majority of the parameter space (cf. Figure 3a). Therefore, markets preventing front-running generally align the private incentives of liquidity providers with the system’s incentives. However, liquidity providers are currently in markets without front-running protections. Thus, an innovative DEX preventing front-running attacks must attract liquidity from other markets. Our analysis highlights that even when placing all liquidity in Pool_N maximizes the private incentives of liquidity providers, the benefit from adjusting a liquidity distribution is often only small. The market, therefore, cannot rely on the inert liquidity providers to revise their liquidity distribution. Therefore, added incentives may be required for the successful adoption of a such market.

Additionally, when there is a significant proportion of retail traders Pool_W is the Nash equilibrium more frequently (cf. Figure 3a). This is a consequence of “irrational” behavior from retail traders who do not respond to the presence of sandwich attacks – possibly due to an information asymmetry.

Thus, for the market to self-regulate it must (1) attract liquidity to novel DEXes, and (2) educate traders.

Attracting Liquidity Providers. In the following, we discuss the possibilities of how a new DEX, implementing a front-running prevention scheme, could attract liquidity providers. One possibility would be for the DEX to distribute its native token to liquidity providers as an added incentive. Similar benefits have been distributed at the launch of new DeFi platforms

(cf. Sushiswap [13]). Note here that it would be important for these distributed tokens to cover at least the gas fees required to migrate the liquidity, as otherwise, it is unlikely that it would motivate many. Another possibility would be to directly cover the migration costs. Such offers are utilized by brokerages in traditional finance (cf. Ally [1]) and could also be implemented by DEXes. Finally, once they reach a certain traction we would expect many liquidity providers to follow.

Educating Traders. In our model, the trade volume retail traders direct to the respective pools is proportional to the liquidity available in that pool, thus they completely ignore the effects of sandwich attacks. Further, this behavior is a major driver in the Nash equilibria lying in pools with sandwich attacks. The behavior of retail traders, who are likely unaware of the presence and consequences of the attacks, can be altered through education. Retail traders adjusting their response to sandwich attacks is crucial for the adoption of a prevention mechanism.

8 Conclusion

Our game-theoretical study of the incentives of traders and liquidity providers to adopt a DEX with a new market design preventing front-running attacks shows that when the vast majority of traders ($\approx 99\%$) are sophisticated, the private incentives of both traders and liquidity providers generally align the market’s social incentives — eliminating front-running attacks. However, this drastically shifts when the proportion of retail traders increases. Even when retail traders only account for 10% of the order flow, the private incentives of liquidity providers oppose the market’s social incentives, i.e., liquidity providers are generally drawn to pools with sandwich attacks.

This finding highlights the struggles of eliminating front-running attacks. Even though, the private incentives of sophisticated traders and liquidity providers align (i.e., without retail traders the Nash equilibrium is always in pools without sandwich attacks), a small proportion of traders who act irrationally, completely changes the picture.

Further, even if the Nash equilibrium is in pools without sandwich attacks there is a further challenge. In the absence of a central authority, market participants must experience a personal benefit for the successful adoption of such a design. The alignment of the liquidity provider’s private incentives and the market’s social incentives is promising. Yet, our analysis also finds that the increase in the liquidity provider’s utility from moving to a market preventing front-running is generally small. Liquidity providers are usually not nimble market participants. Therefore, the prospect of a small utility increase might not suffice.

Successful self-regulation of the market to prevent front-running attacks is likely to not only require intensive education of traders but also additional initial financial incentives to gain the attention of liquidity providers.

References

- [1] Ally. <https://www.ally.com/help/invest/transfers/>, 2024.
- [2] Automata network. <https://www.ata.network/>, 2024.
- [3] bloxroute. <https://bloxroute.com/private-transactions/>, 2024.
- [4] Cowswap. <https://cowswap.exchange/>, 2024.
- [5] DEXs volume. <https://defillama.com/dexs>, 2024.
- [6] Eden. <https://www.edennetwork.io/>, 2024.
- [7] flashbots. <https://docs.flashbots.net/>, 2024.
- [8] Gnosis protocol. <https://gnosis.io/>, 2024.
- [9] Openmev. <https://openmev.xyz/>, 2024.
- [10] Proposer-builder separation, 2024.
- [11] Sandwich overview. <https://eigenphi.io/mev/ethereum/sandwich>, 2024.
- [12] Secretswap. <https://secretswap.net/>, 2024.
- [13] Sushiswap. <https://sushi.com/>, 2024.
- [14] Hayden Adams, Noah Zinsmeister, and Dan Robinson. Uniswap v2 core. 2020.
- [15] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keeper, and Dan Robinson. Uniswap v3 core. Technical report, Uniswap, 2021.
- [16] Avi Asayag, Gad Cohen, Ido Grayevsky, Maya Leshkowitz, Ori Rottenstreich, Ronen Tamari, and David Yakira. A fair consensus protocol for transaction ordering. In *2018 IEEE 26th International Conference on Network Protocols (ICNP)*, pages 55–65, 2018.
- [17] Leemon Baird. The swirls hashgraph consensus algorithm: Fair, fast, byzantine fault tolerance. *Swirls Tech Reports SWIRLDS-TR-2016-01, Tech. Rep.*, 2016.
- [18] Ido Bentov, Yan Ji, Fan Zhang, Lorenz Breidenbach, Philip Daian, and Ari Juels. Tesseract: Real-time cryptocurrency exchange using trusted hardware. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS '19*, page 1521–1538, New York, NY, USA, 2019. Association for Computing Machinery.
- [19] Dan Bernhardt and Bart Taub. Front-running dynamics. *Journal of Economic Theory*, 138(1):288–296, 2008.
- [20] Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. Enter the hydra: Towards principled bug bounties and exploit-resistant smart contracts. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1335–1352, 2018.
- [21] Eric Budish, Robin S. Lee, and John J. Shim. A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market? NBER Working Papers 25855, National Bureau of Economic Research, Inc, 2019.
- [22] Christian Cachin, Jovana Mičić, and Nathalie Steinhauer. Quick order fairness. In *Financial Cryptography and Data Security (FC), Grenada*, 2022.
- [23] Carole Comerton-Forde and Kar Mei Tang. Anonymity, frontrunning and market integrity. *The Journal of Trading*, 2(4):101–118, 2007.
- [24] Andrei Constantinescu, Diana Ghinea, Lioba Heimbach, Zilin Wang, and Roger Wattenhofer. A fair and resilient decentralized clock network for transaction ordering. In *27th International Conference on Principles of Distributed Systems (OPODIS), Tokyo, Japan*, December 2023.
- [25] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Ido Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.
- [26] Jean-Pierre Danthine and Serge Moresi. Front-Running by Mutual Fund Managers: A Mixed Bag. *Review of Finance*, 2(1):29–56, 1998.
- [27] Yael Doweck and Ittay Eyal. Multi-party timed commitments. *arXiv preprint arXiv:2005.04883*, 2020.
- [28] Shayan Eskandari, Mahsa Moosavi, and Jeremy Clark. Sok: Transparent dishonesty: front-running attacks on blockchain. In *Financial Cryptography and Data Security (FC), St. Kitts, Saint Kitts and Nevis*, February 2019.
- [29] Lioba Heimbach and Roger Wattenhofer. Eliminating sandwich attacks with the help of game theory. In *ACM Asia Conference on Computer and Communications Security (ASIA CCS), Nagasaki, Japan*, June 2022.
- [30] Lioba Heimbach and Roger Wattenhofer. SoK: Preventing Transaction Reordering Manipulations in Decentralized Finance. In *4th ACM Conference on Advances in Financial Technologies (AFT), Cambridge, Massachusetts, USA*, September 2022.

- [31] Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. Order-fair consensus in the permissionless setting. *IACR Cryptol. ePrint Arch.*, 2021:139, 2021.
- [32] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *Cryptology ePrint Archive*, Report 2021/1465, 2021. <https://ia.cr/2021/1465>.
- [33] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In *Annual International Cryptology Conference*, pages 451–480. Springer, 2020.
- [34] Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 25–36, 2020.
- [35] Viktor Manahov. Front-running scalping strategies and market manipulation: Why does high-frequency trading need stricter regulation? *Financial Review*, 51(3):363–402, 2016.
- [36] Jerry W. Markham. Front-running - insider trading under the commodity exchange act. *Catholic University Law Review*, 38:69, 1988-1989.
- [37] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. CCS '16, page 31–42, New York, NY, USA, 2016. Association for Computing Machinery.
- [38] Peyman Momeni, Sergey Gorbunov, and Bohan Zhang. Fairblock: Preventing blockchain front-running with minimal overheads. In *Security and Privacy in Communication Networks: 18th EAI International Conference, SecureComm 2022, Virtual Event, October 2022, Proceedings*, pages 250–271. Springer, 2023.
- [39] Imad Moosa. The regulation of high-frequency trading: A pragmatic view. *Journal of Banking Regulation*, 16(1):72–88, 2015.
- [40] Ariel Orda and Ori Rottenstreich. Enforcing fairness in blockchain transaction ordering. *Peer-to-peer Networking and Applications*, 14(6):3660–3673, 2021.
- [41] Andreas Park. The conceptual flaws of constant product automated market making, 2021. Available at SSRN: 3805750.
- [42] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214. IEEE, 2022.
- [43] Michael K Reiter and Kenneth P Birman. How to securely replicate services. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 16(3):986–1009, 1994.
- [44] Chrysoula Stathakopoulou, Signe Rüsçh, Marcus Brandenburger, and Marko Vukolić. Adding fairness to order: Preventing front-running attacks in bft protocols using tees. In *2021 40th International Symposium on Reliable Distributed Systems (SRDS)*, pages 34–45. IEEE, 2021.
- [45] Ana Tatabitovska, Oğuzhan Ersoy, and Zekiraya Erkin. Mitigation of transaction manipulation attacks in uniswap. 2021.
- [46] Ye Wang, Patrick Züst, Yaxing Yao, Zhicong Lu, and Roger Wattenhofer. Impact and User Perception of Sandwich Attacks in the DeFi Ecosystem. In *ACM CHI Conference on Human Factors in Computing Systems (CHI)*, New Orleans, LA, USA, May 2022.
- [47] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. 2014.
- [48] Haoqian Zhang, Louis-Henri Merino, Vero Estrada-Galinanes, and Bryan Ford. Flash freezing flash boys: Countering blockchain front-running. In *2022 IEEE 42nd International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pages 90–95. IEEE, 2022.
- [49] Yunhao Zhang, Srinath Setty, Qi Chen, Lidong Zhou, and Lorenzo Alvisi. Byzantine ordered consensus without byzantine oligarchy. In *14th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 20)*, pages 633–649, 2020.
- [50] Liyi Zhou, Kaihua Qin, and Arthur Gervais. A2mm: Mitigating frontrunning, transaction reordering and consensus instability in decentralized exchanges, 2021.
- [51] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges, 2020.

A Game Equilibria – Heterogeneous Traders

We simulate the Δ_F for a two-point distribution with the following probability mass function:

$$\Psi_A(\alpha) = \begin{cases} \frac{1}{2} & \text{if } \alpha = \alpha_k^- = (1 - \frac{1}{k})\mu\alpha, \\ \frac{1}{2} & \text{if } \alpha = \alpha_k^+ = (1 + \frac{1}{k})\mu\alpha, \end{cases}$$

in Figure 5 for $k = 10$ (cf. Figure 5a) and $k = 3$ (cf. Figure 5b). Note that for these simulations we set $\omega = 0.01$, so retail traders account for only 1% of the order flow. As expected Δ_F resembles the homogeneous case more closely, when the two points of the distribution are close to each other, i.e., for the higher values of k . Note that for $k = \infty$ the two-point distribution becomes a one-point distribution, i.e., the homogeneous case.

We further notice that a more significant area of the parameter space has Pool_N as the Nash equilibrium when the slippage tolerance is large. Half the traders have a lower relative benefit than $\mu\alpha$ and, thereby, these sophisticated traders will only execute transactions in Pool_W for smaller slippage tolerances. We further note that the area of the parameter space, where Pool_W is the Nash equilibrium grows smaller as k decreases. While there remains a small area of the parameter space that has Pool_W as the Nash equilibrium for $k = 10$ (cf. Figure 5a), this is noticeable smaller for $k = 3$ (cf. Figure 5b).

The particular combination of requirements that must be met for Pool_W to be the Nash equilibrium when ω is very small is achieved less frequently as the distance between the two points of the distribution grows.

B Omitted Strategy Proofs

B.1 Sandwich Attack Profitability

Lemma 1. *The sandwich attacker's profit from an attack of size a_x^{in} to the front-running transaction on a victim's transaction $\delta_{x,W}$ can be given analytically.*

Proof. First, the sandwich attacker swaps a_x^{in} and receives


$$a_y = - \int_{(1-p)x}^{(1-p)x+(1-f)a_x^{\text{in}}} \frac{-x \cdot y}{\xi^2} d\xi,$$

in the front-running transaction A_F . Then the trader sells $\delta_{x,W}$ and in return receives


$$\tilde{\delta}_{y,W} = - \int_{(1-p)x+(1-f)a_x^{\text{in}}}^{(1-p)x+(1-f)(a_x^{\text{in}}+\delta_{x,W})} \frac{-x \cdot y}{\xi^2} d\xi.$$

Finally, the sandwich attacker uses a_y Y -tokens to buy a_x^{out} X -tokens in its back-running transaction A^B . Due to the transaction fee f being applied to the input, only $\tilde{a}_y = (1-f)a_y$ of the initially swapped a_y re-enters the pool. Therefore, we write

$$\tilde{a}_y = \int_{(1-p)x+(1-f)(a_x^{\text{in}}+\delta_{x,W})}^{(1-p)x+(1-f)(a_x^{\text{in}}+\delta_{x,W})-a_x^{\text{out}}} \frac{-x \cdot y}{\xi^2} d\xi,$$



(a) We set $k = 10$.



(b) We set $k = 3$.

Figure 5: Visualization of Δ_F across both pools for a heterogeneous trader distribution $\Psi_A(\alpha)$ depending on mean relative benefit and the slippage tolerance. In blue areas the Nash equilibrium is Pool_N , in red areas, it is Pool_W , and in the white area in-between all liquidity distributions are Nash equilibria. Δ_F is cut off at 0.02 for better visibility and the dotted dark blue line visualizes where $\Delta_F = 0.01$. We set $x = 5,000,000 X$, $y = 5,000,000 Y$, $f = 0.003$ and $\omega = 0.01$.

where the sign change in front of the integral is the result of Y -assets being returned to the pool.

The amount of X the attacker holds after the transaction a_x^{out} can be found by equating the two integrals for a_y and \tilde{a}_y , using $\tilde{a}_y = (1-f)a_y$, and solving for a_x^{out} . This yields the profit of the sandwich attacker

$$\begin{aligned} U^A &= a_x^{\text{out}} - a_x^{\text{in}} \\ &= \frac{(1-f)^2 a_x^{\text{in}} ((1-p)x + (1-f)(a_x^{\text{in}} + \delta_{x,W}))^2}{((1-p)x)^2 + (2-f)(1-f)(1-p)x a_x^{\text{in}} + (1-f)^3 a_x^{\text{in}} (a_x^{\text{in}} + \delta_{x,W})} - a_x^{\text{in}}. \end{aligned}$$

□

Lemma 2. *A sandwich attack of size a_x^{in} is only profitable if the trader's transaction size exceeds*

$$\frac{f((1-p)x + a_x^{\text{in}}(1-f))}{(1-f)^2}.$$

Proof. The expression for δ_x^{\min} follows from Lemma 1 by solving $U^A = 0$. The minimum transaction size for which a profitable sandwich attack exists is obtained by setting $a_x^{\text{in}} = 0$. \square

Lemma 3. *The sandwich attacker's maximal input, a_x^s , for a transaction exchanging $\delta_{x,W}$ X-tokens with slippage tolerance s such that the victim's trade still executes can be given analytically.*

Proof. We consider a sandwich attack with initial input a_x^{in} to the front-running transaction. The output of the victim transaction selling $\delta_{x,W}$ becomes

$$\tilde{\delta}_{y,W} = - \int_{(1-p)x+(1-f)a_x^{\text{in}}}^{(1-p)x+(1-f)(a_x^{\text{in}}+\delta_{x,W})} \frac{-x \cdot y}{\xi^2} d\xi.$$

The victim's transaction will, however, only go through, if

$$\begin{aligned} \tilde{\delta}_{y,W} &\geq (1-s)\delta_{y,W} \\ &- \int_{(1-p)x+(1-f)a_x^{\text{in}}}^{(1-p)x+(1-f)(a_x^{\text{in}}+\delta_{x,W})} \frac{-x \cdot y}{\xi^2} d\xi \\ &\geq (1-s) \left(- \int_{(1-p)x}^{x+(1-f)\delta_{x,W}} \frac{-x \cdot y}{\xi^2} d\xi \right). \end{aligned}$$

Thus, the attacker's maximal input a_x^s increases the slippage incurred by the victim to their tolerance, i.e., $\tilde{\delta}_{y,W} = (1-s)\delta_{y,W}$. Solving for a_x^s , we find that the maximal input is

$$a_x^s = \frac{1}{2} \left(\frac{\sqrt{\delta_{x,W}^2(1-f)^2 + \frac{4(1-p)x((1-p)x+\delta_{x,W}(1-f))}{1-s}}}{1-f} - \frac{-2(1-p)x}{1-f} - \delta_{x,W} \right).$$

\square

B.2 Trade Sizes

Lemma 4. *A trade of size $\delta_{x,N}^{\text{opt}} = \max(0, p \cdot x(\sqrt{(1+\alpha)(1-f)} - 1)/(1-f))$ maximizes a sophisticated trader's utility U^R in Pool_N and in Pool_W the optimum is at $\delta_{x,W}^{\text{opt}} = \max(0, (1-p)x(\sqrt{(1+\alpha)(1-s)(1-f)} - 1)/(1-f))$.*

Proof. Without loss of generality, we maximize the trader's utility in each pool independently and start with Pool_N . The trader's utility in Pool_N is given by

$$U_N^T = (1+\alpha) \frac{(1-f)\delta_{x,N} p \cdot y}{(1-f)\delta_{x,N} + p \cdot x} - \frac{y}{x} \delta_{x,N}.$$

We differentiate the trader's utility in Pool_N , U_N^T , with respect

to the transaction size $\delta_{x,N}$ to find the transaction size $\delta_{x,N}^{\text{opt}}$ maximizing the trader's utility. We obtain

$$\partial_{\delta_{x,N}} U_N^T = \frac{(1+\alpha)(1-f)p^2 \cdot x \cdot y}{(\delta_{x,N}(1-f) + p \cdot x)^2} - \frac{y}{x},$$

and the two zero crossing of $\partial_{\delta_{x,N}} U_N^T$ are:

$$p \cdot x(\pm\sqrt{(1+\alpha)(1-f)} - 1)/(1-f).$$

For our parameters, $x, y, \alpha > 0$, $0 < f < 1$, $0 \leq p \leq 1$, the second derivative, $\partial_{\delta_{x,N}}^2 U_N^T$, is only negative for the following zero crossing

$$\delta_{x,N}^{\max} = p \cdot x(\sqrt{(1+\alpha)(1-f)} - 1)/(1-f).$$

Thereby, $\delta_{x,N}^{\max}$ maximizes the traders utility. The trader optimally sells $\delta_{x,N}^{\text{opt}} = \max(0, \delta_{x,N}^{\max})$ in Pool_N .

We proceed analogously as above for Pool_W and find the trader the optimally places $\delta_{x,W}^{\text{opt}} = \max(0, \delta_{x,W}^{\max})$ in Pool_W . In the previous,

$$\delta_{x,W}^{\max} = (1-p)x(\sqrt{(1+\alpha)(1-s)(1-f)} - 1)/(1-f).$$

The trade inputs to maximize the trader's utility can, thus, be determined analytically and are given by $\delta_{x,N}^{\text{opt}} = \max(0, \delta_{x,N}^{\max})$ and $\delta_{x,W}^{\text{opt}} = \max(0, \delta_{x,W}^{\max})$. \square

C Omitted Game Equilibria Proofs

C.1 Homogeneous Traders

Lemma 7. *The only Nash equilibria if $\partial_p F \neq 0$ are $p \in \{0, 1\}$. If $\partial_p F = 0$, all liquidity distributions are ε -equilibria in a homogeneous traders game.*

Proof. If the fees gradient is non-zero ($\partial_p F \neq 0$), the maxima is located at either corner point of the interval, as fees are proportional to p (cf. Lemma 6), and the fees gradient is non-zero, the maxima are located at either corner point of the interval. Otherwise, if the fees gradient $\partial_p F$ is zero, the fees across the entire interval are constant. Therefore, all liquidity distributions are ε -equilibria. \square

Theorem 1. *A liquidity distribution is an ε -equilibrium if there is no liquidity provider with initial liquidity distribution $(p_i l_i, (1-p_i) l_i)$, such that*

$$\frac{\max\{F(f, \alpha, s, y, 0, \omega), F(f, \alpha, s, y, 1, \omega)\}}{F(f, \alpha, s, y, p_i, \omega)} - 1 < \varepsilon.$$

Proof. For a liquidity distribution to qualify as an ε -equilibrium, no liquidity provider must see a possibility to increase their expected fees by more than than a factor $1 + \varepsilon$ through adjusting their liquidity distribution. Further, we know from Lemma 7 that a liquidity provider receives the

most fees either when all their liquidity is in Pool_N or all their liquidity is in Pool_W . Thus, the maximum relative increase to an LP's fees, with current liquidity distribution $(p_i l_i L, (1 - p_i) l_i L)$, is given as

$$\frac{\max\{F(f, \alpha, s, y, 0, \omega), F(f, \alpha, s, y, 1, \omega)\}}{F(f, \alpha, s, y, p_i, \omega)} - 1.$$

In case the previous fraction does not exceed ε for any liquidity provider, the current distribution is a Nash equilibrium. \square

C.2 Heterogeneous Traders

Theorem 2. *A liquidity distribution in a system with heterogeneous traders with distribution $\Psi_A(\alpha)$ is an ε -equilibrium if there is no liquidity provider with initial liquidity distribution $(p_i l_i L, (1 - p_i) l_i L)$, such that*

$$\frac{\max\{\sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, 0, \omega), \sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, 1, \omega)\}}{\sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, p_i, \omega)} - 1 < \varepsilon.$$

Proof. We proceed similarly to Theorem 1 and note that as long as no liquidity provider must see a possibility to increase their expected fees by more than a factor $1 + \varepsilon$ through adjusting their liquidity distribution, the configuration is a Nash equilibrium. The fees received by a liquidity provider LP_i are given by

$$\sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, p_i, \omega),$$

where $F(f, \alpha, s, y, p_i)$ is given by Lemma 6. We, thus, follow that the fees received by liquidity provider LP_i are also proportional to p_i in the heterogeneous case. Further, it holds that the liquidity provider receives the most fees either when all their liquidity is in Pool_N , all their liquidity is in Pool_W or the fees are constant for all liquidity distribution. Thus, the maximum relative increase to an LP's fees, with current liquidity distribution $(p_i l_i L, (1 - p_i) l_i L)$, is given as

$$\frac{\max\{\sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, 0, \omega), \sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, 1, \omega)\}}{\sum_{\alpha} \Psi_A(\alpha) F(f, \alpha, s, y, p_i, \omega)} - 1 < \varepsilon.$$

In case the previous fraction does not exceed ε for any liquidity provider, the current distribution is a Nash equilibrium. \square