Can Smaller Attacks Be More Effective? Understanding Probabilistic Sandwich Attack Strategies in Decentralized Exchanges

Bumho Son[a], Yunyoung Lee[b], Huisu Jang[c,1]

[a]*Department of Business Administration, Chung-Ang University, Heukseok-ro 84, Seoul, 06974, Korea*
[b]*Department of Artificial Intelligence & Data Science, Sejong University, Neungdong-ro 209, Seoul, 05006, Korea*
[c]*School of Finance, Soongsil University, Sangdo-ro 369, Seoul, 06978, Korea*

Abstract

Our study proposes a game-theoretic analysis of sandwich attacks on decentralized exchanges. We account for the uncertainty of state transitions during attackers' decision-making in a two-stage game. We introduce a novel model that modifies existing assumptions about attacker behavior and emphasize that attacker decisions are driven by a strategic trade-off between transaction fees and the unpredictability of their profits. Our findings challenge the belief that attackers always maximize slippage to their victims. In empirical study, we show that attackers often pursue smaller attacks instead of making the largest attack. By incorporating a behavioral approach, we confirm that our model more accurately captures observed attack patterns than existing models. Our research is not limited to blockchain but can be extended to other contexts involving probabilistic decision-making, providing insights into other domains. This broad applicability highlights the value of our model for tackling complex decisions in uncertain environments across various research domains.

Keywords: Decentralized finance, Behavioral approach, Game Theory, Sandwich attack

[1]Corresponding author: Huisu Jang, yej523@ssu.ac.kr
[2]All authors contributed equally to this work.

1. Introduction

In Ethereum networks, block proposers determine both which transactions are included in a block and their execution order, typically selecting transactions from the mempool based on factors like transaction fees (gas fees). Users often compete by offering higher fees to have their transactions processed more quickly. While this fee competition can seem fair, as it links willingness to pay with transaction speed, it can lead to manipulation in transaction ordering. This phenomenon, known as Maximal Extractable Value (MEV), allows block proposers and adversarial users to reorder, include, or exclude transactions for personal gain, potentially compromising fairness. MEV encompasses various forms, including arbitrage, liquidation, and sandwich attacks. While arbitrage and liquidation can enhance market efficiency in decentralized finance (DeFi), sandwich attacks are more problematic. In such attacks, an entity manipulates market prices to extract extra profits, adversely affecting victims and harming market efficiency.

To address the issue of sandwich attacks, we analyze the behavior of involved market participants to understand their strategies and impact on market dynamics. In a sandwich attack, adversaries execute transactions immediately before and after a victim's transaction—known as frontrunning and backrunning—to manipulate asset prices and profit from the price differences. This tactic resembles the frontrunning strategies of high-frequency traders (HFTs) in traditional markets but is more complex in Ethereum due to the block proposer's control over transaction ordering. A critical challenge is the uncertainty of transaction sequences within a block, leading adversaries to pay higher transaction fees to prioritize their transactions. Users can mitigate risks by setting slippage tolerance in Automated Market Makers (AMMs) like Uniswap, limiting how much the price can deviate from the expected value during execution. Analyzing these concepts is crucial for a deeper understanding of sandwich attacks in blockchain networks.

In this research, we present a game-theoretic model to determine the optimal attack volume for adversaries conducting sandwich attacks on decentralized exchanges (DEXs). Our two-stage approach considers that victims first choose their slippage tolerance, aware of potential sandwich attacks, and then adversaries observe this tolerance and strategically select their attack volume to maximize profit. Importantly, our model incorporates the concept of *external slippage*—the variability in attackers' expected profits due to uncertainty about external transactions within the same block. The model can be extended to a more general game-theoretic framework by modifying the scenario as follows. In our model, two players are involved: the victim and the attacker. The victim initiates communication by sending a message to the blockchain's initial state, which the attacker observes before responding with

a counter-message aimed at generating financial gain. However, the challenge arises from the presence of other messages within the blockchain that can also trigger state transitions, complicating the dynamics of the systems. Consequently, the attacker must make profit-maximizing decisions under the uncertainty that the state may change from the observed state by the time the decision is executed. Through this model, contrary to the assertion of Heimbach and Wattenhofer (2022) that attackers exploit maximum slippage, we find that attackers often opt for smaller attack volumes, balancing profit maximization with risk mitigation. This finding suggests that optimal attacker strategies are influenced by prevailing market conditions and the slippage settings chosen by victims.

Our findings highlight external slippage as pivotal in shaping attacker decisions. Attackers face a trade-off: paying higher priority fees increases the likelihood of earlier transaction execution, reducing profit variability but incurring higher costs; paying lower fees reduces costs but increases profit variance due to transaction order uncertainty. Prior studies suggest that attackers often set attack quantities right at the slippage tolerance boundary. However, our study reveals that attackers may not consistently follow this strategy. Our model quantifies the probability of a successful sandwich attack and, supported by empirical evidence, suggests that attackers may prefer smaller-scale attacks. We validate this model using real blockchain data from Uniswap V2 and V3 swap transactions in November 2022. The data confirm that attackers do not consistently choose maximum attack volumes, supporting our premise that they adopt more measured strategies. This real-world validation strengthens the practical relevance of our findings.

Furthermore, we introduce a behavioral approach, acknowledging that market participants may not be entirely rational. Literature reveals that crypto investors are heavily influenced by social factors and public sentiment, often exhibiting irrational, risk-seeking behavior and herding tendencies (Almeida and Gonçalves, 2023). Given this, we find that actual data align more closely with our model than with previous models, reinforcing its relevance in capturing complex market behaviors in cryptocurrency trading.

The structure of this paper is as follows. Section 2 reviews related research to contextualize our contributions. Section 3 outlines our proposed theoretical model and the decision-making process for attackers. Section 4 presents the empirical evidence supporting our model. Section 5 describes the mechanism of decentralized exchanges in detail. Finally, Section 6 summarizes our findings and suggests directions for future research.

2. Related Work

Blockchain technology has significant potential to innovate industries, particularly in reshaping the financial sector through the growth of decentralized finance (DeFi) ecosystems. Baliker et al. (2023) summarize blockchain applications in FinTech, while Nelaturu et al. (2022) highlight unresolved performance, security, and privacy obstacles in blockchain technology for FinTech applications. To address security concerns, Liu et al. (2022) propose a novel information security strategy using blockchain and edge computing, and Song et al. (2023) introduce a blockchain-based FinTech trust evaluation mechanism (BFTEM) to verify user trustworthiness.

MEV is a unique characteristic of cryptocurrency markets arising from ledger transparency and lack of centralized control. Werner et al. (2022) provide a comprehensive overview of DeFi and highlight MEV threats. Foundational work by Qin et al. (2022) and Daian et al. (2020) quantifies blockchain extractable value. In public blockchains, transactions are openly processed and publicly available in the mempool until included in a block, allowing entities controlling transaction order to exploit this privilege for financial gain. MEV types include arbitrage, liquidation, and sandwich attacks. While arbitrage and liquidation are generally benign and enhance market efficiency—arbitrageurs exploit price discrepancies, and liquidators quickly resolve insolvent collateral (Perez et al., 2021)—sandwich attacks are more problematic due to causing market instability.

Sandwich attacks manipulate transaction ordering to profit from price manipulation and have garnered significant attention (Zhou et al., 2021; Heimbach and Wattenhofer, 2022; Züst et al., 2021). These attacks occur on decentralized exchanges (DEXs) using Automated Market Maker (AMM) models with liquidity pools. Zhou et al. (2021) formalize the problem, and studies like Ferreira and Parkes (2022); Alpos et al. (2023) propose cryptographic methods to improve transaction ordering and reduce vulnerability. Züst et al. (2021) demonstrate increasing efficiency of trading bots in executing sandwich attacks and suggest mitigation strategies like splitting large transactions.

The work of Heimbach and Wattenhofer (2022) provides valuable insights into mitigating the sandwich attack problem by considering the optimal decisions of attackers and victims in a game-theoretical framework. This work shows that the proposed algorithm providing effective slippage tolerance outperforms the constant auto-slippage by the AMM, Uniswap. Building upon the findings of Heimbach and Wattenhofer (2022), our study extends the understanding of optimal decision making for attackers and victims. While this study has advanced our understanding, they leave unanswered questions about the variance of expected return of sandwich attack. The variance stems from the fact that the order of transaction in a block could

be different according to the priority fee of the transaction. The attackers attempt to acquire the sandwich attack opportunity by front-running and back-running their victim with higher priority fees to extract MEV (Qin et al., 2022). At this time, the priority fee is included in the cost to the attacker, and depending on how much this cost is spent, the profit from the attack varies along with the order of sandwich attack transactions it is executed in the block. Since the probability of attack success varies depending on the order of transactions within the block, the distribution of expected profits from sandwich attacks is determined by the priority fee, and this study seeks to present a new model that takes this into account.

Canidio and Danos (2024) provides a game-theoretic analysis of front-running attacks within the blockchain ecosystem. The study introduces a novel commit-reveal protocol designed to prevent front-running attacks while maintaining legitimate fee competition among users, thereby ensuring that higher fees result in earlier transaction processing. By doing so, the proposed model mitigates fee competition from malicious actors attempting to carry out front-running attacks, ultimately benefiting honest participants. Although the theoretical framework of this study shares significant parallels with our work, our model specifically addresses sandwich attacks, which encompass both front-running and back-running strategies, which can cause more severe loss to the honest users than front-running attacks. Additionally, the two-stage game model in Canidio and Danos (2024) provides game-theoretic decisions under a deterministic state, whereas our model focuses on stochastic states, where the expected state and its transition is governed by probabilities rather than being deterministic.

Additionally, beyond technical vulnerabilities, the behavior of market participants plays a crucial role in the dynamics of sandwich attacks. Literature reveals that cryptocurrency investors often exhibit irrational and risk-seeking behavior, influenced by social factors and public sentiment. Almeida and Gonçalves (2023) provide a systematic review showing that the crypto investment landscape is dominated by investors driven by the pursuit of high profits, leading to intense herding behavior and market inefficiency. Notably, crypto investors tend to view themselves as superior traders compared to non-crypto participants, and sophisticated investors are more inclined to demand cryptocurrency as a hedge against risks in the real economy (Colombo and Yarovaya, 2024). According to Hackethal et al. (2022), cryptocurrency investors frequently shift their portfolios toward even riskier assets after adopting cryptocurrencies. This irrational behavior can exacerbate the impact of sandwich attacks, as investors may not take optimal protective measures, making them more susceptible to exploitation.

Recognizing that attack success probability and expected profits depend on pri-

ority fees and investor behavior, our study presents a new model that accounts for these factors. By incorporating both the technical aspects of transaction ordering and the behavioral tendencies of investors, we aim to provide a more comprehensive understanding of sandwich attacks and their mitigation.

3. Model

3.1. Problem Definition: Sandwich Attack

The sandwich attack is typically done on the DEX liquidity pool by actual attackers or the predatory trading bots. It aims to attack the traders who want to swap two tokens in the liquidity pool by adding the front- and back- run transactions to the trader's transaction. The basic principles of success of the sandwich attack is to temporarily manipulate the swap price in the liquidity pool. The front-run transaction increase the swap price of the token that the trader want to receive, and the back-run transaction realize the profit of attackers. For the rest of our model analysis, we will consider the sandwich attack on single liquidity pool consisted of tokens X and Y with swap transaction fee f.

We start by formulating the sandwich attack as a two-stage game among two players: Victim (V) and Attacker (A). The first stage of sandwich attack starts with V observing the initial state $t_0 = (x_0, y_0)$ of the liquidity pool, where x_0 and y_0 each denotes the amount of tokens X and Y in the liquidity pool. After observing t, V sends transaction $T_V = (\delta_{v_x}, \delta_{v_y}, s)$ to the mempool, which is the set of pending transactions to be recorded on the block. T_V contains two messages. First message is that V is willing to exchanges δ_{v_x} amount of tokens X to δ_{v_y} amount of tokens Y. Second is the maximum slippage tolerance s, which determines how much loss the trader is willing to accept and proceed with the swap. When the transaction T_V is included in the block, smart contract of the liquidity pool runs and activates the swap if the maximum slippage tolerance condition is satisfied.

The second stage of sandwich attack occurs because mempool is public. When the first stage game ends, A observes T_V in the public mempool, and the second stage game begins. A creates the front-run transaction $T_A^{front} = (\delta_{a_x}^{in}, \delta_{a_y})$ and back-run transaction $T_A^{back} = (\delta_{a_y}, \delta_{a_x}^{out})$ as a countermessage of T_V. Then A binds three transactions $T_A = (T_A^{front}, T_V, T_A^{back})$ and submits it to the mempool with base fee b and priority fee r. If T_A is recorded on the block and run successfully, V and A each earns payoff P_V and P_A, respectively.

The key point of the two-stage game is that the state t_0' when T_A is executed can be different from t_0, the state when T_A is sent to the mempool. While A sends countermessage to gain the profit based on the observed state t_0, the actual profit

can be different since it is based on t_0'. On the other hand, V manages the risk of state difference by setting s. Mathematically, T_V will be only exeucted when the following condition is satisfied:

$$\tilde{\delta}_{v_y} \geq (1-s)\delta_{v_y}, \tag{1}$$

where $\tilde{\delta}_{v_y}$ denotes the actual amount of tokens Y that V will receive based on the t_1.

We can represent the game tree for sandwich attack as in Figure 1. The observed state when both players make decision is different from the state when their strategy is executed and gain profit.

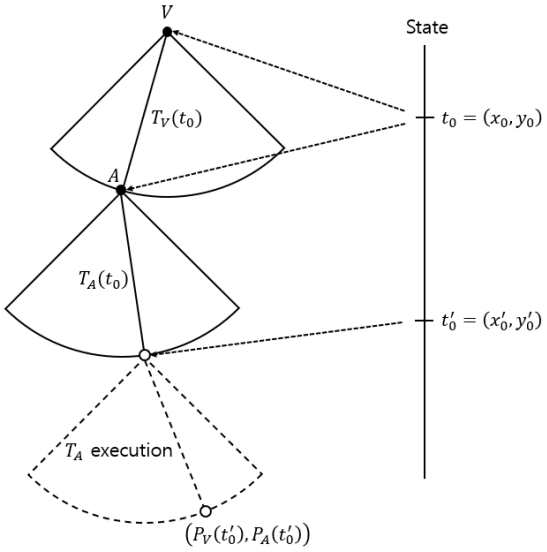

Figure 1: Game tree and corresponding state of the liquidity pool for each stage

Sending a countermessage T_V requires two kinds of cost: base fee $b > 0$ and priority fee $r \geq 0$. While b is exogenous variable, r acts as the auction price. The higher r gives T_V higher probability to be located on the front part of the block. In Section 3.2, we quantitatively define the effect of r on the position of T_V in the block as the difference between t_0 and t_0'.

3.2. External Slippage

At the point that V submits a swap transaction, T_V, to the liquidity pool, he expects to receive δ_{v_y} amounts of token Y as follows:

$$\delta_{v_y} = \frac{y_0(1-f)\delta_{v_x}}{x_0 + (1-f)\delta_{v_x}} \tag{2}$$

However, the amounts of tokens in the liquidity pool will change for every swap transactions. Therefore, the actual states of the liquidity pool when each transaction is executed will be different. We will define the states of the liquidity pool just before when transactions T_A^{front}, T_V, and T_A^{back} are executed as $t'_0 = (x'_0, y'_0), t'_1 = (x'_1, y'_1)$, and $t'_2 = (x'_2, y'_2)$, respectively.

It is important to point out that t'_0 will be different from t_0, which is the crucial difference between our work and Heimbach and Wattenhofer (2022).

In Heimbach and Wattenhofer (2022), the authors assumed that T_A always occupies the very front position in the block. From the perspective of the two-stage game defined in Section 3.1, this assumption can be interpreted as $t_0 = t'_0$. For A, this implies that the amount of token Y received after executing T_A is deterministic. Consequently, A can confidently predict whether T_V will be successfully executed. Thus, A's optimal strategy reduces to choosing between two options: sending the optimal message T_A that satisfies Equation 2 at the boundary, or not sending T_A at all.

However, we argue that A cannot confirm the success of the attack, as it depends on t'_0, which is not observable by A at the time of sending a countermessage. It occurs because of the *external slippage*. We will define *internal slippage* and *external slippage*, which when added together will equal to the total slippage that the transaction of trader will actually face. The sandwich attack of attackers itself causes the *internal slippage* to the transaction of trader who would be a victim of sandwich attack. On the other hand, *external slippage* means the change in the state of liquidity pool that is not expected by the attackers. Even though the attacker who is willing to do the sandwich attack tries to put its transactions at the front of the block, we cannot affirm that it will locate as the first transaction of the newly created block. Therefore, there exists the possibility of other transactions participating in the liquidity pool may precede the T_V^{front}. These other transactions affects the state of the liquidity pool and it becomes the *external slippage*. As a result, the sum of the internal and external slippage should be small enough to satisfy 2, which leads to the execution of T_V.

During the overall procedures of sandwich attack, the external slippage will be only occurred just before the T_A^{front} is executed because the transactions $(T_A^{front}, T_V, T_A^{back})$ will be bundled together and no other transactions will be able to get in between them. We will assume that the external slippage will change (x_0, y_0) to (x'_0, y'_0), where $x'_0 \sim N(x_0, \frac{\sigma^2}{r})$ and $x_0 y_0 = x'_0 y'_0$. σ^2 denotes the magnitude of external slippage. This assumption is plausible in two senses. First, a change in the state of the liquidity pool can cause the amount of X to increase or decrease for same chance. x'_0 following the normal distribution well captures the bi-directional change of the

pool state. Second, as the attacker pays more priority fee r, it becomes more likely to place its transactions at the front part of the block. It reduces the likelihood of sandwich attack transactions being preceded by other transactions participating in the same liquidity pool, which leads to the decrease of the magnitude of external slippage. Even though the normal distribution assumption has the risk that the number of tokens can have negative value, it has significantly small probability because $\frac{\sigma^2}{r}$ will be much smaller than x_0 in most cases.

3.3. Optimal strategy

Given the state of the liquidity pool $t'_0 = (x'_0, y'_0)$ just before T_A are executed, we can derive the quantities of the number of tokens that each transactions will receive. When T_A^{front} is executed, A will receive δ_{a_y} tokens Y as follows:

$$\delta_{a_y} = \frac{y'_0(1-f)\delta_{a_x}^{in}}{x'_0 + (1-f)\delta_{a_x}^{in}} \tag{3}$$

Therefore, (x'_1, y'_1) becomes $x'_1 = x'_0 + \delta_{a_x}^{in}$ and $y'_1 = \frac{x'_0 y'_0}{x'_0 + (1-f)\delta_{a_x}^{in}}$. Consequently, when T_V is executed, V will receive $\tilde{\delta}_{v_y}$ tokens Y as follows:

$$\tilde{\delta}_{v_y} = \frac{\frac{x'_0 y'_0}{x'_0 + (1-f)\delta_{a_x}^{in}}(1-f)\delta_{v_x}}{x'_0 + \delta_{a_x}^{in} + (1-f)\delta_{v_x}} \tag{4}$$

(x'_2, y'_2) becomes $x'_2 = x'_1 + \delta_{v_x}$ and $y'_2 = \frac{x'_1 y'_1}{x'_1 + (1-f)\delta_{v_x}}$. Finally, when T_A^{back} is executed, A will receive $\delta_{a_x}^{out}$ tokens Y as follows:

$$\delta_{a_x}^{out} = \frac{x'_2(1-f)\delta_{a_y}}{y'_2 + (1-f)\delta_{a_y}} \tag{5}$$

When the every procedures are over, the profit of adversary sandwich attackers becomes as follows:

$$P_A = \delta_{a_x}^{out} - \delta_{a_x}^{in} - 2b - 2r \tag{6}$$

However, T_V will not be executed if the total slippage is larger than s. Therefore, we can write the profit more precisely as follows:

$$P_A = \begin{cases} \delta_{a_x}^{out} - \delta_{a_x}^{in} - 2b - 2r, & \text{if } \tilde{\delta}_{v_y} \geq (1-s)\delta_{v_y} \\ -2b - 2r, & \text{otherwise} \end{cases} \tag{7}$$

Since $\delta_{a_x}^{out}$ is the random variable affected by x'_0, we will now consider the expected profit $\mathbb{E}[P_a | \delta_{a_x^{in}}, r]$.

The adversaries aim to maximize its expected profit by controlling $\delta_{a_x}^{in}$ and r. These two parameters that adversaries can control affects the expected profit in the terms of attack success probability and cost. As they increase $\delta_{a_x}^{in}$, $\mathbb{E}[\delta_{a_x}^{out}]$ will increase since the gave more input to the liquidity pool. On the other hand, the victim will face more internal slippage and it will make harder to meet the slippage tolerance condition of T_V. Another parameter r directly affects the expected profit since itself is a cost, and also affects the attack success probability by changing the magnitude of external slippage.

Considering the effects of $\delta_{a_x}^{in}$ and r, we can derive the attack success probability $f(\delta_{a_x}^{in}, r)$ (i.e. probability of slippage tolerance condition is met) as follows:

$$f(\delta_{a_x}^{in}, r) = \mathbb{P}\left(\frac{\frac{x_0' y_0'}{x_0' + (1-f)\delta_{a_x}^{in}}(1-f)\delta_{v_x}}{x_0' + \delta_{a_x}^{in} + (1-f)\delta_{v_x}} \geq (1-s)\frac{y_0(1-f)\delta_{v_x}}{x_0 + (1-f)\delta_{v_x}}\right) \tag{8}$$

We show that $f(\delta_{a_x}^{in}, r)$ has a closed form solution in Lemma 1.

Lemma 1. *If a quadratic equation* $h(x) = -(1-s)x^2 - (1-s)(2\delta_{a_x}^{in} + (1-f)\delta_{v_x})x + x_0^2 + (1-f)\delta_{v_x}x_0 - \delta_{a_x}^{in}(1-s)(\delta_{a_x}^{in} + (1-f)\delta_{v_x}) = 0$ *has two real roots* $h_1 < h_2$*, then* $f(\delta_{a_x}^{in}, r) = \Phi(\frac{h_2 - x_0}{\sigma^2/r}) - \Phi(\frac{h_1 - x_0}{\sigma^2/r})$*, where* $\Phi(\cdot)$ *is the cumulative distribution function of standard normal distribution.* [3]

Proof.

$$\begin{aligned} f(\delta_{a_x}^{in}, r) &= \mathbb{P}(h(x_0) \geq 0) \\ &= \mathbb{P}(-(1-s)(x_0' - h_1)(x_0' - h_2) \geq 0) \\ &= \mathbb{P}(h_1 \leq x_0' \leq h_2') \ (\because 0 \leq s \leq 1) \\ &= \Phi\left(\frac{h_2 - x_0}{\sigma^2/r}\right) - \Phi\left(\frac{h_1 - x_0}{\sigma^2/r}\right) \end{aligned} \tag{9}$$

□

Using Lemma 1, we can derive the expected profit of A as follows:

$$\mathbb{E}[P_A] = \left(\Phi\left(\frac{h_2 - x_0}{\sigma^2/r}\right) - \Phi\left(\frac{h_1 - x_0}{\sigma^2/r}\right)\right)(\delta_{a_x}^{out} - \delta_{a_x}^{in}) - 2b - 2r \tag{10}$$

[3] If a distribution other than the normal distribution is assumed for x_0', the lemma still holds as long as $\Phi(\cdot)$ represents the cumulative distribution function of $\frac{x_0' - x_0}{\sigma^2/r}$.

In the perspective of attackers, A will act to maximize its expected profit. Specifically, he will solve the following two variables optimization problem.

$$\begin{aligned} \max_{\delta_{a_x}^{in}, r} \quad & \mathbb{E}[P_A] \\ \text{s.t.} \quad & \delta_{a_x}^{in}, r \geq 0 \end{aligned} \tag{11}$$

We can compare the optimal solution $\delta_{a_x}^s$ derived from Heimbach and Wattenhofer (2022) with the optimal solution $(\delta_{a_x}^*, r^*)$ of Equation 11. $\delta_{a_x}^*$ will be smaller than $\delta_{a_x}^s$ because $\delta_{a_x}^s$ is based on the belief that optimal adversaries will attack as much property that just satisfies the slippage tolerance. Since we have pointed out the existence of external slippage, adversary should maintain enough safety margin or high level of transaction fee to guarantee high possibility of the attack success.

4. Empirical Study

In this section, we employ both simulated data and actual market data pertaining to sandwich attacks to examine the ability of our proposed model to accurately represent real-world scenarios. Through analysis of simulation results, we confirm that the optimal strategies for attackers, as derived from our model, diverge significantly from those suggested by existing models Heimbach and Wattenhofer (2022). Additionally, we have verified that analysis of actual market data yields results consistent with those from our simulations. From a behavioral economics standpoint, experimental evidence indicates a higher likelihood of the proposed model accurately reflecting optimal decision-making compared to the existing research. Finally, we have explored the implications of the attacker's optimal strategies, as identified through both the proposed and existing models, on the victim's losses and overall social welfare.

4.1. Data Collection

We analyze the sandwich attacks detected on Uniswap V2 and V3 to comprehend the practical nature of how these attacks are executed. Initially, we identify sandwich attacks carried out between block 16,000,000 (Nov 18, 2022) and block 16,010,000 (Nov 20, 2022) employing the detection algorithm proposed in Park et al. (2024). To streamline our analysis, we exclusively focus on attack cases where WETH (Wrapped Ether) is among the underlying assets in the pool. Furthermore, we filter the attacks to include only cases where the attacker's profit is denominated in WETH, and the sum of the other assets amounts to zero. As a result, our empirical analysis comprises of 1463 sandwich attacks, consisting of 4810 swap transactions. To acquire comprehensive details for each transaction within our dataset, we leverage the Erigon node

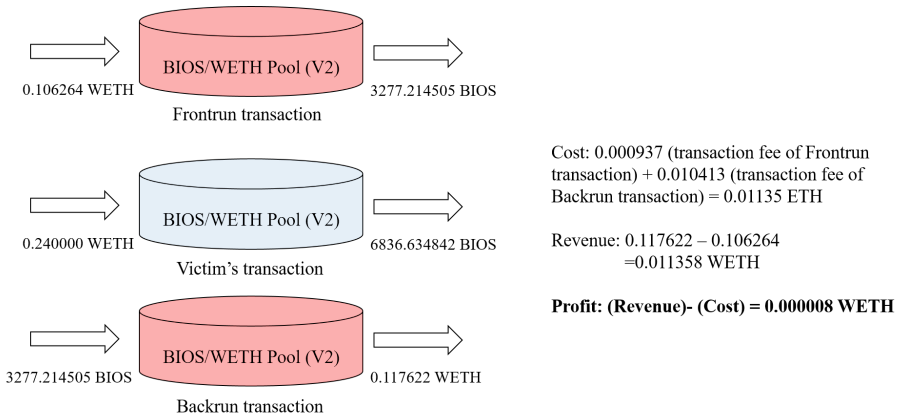

Figure 2: Example of a sandwich attack detected at the Uniswap V2 BIOS/WETH Pool in block 16009112.

of the Ethereum blockchain. Our access to the Erigon archive node facilitated the retrieval of substantial information pertaining to the transactions. This encompassed critical data such as the transaction fee (computed as the product of gas price and used gas), coinbase transfer (direct transfer of Ethereum to the proposer of a block), transaction index, and the quantity of underlying assets exchanged within Uniswap pools.

Table 1: Summary of the sandwich attack dataset

Panel A: Uniswap V2 & V3 Comparison			
	Uniswap V2	**Uniswap V3**	
# of frontrun transactions	1263	423	
# of victim transactions	1089	349	
# of backrun transactions	1263	423	
Panel B: Descriptive Statistics (in WETH)			
	Cost	**Revenue**	**Profit**
# of observations	1463	1463	1463
Mean	0.026640	0.028505	0.001865
Median	0.008073	0.008609	0.000101
Standard Deviation	0.085542	0.089568	0.009376
Min	0.001611	0.001612	-0.000626
Max	1.403588	1.464834	0.184897

Moreover, our dataset includes pools from both Uniswap V2 and V3. Consequently, we compute the marginal price for each pool differently, depending on the respective version of the pool. For Uniswap V2 pools, we denote the amount of each

token deposited at time t as x_t, y_t for two tokens X and Y, and L_t denoting the liquidity of the pool. With the invariant formula $x_t y_t = L_t^2$, the marginal price of token Y respect to token X can be calculated as $p_t = -\frac{\partial x_t}{\partial y_t} = \frac{L_t^2}{y_t} = \frac{x_t}{y_t}$. To apply this formula, we collect the token reserves' amount for Uniswap V2 pools at each block. However, for Uniswap V3 pools, the marginal price cannot be calculated simply as the ratio of token reserves between X and Y since the provided liquidity varies based on the given price interval. Therefore, we additionally obtain SqrtPriceX96 at time t, where SqrtPriceX96 $= 2^{96} \times \sqrt{p_t}$. Thus, the marginal price can be computed as $p_t = (\text{SqrtPriceX96})^2/2^{192}$ for Uniswap V3 pools. Table 1 represents the descriptive statistics of sandwich attack dataset used for our empirical analysis.

To enhance readers' understanding of the sandwich attack, we provide an illustration of an actual sandwich attack detected at block 16009112 in Figure 2. The attack was executed at BIOS/WETH pool of Uniswap V2, where the attacker initially inflated the price of BIOS by buying 3277.214505 BIOS with 0.106264 WETH. Subsequently, the victim had to purchase 6836.634842 BIOS at an elevated price, further driving up the price OF BIOS in the pool. Finally, the attacker sold the 3277.214505 BIOS obtained in the frontrun transaction at a higher price, resulting in an arbitrage revenue of 0.001358 WETH. After deducting the transaction fees of 0.00135 ETH for both the frontrun and backrun transactions, the total profit from the attack amounted to 0.000008 WETH.

4.2. Simulation Results

We have done numerical analysis of our proposed model framework. For the simulation, we set parameters $s = 0.001, x_0 = 1000, y_0 = 1000, f = 0.003, \delta_{v_x} = 100, \sigma = 0.01, b = 0$ [4]. Figure 3 shows the surface of expected profit by varying $\delta_{a_x}^{in}$ and r. The red dot is the optimal $(\delta_{a_x}^*, r^*) = (0.405, 0.005)$, while the dashed red line is the line with $\delta_{a_x}^{in} = 0.405$. In this parameter setting, $\delta_{a_x}^s$ is 0.524. We can check that there exists an optimal solution at a point smaller than when we do not account for the external slippage.

Our study differs from the previous work by Heimbach and Wattenhofer (2022) in that we consider the impact of the priority fee, r, on the probability of success. As a result, for the same value of r, the expected profit becomes a convex function of $\delta_{a_x}^{in}$. This distinction causes the optimal δ_{a_x} in our study to be smaller compared

[4]The parameters used in simulation are based on that used in the Heimbach and Wattenhofer (2022). It is designed to reflect the conditions of a realistic swap pool as closely as possible, and is also set to show how the optimal solution of the model presented in this study differs from previous studies.

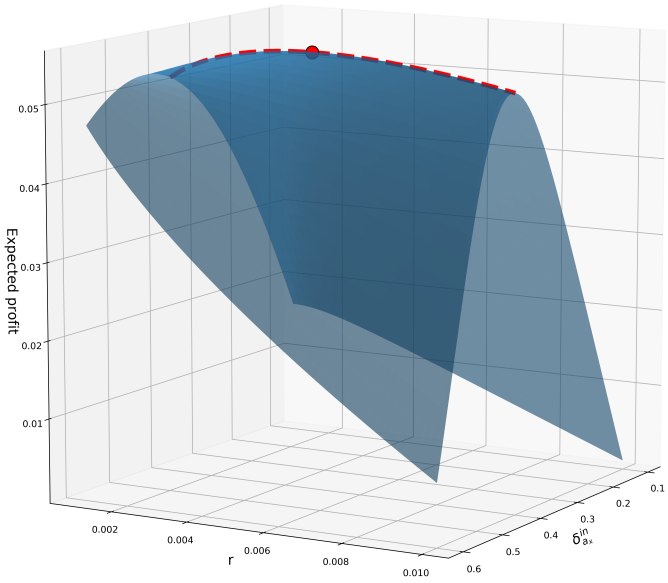

Figure 3: Simulation result of the optimal solution of Equation 11

to the optimal δ_{a_x} found in previous research.

However, $\delta^*_{a_x}$ is far smaller than the attack amount from the real-world transaction data. This is a major limitation of our model structure. We assumed that x'_0 will follow a normal distribution, which has the advantage of fully reflecting changes in both directions. However, the problem with this assumption is that it inherently assumes a low probability of attack success. For example, applying the optimal attack volume derived from Heimbach and Wattenhofer (2022) to the current model results in only a 50% chance of attack success.

4.3. Empirical Analysis of Sandwich Attacks

To support our theoretical model that the amount of fee that attackers are willing to pay to the block proposers (or validators) can significantly influence the position of their transactions within a block, we empirically investigate the relationship between transaction costs and transaction indices. Accordingly, we compute the correlation between transaction index and the cost-related variables(gas price, transaction fee,

and the overall cost=transaction fee + coinbase transfer). Our attack model is based on the assumption that block proposers determine the transaction order within the block based on the potential revenue they can receive from each transaction. As a result, in addition to the absolute value of each cost-related variables, we also compute the relative rankings of transaction index, gas price and overall cost within the block in descending order. For instance, if a transaction has a lower gas price ranking, it indicates a higher gas price compared to other transactions. Using these data, we begin by calculating the Pearson's correlation matrix among transaction ranking, gas price, overall cost, cost ranking, and gas price ranking. As shown in Table 2, the correlation matrix reveals a high correlation between the transaction order and the relative ranking of gas price (0.7912). This suggests that block proposers accord higher priority to gas price compared to other variables when determining the sequence of transactions within the block. This behavior of block proposers seems quite reasonable, given that the actual transaction fee cannot be precisely predicted in advance, as block proposers face challenges in estimating the exact amount of gas used for each transaction. Instead, they utilize the relative gas price ranking of the transactions in the mempool as a criterion for establishing the order of transactions. Additionally, Figure 4 plots the gas price ranking and transaction order ranking of 5,000 transactions, visually confirming the high correlation between them. It clearly demonstrates that transactions with a higher gas price compared to others are positioned earlier in the block. We also provide the Spearman and Kendall's correlation between transaction order and the cost-related variables in Table 3. The results closely align with those in Table 2.

Table 2: Pearson's Correlation Matrix between transaction order and cost-related variables

	tx ranking	**gas price**	**gas price ranking**	**overall cost**	**cost ranking**
tx ranking	1.0	-0.1160	**0.7912**	-0.0147	0.1603
gas price	-0.1160	1.0	-0.1280	-0.1262	-0.0803
gas price ranking	**0.7912**	-0.1280	1.0	0.0064	0.0576
overall cost	-0.0147	0.1262	0.0064	1.0	-0.0533
cost ranking	0.1603	-0.0803	0.0576	-0.0533	1.0

To investigate the practical evidence of our sandwich attack model proposed in Section 3, we assess whether the attackers attempt to maximize their profit from the attack by selecting the highest possible input amount ($\delta_{a_x}^{in}$), thereby reaching the maximum slippage tolerance of the victim's transaction. However, in order to correctly compare the slippage tolerance chosen by the victim with the actual slippage observed through the transaction, we require the historical memory pool data of the Ethereum blockchain, as the slippage tolerance itself is not recorded in the main

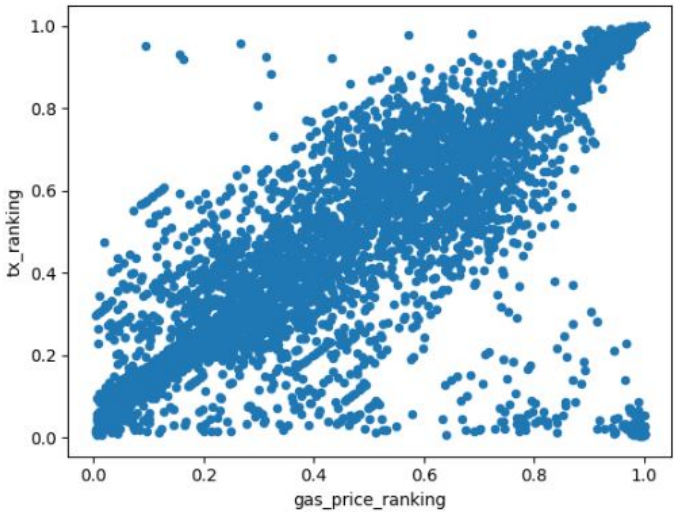

Figure 4: The transaction order and the gas price ranking of 5,000 transactions.

Table 3: Spearman and Kendall's Correlation Matrix between transaction order and cost-related variables

Panel A: Spearman Correlation					
	tx ranking	**gas price**	**gas price ranking**	**overall cost**	**cost ranking**
tx ranking	1.0	-0.4351	**0.6842**	-0.1358	0.1464
Panel B: Kendall Correlation					
	tx ranking	**gas price**	**gas price ranking**	**overall cost**	**cost ranking**
tx ranking	1.0	-0.3978	**0.6814**	-0.09678	0.1187

blockchain. Nevertheless, due to constraints in accessing historical memory pool data, which are not stored on the blockchain, our analysis uses Uniswap's default maximum slippage thresholds of 0.5% and 5.5% as proxies for these values, rather than the actual historical pool data. In specific instances, users may opt for the highest permissible slippage, with Uniswap's maximum slippage threshold set at 20%. Consequently, in our empirical analysis, we estimated slippage by examining the transaction volume between the victim and the attacker, employing the most relevant slippage benchmarks of 0.5%, 5.5%, and 20%.

We introduced a variable termed "ratio", representing the proportion of the actual attack input volume ($\delta_{ax}^{in}(actual)$) executed by the sandwich attacker relative to the maximum feasible attack input volume ($\delta_{a_x}^{in}(feasible)$), constrained by the victim's maximum slippage allowance. Figure 5 presents a box plot for the 'ratio' variable in our sandwich attack dataset, clearly indicating that the ratio value for 75% of the total data is below 1. This suggests that attackers do not use the full

maximum possible attack volume suggested by Heimbach and Wattenhofer (2022)'s study. However, in approximately 25% of the instances, the ratio exceeds 1, which may be considered an anomaly resulting from the actual slippage value reported by the victim being unknown.

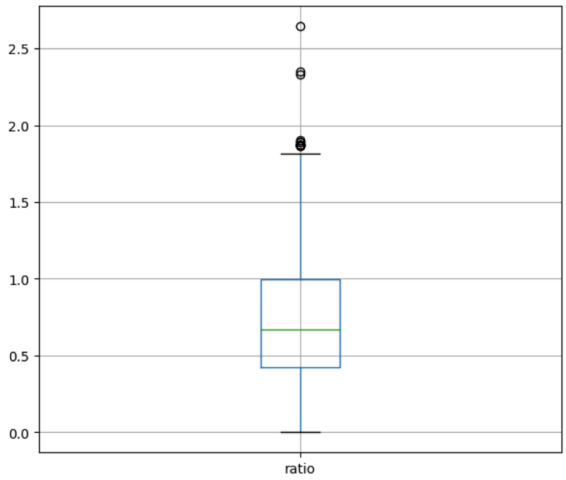

Figure 5: A box plot of the ratio of the actual attack input volume ($\delta_{a_x}^{in}(actual)$) executed by the sandwich attacker relative to the maximum feasible attack input volume ($\delta_{a_x}^{in}(feasible)$)

Our analysis confirms that the optimal attack volume proposed by our theoretical model is consistent with actual data. However, attacks aiming to achieve maximum slippage in a victim's transaction, as described in existing research (Heimbach and Wattenhofer (2022)), constitute less than 10% of all documented successful sandwich attacks. Although there are some inconsistencies between empirical data and our model, refining the model's assumptions could lead to a more accurate reflection of real-world data.

We have examined the function call parts within each transaction on the data collected in Section 4.1 to collect the slippage information reported by the victims. Consequently, we have selectively gathered the transaction data from the original data in Section 4.1 where information on 'amountInMax' or 'amountOutMin' is available. Following this refined extraction method, we obtained a subset of data from the original dataset, resulting in a total of 148 transactions with slippage information. As previously indicated, the actual subset data reveal that the victim's exchange volumes stay over the maximal slippage threshold depending on the front-running attacker's transaction not fully attacking the maximal slippage threshold. Table 4 shows the descriptive statistics of the slippage s submitted by victims and the ratio between the actual exchange amount $\tilde{\delta_{v_y}}$ to the minimum exchange amount $(1-s)\delta_{v_y}$

with the maximum slippage threshold. On average, traders submit a slippage rate of approximately 10 percent, while the actual volume exchanged from the liquidity pool exceeds the required minimum volume under the submitted slippage by an average of 4.16 percent. In Heimbach and Wattenhofer (2022), the discussion revolved around the strategy of a sandwich attacker exploiting the total of the maximum slippage of the victim, thus limiting the victim to exchanging only the minimal amount permissible. However, empirical evidence has shown that victims, on average, exchange amounts that exceed this minimum threshold. The model introduced, which incorporates external slippage, offers a better theoretical and experimental explanation for this discrepancy compared to the prevailing model.

Table 4: Descriptive statistics of the slippage s submitted by victims and the ratio between the actual exchange amount $\tilde{\delta_{v_y}}$ to the minimum exchange amount $(1-s)\delta_{v_y}$ with the maximal slippage threshold.

	Slippage s	Amount ratio $\tilde{\delta}_{v_y}/(1-s)\delta_{v_y}$
Mean	0.1175	1.0416
Median	0.0932	1.0001
Standard Deviation	0.0967	0.0997
Min	0.0037	1.0000
Max	0.3863	1.4715

To demonstrate that real sandwich attackers can achieve greater profits by optimizing their strategies using the proposed model, we present empirical evidence in Table 4. This table compares the profits realized by attackers from real data with the potential profits obtained by applying our proposed model. The results show that the profits achieved using the proposed model are statistically significantly higher than those observed in the real data, as confirmed by a t-test analysis. By adjusting the cost of the sandwich attack in response to external slippage, the proposed model yields a higher expected return from the perspective of the Sharpe ratio. Furthermore, the model increases the likelihood that victims receive a more favorable exchange amount compared to the minimum exchange amount determined by their submitted slippage.

4.4. Behavioral Approach

In Section 3.3, we delineates the derivation process for ascertaining the optimal quantity of the token X alongside the priority fee required by an attacker. Utilizing the developed profit model, we have conceptualized the probabilities of an attacker's decision-making process within the ambit of a behavioral model. This includes the

Table 5: Descriptive statistics of the realized attacker's profit and the optimal profit based on the proposed model

	Realized profit (in WETH)	Optimal profit (in WETH)
Mean	0.00131	0.02028
Median	0.00002	0.01310
Standard Deviation	0.00835	0.02352
Min	-0.00063	-0.00099
Max	0.09590	0.13338
Sharpe Ratio	0.15689	0.86383
T-statistic	-9.24595	
p-value	< 0.001	

integration of a rationality parameter for the attacker, enhancing the model's precision in predicting the attacker's potential actions, as expounded upon in Kim et al. (2019). This provides a basis for modeling the realistic behavior of sandwich attack participants, as well as evaluating the likelihood of the model proposed in this study.

We propose a utility function for the attacker, premised on the expectation of profit, to estimate the probabilities of the attacker selecting specific actions, $\delta_{a_x}^{in}$ and r. This is formally presented as follows:

$$Pr(\delta_{a_x}^{in}, r) = \frac{exp(\lambda \mathbb{E}[P_a | \delta_{a_x}^{in}, r])}{\iint_{\delta_{a_x}^{in}, r} exp(\lambda \mathbb{E}[P_a | \delta_{a_x}^{in}, r])}, \tag{12}$$

where λ represents the rationality parameter of the attacker. Notably, λ is a non-negative value that correlates with the attacker's level of rationality, influencing the likelihood of selecting the optimal $\delta_{a_x}^{in}$ and r. A λ value at zero suggests that the attacker's selections are entirely random, whereas an infinite λ indicates a deterministic approach towards selecting the optimal action.

Table 6 presents a comparative analysis of the log likelihood estimates for our proposed model and a nested model that excludes external slippage, based on UniSwap data. We evaluated these models across four distinct rationality parameters ($\lambda = 1, 0.5, 0.1, 0.01$) to determine how the inclusion of external slippage impacts the model's alignment with real-world scenarios.

The empirical results clearly show that our proposed model, which incorporates external slippage, yields higher log likelihood values than the nested model for λ values of 1, 0.5, and 0.1. This improvement is statistically significant, as confirmed by Likelihood Ratio (LR) tests, which indicate that the proposed model offers a more realistic framework by capturing the influence of external slippage on attacker

Table 6: Comparison of log likelihood estimates between proposed model and nested model

		Proposed model	Nested model
$\lambda = 1$	Log Likelihood	-303.00	-326.51
	Average probability	0.196	0.175
	LR-test	p<0.001	
$\lambda = 0.5$	Log Likelihood	-430.09	-441.66
	Average probability	0.099	0.093
	LR-test	p<0.001	
$\lambda = 0.1$	Log Likelihood	-700.95	-703.23
	Average probability	0.023	0.022
	LR-test	p<0.05	
$\lambda = 0.01$	Log Likelihood	-838.31	-838.54
	Average probability	0.011	0.011
	LR-test	p>0.05	

Notes. The LR-test row presents the log likelihood ratio test for each λ value, comparing the proposed model with external slippage and the nested model without external slippage. The significant p-values indicate that the proposed model has a statistically higher log likelihood than the nested model, underscoring the practical relevance of accounting for external slippage.

decision-making. However, at $\lambda = 0.01$, the difference between models is no longer statistically significant (Kim et al. (2019); Chung et al. (2020)). This outcome aligns with the theoretical expectation that, under a low rationality parameter, the distinction between optimal and sub-optimal choices diminishes, resulting in similar likelihoods for both models.

These results support our assertion that a model incorporating external slippage provides a more accurate and realistic estimation of an attacker's behavior within DEX environments. Consequently, the proposed model contributes a valuable behavioral framework for understanding sandwich attacks, setting it apart from existing models that overlook external market factors.

5. Automated Market Maker (AMMs)

In this study, we focus on sandwich attacks within Decentralized Exchanges (DEX) and examine the vulnerabilities they expose. This section provides a comprehensive overview of the fundamental structure of DEXs, with particular emphasis on slippage, a crucial factor in our research. Slippage, the difference between the expected price of a trade and the price at which it is executed, is more pronounced in DEXs due to the reliance on liquidity pools and the decentralized nature of the trading process. This makes slippage a critical aspect to consider when analyzing the mechanics of sandwich attacks.

As highlighted in Mohan (2022), the decentralized finance (DeFi) ecosystem is rapidly evolving, and the distinctions between Centralized Exchanges (CEX) and DEX are becoming increasingly important. CEXs typically offer faster trades because they process transactions off-chain, which also leads to lower transaction fees. In contrast, DEXs rely on on-chain transactions, which, while providing greater transparency, result in slower and more expensive trades due to blockchain fees such as gas costs. These factors contribute to higher slippage on DEXs, especially during periods of high volatility or low liquidity.

One of the most significant differences between CEXs and DEXs is the automation of market-making. CEXs rely on various market makers who supply liquidity, ensuring that trades can be executed continuously. These market makers actively manage order books to match buyers and sellers. In contrast, DEXs face technical challenges in implementing traditional market makers due to the decentralized nature of blockchain technology. To overcome this, DEXs have adopted Automated Market Makers (AMMs), which are designed to enable liquidity provision without external market makers. AMMs, through their algorithmic pricing models, allow users to trade directly with liquidity pools, ensuring decentralized and continuous trading.

Many DEXs implement AMMs as a core mechanism, where trades are facilitated without the need for conventional order books. Most AMM-based DEXs use Constant Function Market Makers (CFMMs), which calculate asset prices based solely on the ratio of assets within a liquidity pool. This fixed pricing mechanism simplifies the trading process by eliminating the need for complex algorithms or external price oracles, streamlining price discovery and increasing transparency. Uniswap's implementation of a constant function AMM has garnered significant attention for its simplicity and resistance to certain forms of market manipulation. However, this model is not without its challenges. Issues such as impermanent loss and suboptimal pricing during extreme market conditions remain areas of concern, highlighting the need for further research and optimization.

These distinctions between CEXs and DEXs, especially regarding slippage and market-making mechanisms, provide essential context for understanding the vulnerabilities and trade-offs in decentralized markets. Our research emphasizes these dynamics, particularly the implications of slippage, as a key factor in sandwich attacks, further contributing to the broader discourse on the security and efficiency of decentralized exchanges.

Uniswap V2. In Uniswap V2 Adams et al. (2020), we denote the amount of tokens X and Y reserved in the liquidity pool as x and y respectively, and the overall liquidity of the pool as L. Assuming that there is no additional liquidity provision, the amount of tokens in the pool should always follow the CFMM formula as follows:

$$x \cdot y = L^2 \tag{13}$$

The marginal price of the asset Y with respect to X at time t can be computed as :

$$p_t = -\frac{\partial x_t}{\partial y_t} = \frac{L_t^2}{y_t^2} = \frac{x_t}{y_t} \tag{14}$$

A token swap within the liquidity pool triggers a state change in the token reserves. However, the pool must consistently adhere to the invariant formula expressed in Equation 13. Consequently, when there is an alteration in quantities denoted by Δx and $-\Delta y$, the resulting amounts $(x + \Delta x)$ and $(y - \Delta y)$ must still satisfy:

$$(x + \Delta x) \cdot (y - \Delta y) = L^2 \tag{15}$$

Uniswap V3 & Concentrated Liquidity. Uniswap V3 Adams et al. (2021), represented a significant advancement beyond Uniswap V2 by introducing the concept of concentrated liquidity, deviating from the equal distribution of liquidity across the entire price spectrum seen in Uniswap V2. Formally, liquidity providers in Uniswap V3 possess the capability to concentrate their liquidity within specific price ranges, referred to as "ticks," rather thanuniformly providing liquidity across the entire price curve. This design allows liquidity providers to strategically target specific price ranges where they anticipate more favorable trading opportunities or reduced impermanent loss. The introduction of concentrated liquidity in Uniswap V3 aims to enhance capital efficiency, providing liquidity providers with more nuanced control over their assets within the trading range. This innovation reflects a sophisticated approach to liquidity provision, catering to a broader spectrum of user preferences and risk profiles within the dynamic landscape of DeFi.

As each Uniswap V3 liquidity provider has unique liquidity positions characterized by distinct tick ranges, consider a liquidity position with liquidity L, the lower

price boundary p_l, and the upper price boundary p_u. In this context, the following equation should hold:

$$(x + \frac{L}{\sqrt{p_u}})(y + L\sqrt{p_l}) = L^2 \tag{16}$$

The description of Uniswap V3 provided here is inherently localized, focusing on trade dynamics within specific price intervals. Integrating these local dynamics across all price points results in the formation of an aggregate reserve curve, governing trades across the entire spectrum of possible prices.

The swap mechanism employed in Uniswap V3 adheres CPMM model, which is in line with the approach employed in Uniswap V2. Assuming that the current price is P_c and a trader endeavors to input Δy of token Y and receive Δx of token X in return. We know the fact that when swapping within a price range, only P_c changes and the liquidity L remains unchanged. Then, we can find the post-swap price by using:

$$\Delta\sqrt{P} = \frac{\Delta y}{L} \tag{17}$$

As we know the input amount Δy, the post-swap price P_a is:

$$\sqrt{P_a} = \sqrt{P_c} + \frac{\Delta y}{L} \tag{18}$$

After calculating the post-swap price, we can calculate the token amounts by using the amount calculations functions:

$$\begin{aligned} x &= \frac{L(\sqrt{P_c} - \sqrt{P_a})}{\sqrt{P_c}\sqrt{P_a}} \\ y &= L(\sqrt{P_c} - \sqrt{P_a}) \end{aligned} \tag{19}$$

6. Conclusion

In this paper, we have explored the multifaceted nature of Maximal Extractable Value (MEV) in the DeFi ecosystem, with a specific focus on sandwich attacks on decentralized exchanges (DEXs). Our research contributes a novel game-theoretic model for determining optimal attack volumes, challenging traditional assumptions about attacker behavior in decentralized finance.

Our findings challenge the conventional belief that attackers consistently exploit the maximum slippage allowed by victims. Instead, our model reveals a more nuanced

strategy, where attackers must weigh transaction fees against the unpredictability of profits, largely influenced by the concept of *external slippage*. This balance reflects the need to consider the variability introduced by the presence of other transactions within the same block, which can significantly affect the profitability of sandwich attacks. The introduction of *external slippage* is a key contribution that helps in understanding how attackers strategically navigate the uncertain environment of blockchain state transitions.

Empirical analysis, based on Uniswap liquidity pool transaction data, provides strong support for our theoretical findings. We confirm that attackers do not always seek to maximize attack volumes but often opt for smaller-scale attacks to optimize risk versus reward. The evidence also shows that higher transaction priority fees correlate with earlier execution within a block, thereby influencing both the success probability and variability in profit. This alignment of theory and practical observation strengthens the robustness of our model.

By incorporating a behavioral perspective, we further demonstrate that market participants, particularly in the context of cryptocurrencies, often act in ways that deviate from purely rational behavior. Irrationalities and risk-seeking tendencies in the behavior of cryptocurrency investors play a critical role in shaping market dynamics. Our model, which integrates these behavioral elements, offers a more accurate explanation of real-world sandwich attack scenarios compared to existing purely rational models.

Overall, this paper not only enhances the theoretical understanding of sandwich attacks within the MEV framework but also provides practical insights that are relevant to policymakers, blockchain developers, and market participants interested in the integrity of DeFi systems. Future research could focus on defensive strategies against sandwich attacks, potentially involving improved transaction sequencing or cryptographic techniques to mitigate the effects of *external slippage*. Moreover, examining the broader impact of these attacks on market efficiency and the role of regulatory frameworks could yield valuable insights. Lastly, refining model assumptions, such as attack volumes and their influence on market conditions, will help bridge the gap between theoretical analysis and empirical observations. This work lays a strong foundation for further exploration of adversarial strategies in decentralized finance and the ongoing evolution of market mechanisms to address inherent vulnerabilities.

Acknowledgements

This work was supported by the faculty research fund of Sejong University in 2023, by Chung-Ang University in 2023, and by the National Research Foundation

of Korea (NRF) grants funded by the Korean government (MSIT: Ministry of Science and ICT) under grant number NRF-2022R1F1A1074008.

References

Adams, H., Zinsmeister, N., Robinson, D., 2020. Uniswap v2 core, 2020. URL: https://uniswap. org/whitepaper. pdf .

Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D., 2021. Uniswap v3 core. Tech. rep., Uniswap, Tech. Rep. .

Almeida, J., Gonçalves, T.C., 2023. A systematic literature review of investor behavior in the cryptocurrency markets. Journal of Behavioral and Experimental Finance 37, 100785.

Alpos, O., Amores-Sesar, I., Cachin, C., Yeo, M., 2023. Eating sandwiches: Modular and lightweight elimination of transaction reordering attacks. arXiv preprint arXiv:2307.02954 .

Baliker, C., Baza, M., Alourani, A., Alshehri, A., Alshahrani, H., Choo, K.K.R., 2023. On the applications of blockchain in fintech: advancements and opportunities. IEEE Transactions on Engineering Management .

Canidio, A., Danos, V., 2024. Commitment against front-running attacks. Management Science 70, 4429–4440.

Chung, K., Kim, K., Lim, N., 2020. Social structures and reputation in expert review systems. Management Science 66, 3249–3276.

Colombo, J.A., Yarovaya, L., 2024. Are crypto and non-crypto investors alike? evidence from a comprehensive survey in brazil. Technology in Society , 102468.

Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A., 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability, in: 2020 IEEE Symposium on Security and Privacy (SP), IEEE. pp. 910–927.

Ferreira, M.V., Parkes, D.C., 2022. Credible decentralized exchange design via verifiable sequencing rules. arXiv preprint arXiv:2209.15569 .

Hackethal, A., Hanspal, T., Lammer, D.M., Rink, K., 2022. The characteristics and portfolio behavior of bitcoin investors: Evidence from indirect cryptocurrency investments. Review of Finance 26, 855–898.

Heimbach, L., Wattenhofer, R., 2022. Eliminating sandwich attacks with the help of game theory, in: Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security, pp. 153–167.

Kim, K., Chung, K., Lim, N., 2019. Third-party reviews and quality provision. Management Science 65, 2695–2716.

Liu, S., Wang, C., Zhou, Y., 2022. Analysis of financial data risk and network information security by blockchain technology and edge computing. IEEE Transactions on Engineering Management .

Mohan, V., 2022. Automated market makers and decentralized exchanges: a defi primer. Financial Innovation 8, 20.

Nelaturu, K., Du, H., Le, D.P., 2022. A review of blockchain in fintech: taxonomy, challenges, and future directions. Cryptography 6, 18.

Park, S., Jeong, W., Lee, Y., Son, B., Jang, H., Lee, J., 2024. Unraveling the mev enigma: Abi-free detection model using graph neural networks. Future Generation Computer Systems 153, 70–83.

Perez, D., Werner, S.M., Xu, J., Livshits, B., 2021. Liquidations: Defi on a knife-edge, in: Financial Cryptography and Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part II 25, Springer. pp. 457–476.

Qin, K., Zhou, L., Gervais, A., 2022. Quantifying blockchain extractable value: How dark is the forest?, in: 2022 IEEE Symposium on Security and Privacy (SP), IEEE. pp. 198–214.

Song, Y., Sun, C., Li, L., Wei, F., Liu, Y., Sun, B., 2023. Research on blockchain-based fintech trust evaluation mechanism. IEEE Access .

Werner, S., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., Knottenbelt, W., 2022. Sok: Decentralized finance (defi), in: Proceedings of the 4th ACM Conference on Advances in Financial Technologies, pp. 30–46.

Zhou, L., Qin, K., Torres, C.F., Le, D.V., Gervais, A., 2021. High-frequency trading on decentralized on-chain exchanges, in: 2021 IEEE Symposium on Security and Privacy (SP), IEEE. pp. 428–445.

Züst, P., Nadahalli, T., Wattenhofer, Y.W.R., 2021. Analyzing and preventing sandwich attacks in ethereum. ETH Zürich .