Maximal Extractable Value in Batch Auctions

MENGQIAN ZHANG, Yale University, USA
YUHAO LI, Columbia University, USA
XINYUAN SUN, Flashbots, USA
ELYNN CHEN, New York University, USA
XI CHEN, New York University, USA

In the ever-evolving blockchain ecosystem, decentralized exchanges (DEXs) have seen significant growth, which, however, has also brought challenges of Maximal Extractable Value (MEV). DEXs offer a decentralized platform for cryptocurrency trading. Such trading mechanisms primarily include Constant Function Market Makers (CFMMs) and batch auctions.

We first examine MEV in batch auctions. By treating assets and transactions as goods and traders within a pure exchange market, batch auctions can be formulated as a linear market, allowing exchange rates and trading outcomes in a batch to be derived from the prices and allocation in its market equilibrium. This design ensures fairness and efficiency from the perspective of economics. Despite the general belief that batch auctions are less vulnerable to MEV due to uniform pricing and order independence of transactions, we highlight that the block builder's current ability to rearrange the batch content is certainly sufficient to extract MEV in batch auctions, where the strategic behavior is a novel *market equilibrium manipulation*. We further explore the computation and show that: When the transactions in a batch form a Fisher market, an optimal attack can be computed in polynomial time. When the transactions form a general Arrow-Debreu market, under the Unique Games Conjecture, obtaining a 50.01%-approximation is NP-hard.

In addition, we show it is NP-hard to compute a strategy that obtains even 0.01% of optimal MEV in CFMMs when the fraction of swap fee is any constant larger than zero (e.g., 0.3%). This holds even for almost *any* CFMM, including the simple CFMM with constant product function (e.g., Uniswap v2).

This paper provides a framework to study MEV in DEXs, and our results resolve the computation of optimal MEV therein, contributing valuable insights to the ongoing discourse on DEX security.

Additional Key Words and Phrases: Blockchain, Market Equilibrium Manipulation, MEV, Batch Auction

Manuscript submitted for review to the 26th ACM Conference on Economics & Computation (EC'25).

1 Introduction

Within the evolving landscape of blockchain technology, the decentralized finance (DeFi) space has witnessed unprecedented growth, with various financial services available on-chain. As a cornerstone of this movement, Decentralized Exchanges (DEXs) mark a significant shift in the paradigm of asset exchange within the blockchain ecosystem.

Automated Market Makers (AMMs) represent the predominant class of DEXs by replacing traditional order book mechanisms with liquidity pools. In this framework, traders interact directly with the pools to swap their tokens, and the exchange rate is algorithmically determined based on the pool's token reserves. Many AMMs utilize constant functions [27], notably the constant product formula popularized by Uniswap [1], and are therefore referred to as constant function market makers (CFMMs). To mitigate the risk of unfavorable price movements, users typically set a slippage tolerance, effectively establishing a minimum acceptable exchange rate for their transactions. Additionally, a small swap fee (say, 0.3%) is charged on each transaction to reward liquidity providers for their contributions.

Despite the huge advantages offered by DEXs, they also introduce the challenge of Maximal Extractable Value (MEV) [10], a phenomenon that underscores the potential for block builders (*e.g.*, miners, validators) to exploit their position by transaction insertion, deletion, and reordering for financial gain. In the context of AMMs, the execution order of transactions plays a crucial role, making them particularly susceptible to MEV exploitation, notably through front-running [10] and sandwich attacks [34]. Beyond degrading the user experience, these unexpected MEV behaviors have also caused broader systemic issues, such as network congestion, high gas prices, and consensus instability [10, 33].

To address these vulnerabilities, the novel exploration of *batch auctions* offers a promising DEX design. Instead of processing transactions individually, DEXs utilizing batch auctions aggregate orders over a short period and settle them at a set of *uniform prices*. Here, "uniform" means that all successfully executed transactions in the same trading direction, say $\mathcal{X} \to \mathcal{Y}$, are under the same exchange rate $p_{\mathcal{X}}/p_{\mathcal{Y}}$. Then, the critical challenge of this approach lies in setting prices for all involved tokens. This problem turns out to be closely related to the *market equilibrium*. By modeling tokens and transactions respectively as goods and traders in a pure exchange market, each batch forms a linear market where each player's utility function is linear, and the unique price vector under market equilibrium is exactly what batch auctions need [25]. This feature of uniform exchange rates eliminates the execution sequence in batch auctions and makes such DEXs resistant to front-running, sandwich attacks, and internal arbitrage. While the batch auction design appears fundamentally less vulnerable to MEV, as we are going to show, it still leaves enough non-trivial space for a block builder to exploit batch transactions and extract MEV.

In this paper, we investigate the MEV optimization problem in these two prominent DEXs: CFMMs and batch auctions. Specifically, we study the following natural questions: (i) How can a strategic player extract MEV from batch auctions? (ii) Can the player efficiently obtain the optimal MEV from orders in CFMMs or batch auctions? (iii) If achieving the "optimal MEV" is intractable, is there an efficient algorithm that always achieves a good approximation?

These questions are of natural interest to block builders/searchers/solvers in the MEV supply chain, as even marginal improvements in extraction efficiency can translate into significant financial gains in the current highly competitive environment [13, 31]. Moreover, a clear characterization of MEV extraction can equip DEX designers with essential insights to develop countermeasures against adversarial exploitation [22]. Beyond builders and DEX designers, our findings can also benefit the broader blockchain community, e.g., by informing policy debates and technical efforts at more equitable MEV redistribution [32].

1.1 Our Model and Contributions

This paper studies the MEV optimization problem in CFMMs and batch auctions. The analysis is under a unified framework: both scenarios share the same format of transactions, the same strategy space, and the same utility function.

The Model. We consider the exchange transactions among n tokens $\{\tau_1, \cdots, \tau_n\}$. In both scenarios, a transaction is trying to swap one type of token $\mathcal{X}$ for another $\mathcal{Y}$, represented in the format of $(\mathcal{X} \to \mathcal{Y}, \delta_{\mathcal{X}}, r)$ but denoted by SWAP and BATCH, respectively. In other words, each SWAP or BATCH transaction is composed of three components, *i.e.*, the exchange direction $\mathcal{X} \to \mathcal{Y}$, the number of tokens willing to sell $\delta_{\mathcal{X}}$, and the threshold for the exchange rate r which means the user should receive at least $\delta_{\mathcal{X}} \cdot r$ amount of token $\mathcal{Y}$ in return.

For a SWAP transaction, tokens are exchanged by a certain CFMM, which supports the exchange between two involved tokens. Without loss of generality, we explore the optimal MEV among the SWAP transactions between two tokens $\mathcal{X}, \mathcal{Y} \in \{\tau_1, \cdots, \tau_n\}$. Instead, in the batch auction scenario, all BATCH transactions involving multiple tokens can be executed in the same batch.

Regarding the strategy space and utility function, we follow the consensus of the community that MEV refers to the *additional* value that can be extracted from block production by *including, excluding, and reordering* transactions in a block. From now on, we use attacker/mediator interchangeably to refer to the role (*e.g.*, block builders, miners) who can extract MEV. Given a set of m user transactions $\left\{\text{SWAP}^i\right\}_{i\in[m]}$ or $\left\{\text{BATCH}^i\right\}_{i\in[m]}$, the attacker is able to insert some transactions of the same type, select a subset of user transactions, and compute an order for these selected transactions, which together form an MEV strategy. For the batch auction scenario, the attacker can ignore the last reordering step because the order has no influence on the execution outcome of transactions.

Under a set of user transactions, once given an MEV strategy, the transactions' outcomes as well as the attacker's profits are determined. In this paper, the utility of a strategy is measured by the overall value of the attacker's final token holdings. Note that the attacker's profits are all from its own transactions. What utility functions of both scenarios (formally defined in Equation (1) and Equation (7)) do is to enumerate all attacker's newly added transactions, and for each transaction, its utility is the value of tokens finally received minus the value of tokens it brought. Here, the value of a token is measured by its exogenous price, which represents the attacker's self-belief – it may be the price in another DEX, another domain, or even the off-chain information (*e.g.*, the price of tokens in a centralized exchange like Coinbase). Throughout this paper, we assume that the exogenous prices remain the same during the attack, which is around 12 seconds in Ethereum.

Our Contributions. We initiate the study of MEV on batch swaps. It is a widespread belief that batch auctions are fundamentally less vulnerable to MEV since the outcome doesn't depend on the order of transactions. However, we first observe that the ability to insert and delete transactions is already sufficient for the mediator to extract MEV (see Example 3.3 for a very concrete example)! By adopting the formulation of market equilibrium, such behavior is a *novel market equilibrium manipulation*, where a strategic player in the market has a very strong power – they can arbitrarily kick other participants out of the market and insert several fake identities. Although seemingly unreasonable, this is what could happen with block builders in the current blockchain system. This seeming unreasonableness exactly reflects the potential vulnerability of batch auctions, which, to the best of our knowledge, was not discussed before. One plausible reason is that while one *can* manipulate the batch auction in such a way, it is not clear *how* to manipulate since the outcome is not as easily predictive as that in AMMs. To this end, as our technical contributions, we discover many underlying combinatorial structures of the optimal attacks, the insights of which are going to be shown in the next subsection (Section 1.2).

As our second main contribution, we fully characterize the computational complexity of the MEV optimization problem in batch auctions based on the structure of market: When the transactions in a batch form a Fisher market, an optimal attack can be computed in polynomial time (Theorem 3.13); When the transactions form a general Arrow-Debreu market, it is NP-hard to find such one (Theorem 3.14). Furthermore, we strengthen the hardness result by showing that, in Arrow-Debreu markets, computing a strategy that achieves even 50.01% of the optimal revenue remains NP-hard, assuming the Unique Games Conjecture (a well-known conjecture in hardness of approximation) [16].

MEV issues in AMMs have received a lot of attention; see *e.g.* [3, 10, 17, 24, 33, 34]. From the computational aspect, previous work [3] shows a polynomial time algorithm that can obtain optimal MEV for a special case where the AMM has no swap fees. In sharp contrast, as our third contribution (Theorem 4.4), we show that it is NP-hard to compute a strategy that obtains even 0.01% of optimal MEV in CFMMs when the fraction of swap fee is any constant larger than zero (*e.g.*, 0.3%). This holds even for almost *any* CFMM, including the simplest CFMM with constant product function (i.e., Uniswap v2).

1.2 Overview of Insights in the Proofs

While optimizing MEV is computationally hard under both popular decentralized exchange scenarios (CFMMs and Batch auctions), the routines towards these two hardness results, in fact, provide many distinct insights, which we summarize below.

Batch auctions. Batch auctions have a more sophisticated global structure than AMMs, as they bring the concept of *equilibrium*. Thus, one small change of the batch (insert, delete, or modify a transaction) may result in dramatically different outcomes for every transaction, making it difficult to analyze compared to AMMs. Nevertheless, we observe many underlying combinatorial structures for optimal attacks and highlight the following two points, which we found insightful:

- In general, the directed graph of selected user transactions is acyclic (Lemma 3.8). The proof of this lemma follows a simple trick: If there was some cycle, then the mediator can locate one user transaction that can obtain profits, and replace it with their own transaction of exactly the same content. Despite being simple and intuitive, *this actually illustrates a kind of front-running in batch auctions!* The acyclic property also leads to identifying the hardness structure from the *Maximum Acyclic Subgraph* problem, which we reduce to the MEV optimization problem for the general Arrow-Debreu market.
- In general, the directed graph of the mediator's transactions is acyclic (Lemma 3.9), and the transactions are inserted along the edges of the undirected economy graph of the initial users' transactions (Lemma 3.10). In other words, the mediator is never necessary to complicate the market structure. In particular, when the mediator is attacking a set of user transactions that form a Fisher market, it suffices to consider attacks that remain a simple Fisher market, which is crucial to our efficient optimal algorithm for Fisher markets.

Note that all lemmas above assert structural statements that can be applied to the analysis of both optimal and approximately optimal attack strategies. Thus they are also crucial for us to obtain the 50.01%-inapproximability result.

In fact, we couldn't come up with a polynomial time approximation algorithm that achieves even 1% approximation guarantee. Thus, we conjecture even a 1%-approximation algorithm for that in batch auctions is computationally hard. However, the Maximum Acyclic Subgraph problem has a very simple 50%-approximation algorithm. Thus to strengthen the hardness results, new ideas to exploit the structure of batch auctions are needed.

CFMMs. Let's first get intuition on why the mediator could get some MEV in CFMMs and then explain why it is hard to obtain a good approximation of MEV. Imagine that at the latest state of a CFMM pool, the exchange rate of two tokens is exactly the same as the ratio of their exogenous prices. Then, the mediator could execute an arbitrary user transaction, after which the exchange rate in the CFMM must deviate. Thus, this leaves a space for the mediator to back-run and obtain some profits.

This is an ideal argument which, however, is not always true when there is a constant fraction of swap fees (in particular, $f = 0.3\%$ of tokens is charged in common Uniswap pools). Specifically, when the volume of a user's transaction is small, the profits obtained by back-running may not be able to beat the swap fees! Thus, we should adapt our intuition to try to back-run some large user transactions (*i.e.*, transactions that want to swap a large amount of tokens), and there is where another constraint comes in: the slippage tolerance of these transactions.

Suppose that there is a transaction Swap that wants to swap a relatively large number of $\mathcal{Y}$ tokens for some $\mathcal{X}$ tokens, but with a non-trivial slippage tolerance requirement. The mediator would meet the following challenge to finish the best back-running: The mediator would like to find a set of users' $\mathcal{X} \to \mathcal{Y}$ transactions to reach a state at which the exchange rate of Swap exceeds but is closest to its slippage tolerance. Thus, the NP-hard problem used in the reduction is the *Partition problem*, which exactly reflects the hardness of achieving the goal above. We will provide more intuition about this argument in Section 4.3 before the formal proof of Theorem 4.4.

1.3 Organization

The rest of the paper is organized as follows. In Section 2, we provide more background on the MEV, its concrete behaviors in AMMs, and the design of batch auctions. In Section 3, we first provide a light preliminary of market equilibrium and formulate batch auctions in this framework. We then argue that the block builder's ability to manipulate the block content is already sufficient to extract MEV and formulate the optimal MEV problem. Then, we provide a few structural lemmas that reveal combinatorial structures in attacks. They are crucial for us to prove our computational results regarding the Fisher market and the general Arrow-Debreu market. Finally, in Section 4, we provide the formulation of the MEV optimization problem in AMMs and show its (any) constant inapproximability. We discuss some future directions in Section 5.

2 Background and Related Work

2.1 Maximal Extractable Value

Although DEXs allow users to directly interact with on-chain smart contracts through a trading transaction when they want to exchange cryptocurrencies, such a transaction only represents the individual's trading intent. The trading is truly executed when the transaction is included in a block on the canonical chain, which is managed by miners (in Proof-of-Work networks) or validators (in Proof-of-Stake networks). In the block-building process, they have the authority to decide which transactions are included in a block and in what order. It's found that block builders are able to extract additional value from block production in excess of the standard block rewards and gas fees by manipulating the block content. This additional value was initially referred to as "miner extractable value" and modified to be "maximal extractable value" since the transition from proof-of-work to proof-of-stake via The Merge. In practice, besides block builders, a large portion of MEV is extracted by independent network participants called "searchers." As the name suggests, they run complex algorithms to search profitable MEV opportunities and have bots to submit MEV-capture transactions automatically. The studies of MEV strategies naturally interest these roles, which are collectively referred to as the attacker or mediator in this paper.

2.2 MEV in AMMs

In the context of AMMs, the MEV phenomenon arises primarily due to the arbitrage opportunities caused by the price movement within a DEX or price discrepancy among DEXs. *DEX arbitrage* is the simplest and most well-known MEV opportunity. Zhou *et al.* [33] translate the detection of DEXs arbitrage into a negative cycle detection problem, and Wang *et al.* [30] further analyze the profitability conditions and optimal trading strategies of cyclic arbitrages among multiple DEXs. After detecting such a profitable MEV opportunity, someone can submit their own transaction with the same arbitrage logic but a higher gas price to steal the profit by *front-running*. Torres *et al.* [29] perform a large-scale analysis of the real profits made by front-running attacks on Ethereum, providing evidence that front-running is highly lucrative. Qin *et al.* [24] provide a generalized transaction replay algorithm to clone and front-run a victim transaction without the need to understand the underlying transaction logic. *Sandwich attack* [34] is another common MEV behavior where a trader can "manually" create the arbitrage opportunity by exploiting a large order, the profit of which is quantified by Qin *et al.* [24]. Many other behaviors like back-run arbitrage [19] and cross-domain MEV [23] are also discussed in the literature.

The most related to us are the following papers. Bartoletti *et al.* [3] explore the same MEV optimization problem in AMMs but without swap fees, where the optimal attack is a multi-layer sandwich. [15] considers the practical scenario with swap fees and computes the optimal strategy to sandwich a single transaction. Our work fills this research gap by providing the computational hardness of attacking multiple transactions in AMMs with swap fees. Another related work is [18], which studies both with and without fees in a special AMM with greedy sequencing rule.

2.3 Batch Auctions

In traditional finance, batching trades has been proposed as a solution to the high-frequency trading arms race [4]. In the cryptocurrency domain, batch auctions have emerged as an attractive approach to mitigate inefficiencies and MEV challenges inherent in AMMs, with several systems already proposed and deployed. For instance, in SPEEDEX [25], the batch execution is triggered when the core SPEEDEX engine receives a new block, followed by an algorithm query to compute clearing valuations and, thus, the execution outcomes. Similarly, CoWSwap [28], which uses mixed-integer programming to clear offers in batches, also contains two steps: a centralized entity known as the "driver" aggregates all user orders; these orders are then relayed to specialized centralized actors termed "solvers", who compete to solve an optimization problem, and as a result, determine the trading price (and over half [11] of non-stablecoin CoWSwap orders are traded against private liquidity). As we later demonstrate, organizing batches through a mediator brings the risk of MEV.

This paper mainly follows the mathematical model in SPEEDEX [25], which maps the computation of batch prices to a well-studied equilibrium computation problem of pure exchange markets (details are introduced later in Section 3.1). In their paper, the authors mentioned a conceptually similar MEV risk, remarking that "one might estimate the clearing prices in a future batch and arbitrage the batch against low-latency markets." To the best of our knowledge, we are the first to *formally* study MEV in batch auctions. Although our analysis focuses on the market equilibrium model, the resulting insights — especially regarding MEV strategies — have broader applicability.

A recent direction of work [5–7, 26] has also explored combining batch auctions with CFMMs, which is orthogonal to the subject of study of this paper.

3 MEV in Batch Auctions

Batch auction is a trading mechanism that computes a uniform exchange price and executes all transactions simultaneously. So the prices and outcome of the execution do not depend on the order of the transactions. Because of this, batch auctions are believed to be fundamentally less vulnerable to MEV compared to order-dependent mechanisms (such as constant function automated market makers).

However, even though batch auctions are robust against front-running and sandwich attacks discussed above, we observe new strategy space for a mediator to obtain MEV by manipulating the block content. In the rest of this section, we first introduce the concept of market equilibrium for efficiency of the batch price, which was also introduced in the design of SPEEDEX as the mathematical foundation [25]. Then we *initial and formalize* the study of MEV therein.

3.1 Market Equilibrium Formulation of Batch Auctions

Pure Exchange Market. A pure exchange market consists of n divisible goods, denoted by $\{\tau_1, \cdots, \tau_n\}$ and m traders, denoted by $\{T_1, \cdots, T_m\}$. Every trader T_i initially owns some endowment $\mathbf{w}_i \in \mathbb{R}^n_{\geq 0}$, where w_{ij} represents the amount of τ_j that T_i owns. Every trader i has its own utility function $\mathcal{U}_i : \mathbb{R}^n_{\geq 0} \to \mathbb{R}_{\geq 0}$, where $\mathcal{U}_i(\mathbf{x}_i)$ means T_i's utility if she gets a bundle of goods with amount x_{ij} for τ_j. In this paper, we focus on linear market, meaning that the utility function $\mathcal{U}_i = \sum_j u_{ij} x_{ij}$, where $u_{ij} \geq 0$ is the utility (or say preference) of trader i for a unit amount of goods τ_j. In economics, it has long been understood that prices are determined by the interplay of supply and demand where under a *competitive price*, supply precisely meets demand. Such an idea is captured by *market equilibrium.*

Definition 3.1 (Market Equilibrium). A price vector $\mathbf{p}$ along with an allocation $\mathbf{x}$ is a market equilibrium if the following conditions meet:

- Market Clearance: $\sum_{i \in [m]} x_{ij} = \sum_{i \in [m]} w_{ij}$ for all $j \in [n]$;
- Budget Constraint: $\sum_{j \in [n]} x_{ij} p_j = \sum_{j \in [n]} w_{ij} p_j$ for all $i \in [m]$; and
- Individual Optimality: $\mathbf{x}_i \in \mathtt{OPT}_i(\mathbf{p})$, where $\mathtt{OPT}_i \coloneqq \arg\max \mathcal{U}_i(\mathbf{x}_i)$ among all $\mathbf{x}_i$ that satisfies the budget constrain $\sum_{j \in [n]} x_{ij} p_j \leq \sum_{j \in [n]} w_{ij} p_j$.

The existence of market equilibrium (sometimes called general equilibrium) under mild conditions was proved by Arrow and Debreu [2] and independently by McKenzie [21], which is regarded by many as the crown jewel of Mathematical Economics. Our world is much easier to navigate: We will show that the batch prices can be computed by a linear market equilibrium.

Consider a set of Batch transactions $\left\{\text{Batch}^1, \cdots, \text{Batch}^m\right\}$ among tokens $\{\tau_1, \cdots, \tau_n\}$, where each $\text{Batch}^i = (\mathcal{X}^i \to \mathcal{Y}^i, \delta_{\mathcal{X}^i}, r_i)$ would like to swap $\delta_{\mathcal{X}^i}$ amount of token $\mathcal{X}^i$ for some token $\mathcal{Y}^i$ and the exchange rate between $\mathcal{Y}^i$ and $\mathcal{X}^i$ should be at least r_i (*i.e.*, the amount of received $\mathcal{Y}^i$ should be no less than $r_i \delta_{\mathcal{X}^i}$). Here, tokens naturally correspond to the goods in the market. Every transaction Batch^i can be viewed as a trader, where its endowment is $w_{i,j} = \delta_{\mathcal{X}^i}$ for $\tau_j = \mathcal{X}^i$ and $w_{i,j} = 0$ otherwise. Its utility function is defined as $\mathcal{U}_i(\mathbf{x}_i) = r_i \cdot x_{ij} + x_{ik}$ for $\tau_j = \mathcal{X}^i$ and $\tau_k = \mathcal{Y}^i$. Given any set of Batch transactions, we refer to the corresponding market as an Arrow-Debreu market if there is no more specified structure (this is mainly for comparison of the Fisher market that we will define below).

The market constructed above is a linear market in which every participant's utility function is linear (in fact, it is even more succinct – only two goods have non-zero coefficients). The linear market enjoys many nice properties: First, previous work [12] implies that our batch auction structure always has a *unique* competitive price vector[1], so that we do not face the equilibrium selection problem. In the distributed system where the BATCH transactions are eventually processed by an arbitrary node, it releases us from the concern that certain nodes may strategically decide the exchange rates. Second, it is polynomial-time computable, while for some other classic markets like Leontief or CES (stands for constant elasticity of substitution) functions, the computation of market equilibrium is computationally hard (PPAD-complete [8, 9]). Third, it is generally well understood from the works of past decades by researchers from economics, computer science, and operation research.

PROPOSITION 3.2 ([25]). *By modeling the tokens $\{\tau_j\}_{j\in[n]}$ and transactions $\{\text{BATCH}^i\}_{i\in[m]}$ as goods and traders, under the (unique) competitive price vector* $\mathbf{p}$ *and arbitrary equilibrium allocation* $\mathbf{x}$*, it satisfies*

- *Internal arbitrage-free: The prices are internal arbitrage-free (simply because we have a uniform price vector);*
- *Soundness: Any exchange (allocation) follows the price vector* $\mathbf{p}$ *and meets its threshold requirement for the exchange ratio;*
- *Completeness: For any* BATCH^i *with its threshold* r_i*,* $\tau_j = \mathcal{X}^i$ *and* $\tau_k = \mathcal{Y}^i$*, if* $p_j/p_k > r_i$*, then* BATCH^i *sells all its* $\mathcal{X}^i$ *and gets* $\delta_{\mathcal{X}^i} \cdot p_j/p_k$ *amount of* $\mathcal{Y}^i$*.*

3.2 MEV Optimization Problem

The feature of uniform price in batch auctions makes it resistant to several DEX-related MEV behaviors like sandwich attacks and internal arbitrage. As a result, MEV seems impossible in batch auctions. We first note that this is not the case – we observe new strategic behavior, which is highlevelly in line with block content manipulation, but different from well-known MEV strategies like front-running or sandwich attack. The point is the execution result of a BATCH transaction depends on which other transactions are included in the batch. Therefore, a strategic mediator can affect the outcome of exchanges by manipulating the batch contents (inserting and deleting transactions). The mediator still respects the market equilibrium outcomes, but that of the *manipulated block content*, during which it may gain profits. We give an example to provide more intuition.

Example 3.3 (MEV in Batch Auctions). Consider the scenario with three tokens $\{A, B, C\}$ and three user transactions $\{\text{BATCH}^i\}_{i\in[1,3]}$ where $\text{BATCH}^1 = (A \to B, 2, 0.5)$, $\text{BATCH}^2 = (B \to C, 1, 4)$, and $\text{BATCH}^3 = (C \to A, 4, 0.5)$. Assume the exogenous prices of tokens are all one, namely, $p^*_A = p^*_B = p^*_C = 1$. Given this batch of transactions, an honest mediator makes them exchange with each other (i.e., three users receive $1B, 4C, 2A$, respectively) under the price equilibrium $p_A = 0.5, p_B = 1, p_C = 0.25$. Nevertheless, a strategic mediator can extract some additional value by re-organizing the batch. One method is to exclude BATCH^2 and insert two attacking transactions $\{\text{BATCH}^4 = (B \to A, 1, 2),\ \text{BATCH}^5 = (A \to C, 2, 2)\}$. By executing them in the batch $\{\text{BATCH}^i\}_{i\in\{1,3,4,5\}}$, the attacker receives $2A$ and $4C$ at the cost of $1B$ and $2A$, obtaining a net benefit of 3 (recall that their exogenous prices are all 1). Another way is to directly replace the BATCH^2 with the same attacking transaction $\text{BATCH}^{2'} = (B \to C, 1, 4)$. In this way, the attacker can also get a profit of 3.

Next, we give formal definitions of an attacker's strategy space and utility function.

[1]The uniqueness is in terms of scaling, which means scaling the equilibrium price vector by a constant is also an equilibrium. It does not matter because what is really needed is the ratio of prices, which remains the same no matter how it is scaled.

Definition 3.4 (Strategy Space). Given a batch of user transactions $\{\textsc{Batch}^i\}_{i\in[m]}$, a mediator could select a subset of all these transactions $S \subseteq [m]$, create k its own Batch transactions $\{\textsc{Batch}^i\}_{i\in[m+1:m+k]}$, and execute them in a same batch.

Without loss of generality, we assume that all mediator's newly inserted transactions will be successfully executed; otherwise, removing failed ones from the set $\{\textsc{Batch}^i\}_{i\in[m+1:m+k]}$ makes no difference in results.

Definition 3.5 (Utility Function). Mediator's profit $U(S, \{\textsc{Batch}^i\}_{i\in[m+1:m+k]})$ is defined as

$$\max_{\mathbf{x}\in\mathbf{X}} \sum_{i\in[m+1:m+k]} -w_{ij}\cdot p_j^* + x_{i\ell}\cdot p_\ell^*, \tag{1}$$

where $\textsc{Batch}^i = (\tau_j \rightarrow \tau_\ell, \delta_{\tau_j} = w_{ij}, r_i)$, $\mathbf{X}$ is the set of feasible allocations under equilibrium, and vector $\mathbf{p}^*$ represents the exogenous prices of tokens.

Remark. Note that under market equilibrium, there may exist multiple equilibrium allocations although the price vector is unique. For instance, consider the batch of user transactions $\{\textsc{Batch}^i\}_{i\in[1,3]}$ in Example 3.3. All execution and none execution (as well as partial executions) are all equilibrium allocations. This is not a problem in an "honest" market where everyone reports their true preferences because each equilibrium allocation maximizes their utility, which is defined according to the reported preferences (i.e., the pair of trading tokens $\mathcal{X}^i$ and $\mathcal{Y}^i$ and the lowest acceptable exchange rate r_i). However, for the mediator who may report fake preferences to manipulate the market, this is not the case; namely, maximizing the utility of newly inserted transactions is not equal to the maximization of the mediator's utility, which is defined based on its true belief p^*. Considering that different equilibrium allocations $\mathbf{x} \in \mathbf{X}$ may bring different profits, we (optimistically) define the utility as the maximal one as shown in Equation (1). There may be other reasonable definitions, say, the minimal one (thus pessimistically). In fact, this relates to the equilibrium selection problem in the batch auction design, which is an interesting topic but not the focus of this paper. More importantly, our following analysis is independent of the specific selection rule (say, "max", "min", or others). Intuitively, multiple (different) allocations under equilibrium only happen to transactions satisfying that their specified exchange rates are exactly the same as the ratio of equilibrium prices. However, it is easy to bypass this case and guarantee a unique allocation by perturbing the mediator's attacking strategy with an arbitrarily small loss on the utilities.

Definition 3.6 (Batch-MEV). The Batch-MEV problem refers to the following computational problem: Given user transactions $\{\textsc{Batch}^i\}_{i\in[m]}$, compute a strategy $(S, \{\textsc{Batch}^i\}_{i\in[m+1:m+k]})$ that achieves the optimal revenue. For $c \in (0, 1)$, we say a strategy is a c-approximation to Batch-MEV if the strategy obtains at least c-fraction of the optimal MEV revenue.

This section asks the following research question: Given a batch of user transactions, how to compute an (approximately) optimal strategy for a mediator? Based on the structure of the given user transactions, we discuss this problem for the Fisher Market (Section 3.4) and the general Arrow-Debreu market (Section 3.5), respectively. We show that in the Fisher Markets, an optimal strategy can be found in polynomial time (Theorem 3.13). However, in general, Arrow-Debreu markets, we show that it is NP-hard to compute an optimal strategy (Theorem 3.14). Furthermore, we strengthen the hardness result by showing that, in Arrow-Debreu markets, computing a strategy that achieves even 50.01% of the optimal revenue is computationally hard (Theorem 3.17), assuming the Unique Games Conjecture (a well-known conjecture in hardness of approximation) [16].

3.3 Combinatorial Structures in the Batch-MEV problem

In this section, we present a few general lemmas that provide more intuition about how to tackle the MEV Optimization Problem in Batch auctions, and are very useful to the analysis of the computational results. Before that, from a technical perspective, we recall the definition of the economy graph of a market, which was defined by Maxfield [20]. In our context, we will use the terminology of transactions, and define both the directed and undirected versions of economy graphs for our purpose.

Definition 3.7 (Economy Graph). Given a set of transactions $\left\{\text{Batch}^i\right\}_{i\in[m]}$, we define a directed graph as follows. Each vertex corresponds to a token τ_i. For two tokens τ_i and τ_j, we add a directed edge from τ_i to τ_j if there is a transaction Batch^k that swaps τ_i for τ_j. We call this (directed) graph G the economy graph of $\left\{\text{Batch}^i\right\}_{i\in[m]}$.

The undirected economy graph H is defined as the same graph as the directed economy graph G but all edges are changed to undirected.

Intuitively, if there is an undirected edge between two tokens in H, this means there is a user who is interested in these two tokens (would like to swap one token for the other).

Now we are ready to introduce the lemmas.

At a high level, the first lemma works on the side of selecting users' transactions, showing that for an attack, it suffices to select users' transactions such that the directed economy graph of them is acyclic. The second lemma and third lemma concern the side of inserting attacking transactions, stating that for an attack, it is not necessary to make the *directed* economy graph more complicated. More precisely, it suffices to only insert transactions such that the following two conditions are met: (1) the directed economy graph of the inserted transactions is acyclic, and (2) the transactions are inserted along the edges of the undirected economy graph of the initial users' transactions $\left\{\text{Batch}^i\right\}_{i\in[m]}$. These lemmas bring many insights for the mediator to search for an (approximately) optimal attack and are crucial for both our efficient algorithm design, and NP-hard analysis, and the hardness of approximation analysis in subsequent sections.

Lemma 3.8. *Let $H = (V, E)$ be the undirected economy graph of users' transactions $\left\{\text{Batch}^i\right\}_{i\in[m]}$. It is sufficient to consider only attacks of the form $(S, \left\{\text{Batch}^i\right\}_{i\in[m+1:m+k]})$ such that the economy graph of selected users' transactions $\left\{\text{Batch}^i\right\}_{i\in S}$ is acyclic.*

Proof. Without loss of generality, assume that all transactions in this batch, namely, all transactions $\left\{\text{Batch}^i\right\}_{i\in S} \cup \left\{\text{Batch}^i\right\}_{i\in[m+1:m+k]}$ are successfully executed. Otherwise the mediator can remove all transactions that are not executed and the profit will be the same.

Suppose there is a directed cycle in the economy graph of selected users' transactions. Without loss of generality, assume the cycle is $\left(\text{Batch}^1, \text{Batch}^2, \cdots, \text{Batch}^t\right)$ with a vertex sequence $(\tau_1, \tau_2, \cdots, \tau_t, \tau_1)$, where $t \in [2, m]$. Then, under equilibrium, there are only two cases.

- Case 1: the exchange rates of all transactions in the cycle are exactly the same as the ratio of corresponding tokens' exogenous prices. In this case, replacing one of the user's transactions with an attacking transaction of the same content has no impact on the equilibrium nor the mediator's profit (as the profit of this replacement transaction is 0).
- Case 2: the exchange rates of transactions in the cycle are not all the same as the ratio of corresponding tokens' exogenous prices. Note that $\prod_{j=1}^{t-1} \frac{p_{\tau_j}}{p_{\tau_{j+1}}} \cdot \frac{p_{\tau_t}}{p_{\tau_1}} = 1 = \prod_{j=1}^{t-1} \frac{p^*_{\tau_j}}{p^*_{\tau_{j+1}}} \cdot \frac{p^*_{\tau_t}}{p^*_{\tau_1}}$, where $\frac{p_{\tau_j}}{p_{\tau_{j+1}}}$ is the equilibrium exchange rate of the transaction swapping τ_j for τ_{j+1} and $\frac{p^*_{\tau_j}}{p^*_{\tau_{j+1}}}$ is the exogenous rate of these two tokens. In this case, at least one transaction's equilibrium

exchange rate is better than the exogenous one, say Batch^j with the direction $\tau_j \to \tau_{j+1}$. In other words, $p_j/p_{j+1} > p_j^*/p_{j+1}^*$. Then, replacing *this user's transaction* by an attacking transaction with the same content (namely, remove this user's transaction from the batch and insert a mediator's transaction with the same content) does not influence the equilibrium but increases the mediator's utility.

In this way, for any attacking strategy with a directed cycle in the economy graph of selected user transactions, there is another strategy with no such cycle, such that the attacker receives no less profit than the previous one.

This finishes the proof. □

Lemma 3.9. *Let $H = (V, E)$ be the undirected economy graph of users' transactions $\{\text{Batch}^i\}_{i\in[m]}$. It is sufficient to consider only attacks of the form $(S, \{\text{Batch}^i\}_{i\in[m+1:m+k]})$ such that the economy graph of mediator's transactions $\{\text{Batch}^i\}_{i\in[m+1:m+k]}$ is acyclic.*

Proof. Without loss of generality, assume that all transactions in this batch, namely, all transactions $\{\text{Batch}^i\}_{i\in S} \cup \{\text{Batch}^i\}_{i\in[m+1:m+k]}$ are successfully executed. Otherwise the mediator can remove all transactions that are not executed and the profit will be the same.

Suppose there is a directed cycle in the economy graph of the mediator's transactions. Without loss of generality, assume the cycle is $(\text{Batch}^{m+1}, \text{Batch}^{m+2}, \cdots, \text{Batch}^{m+t})$ with a vertex sequence $(\tau_1, \tau_2, \cdots, \tau_t, \tau_1)$, where $t \in \{2, \cdots, k\}$. Note that these transactions are executed at a uniform clearing price. Thus, the mediator is cyclically exchanging τ_1 tokens to τ_2 tokens, τ_2 tokens to τ_3 tokens, and so on, ultimately converting τ_t tokens to τ_1 tokens. Since this forms a redundant cycle, we can proportionally reduce all transactions in the cycle by a common amount, allowing the "smallest" transaction to be eliminated without affecting the mediator's profit.

By this procedure we can decrease the number of mediator's transactions by one. Repeating this procedure, we can make sure that the economy graph of the mediator's transactions is acyclic without affecting the mediator's profit. □

Before presenting the last lemma, we first show the following intuitive claim for convenience.

Claim 1. Let $H = (V, E)$ be the undirected economy graph of users' transactions $\{\text{Batch}^i\}_{i\in[m]}$. If $H = (V, E)$ is not connected, then the mediator can attack each connected component separately.

Proof. Suppose that the connected components are $V_1, \cdots, V_t$, where $V_1 \cup \cdots \cup V_t = V$. Fix an attack $\texttt{att} = (S, \{\text{Batch}^i\}_{i\in[m+1:m+k]})$. By Lemma 3.9, we know that the directed economy graph of $\{\text{Batch}^i\}_{i\in[m+1:m+k]}$ is acyclic. Thus, in the new market formed by att, the vertex set of every strongly connected component must be contained in V_i for some $i \in [t]$. Noting that the edges between strongly connected components are essentially irrelevant to the (unique) market equilibrium, we conclude that att is attacking each connected component V_i separately. So we conclude the claim. □

Lemma 3.10. *Let $H = (V, E)$ be the undirected economy graph of users' transactions $\{\text{Batch}^i\}_{i\in[m]}$. It is sufficient to consider only attacks of the form $(S, \{\text{Batch}^i\}_{i\in[m+1:m+k]})$ such that the undirected economy graph $H' = (V, E')$ of $\{\text{Batch}^i\}_{i\in[m+1:m+k]}$ satisfies $E' \subseteq E$.*

Proof. Assume from Claim 1 that the original undirected economy graph $H = (V, E)$ is connected. Given any general strategy $(S, \{\text{Batch}^i\}_{i\in[m+1:m+k]})$, we will construct a new strategy such that (1) the undirected economy graph of the attacking transactions in the new strategy is a subgraph of H; and at the same time (2) the mediator's profit is the same as using $(S, \{\text{Batch}^i\}_{i\in[m+1:m+k]})$.

By doing so, we can conclude this lemma. In particular, the selected users' transactions are the same, so we focus on the attacking transactions next.

Consider any transaction $\text{Batch}^i = (\tau_j \longrightarrow \tau_{j'}, w_{ij}, r_i)$ for some $i \in [m+1 : m+k]$. Suppose that Batch^i is successfully executed and $(j, j') \notin E$. Note that if Batch^i is not successfully executed, we can simply remove it from the strategy and everything will remain unchanged.

We will replace the transaction Batch^i with a set of new transactions such that the edges that correspond to new transactions are a subset of E, and the mediator's profit remains unchanged.

Since H is connected, there is at least one simple path from j to j'. Denote such path by $v_0, \cdots, v_t$, where $v_0 = j$ and $v_t = j'$. Let $\mathbf{p}$ be the equilibrium price vector of $(S, \{\text{Batch}^i\}_{i \in [m+1:m+k]})$. We construct t new transactions based on these information as follows. For every $\ell \in [t]$, we let the new transaction be $(\tau_{v_{\ell-1}} \longrightarrow \tau_{v_\ell}, w_{ij} \cdot \frac{p_{v_{\ell-1}}}{p_0}, \frac{p_{v_\ell}}{p_{v_{\ell-1}}})$. Note that by construction all these new transactions are with respect to E, i.e., the edges that correspond to these newly constructed transactions are subset of E.

Now we show that the mediator's profit remains unchanged. Note that mediator brings $w_{ij} \cdot \frac{p_{v_{\ell-1}}}{p_0}$ more $\tau_{v_{\ell-1}}$ tokens. But it is easy to verify that under new market equilibrium, each transaction $(\tau_{v_{\ell-1}} \to \tau_{v_\ell})$ will be successfully executed, and obtain $w_{ij} \cdot \frac{p_{v_\ell}}{p_0}$ many taken τ_ℓ. In total, the payoff of middle transactions will be canceled and the mediator will pay w_{ij} many token τ_j and get $w_{ij} \cdot \frac{p_{v_t}}{p_0} = w_{ij} \cdot \frac{p_{j'}}{p_j}$ many token τ_j, which is exactly the same as before.

This finishes the proof. □

Remark. Note that the lemmas above assert structural statements that can be applied to the analysis of both optimal and approximately optimal attack strategies. We will use these lemmas throughout the remaining section.

3.4 Optimal MEV under Fisher Market

Our main result in this section is an efficient (in fact, almost linear-time) algorithm for the mediator to compute an optimal strategy for MEV under *Fisher market* structure, that is, where trades occur between τ_1 and τ_j for all $j \in [2 : n]$, with no trades taking place between τ_j and $\tau_{j'}$ for any $j, j' \in [2 : n]$. An interpretation of this model is to view τ_1 as USDC and every user's transaction is trying to trade between other cryptocurrencies with USDC.

We first give two lemmas to prepare for our optimal MEV algorithm.

The first lemma says that the mediator can independently attack the user transactions that occur between τ_1 and τ_j for all $j \in [2 : n]$.

Lemma 3.11. *For every $j \in [2 : n]$, let $S_{1\leftrightarrow j} \subseteq [m]$ be the indices of all users' transactions that trade between τ_1 and τ_j, and ATT_j be an optimal strategy to attack $\{\text{Batch}^i\}_{i \in S_{1\leftrightarrow j}}$. Then, $\cup_{j \in [2:n]} \mathsf{ATT}_j$ is an optimal attack to $\{\text{Batch}^i\}_{i \in [m]}$, where a union of two strategy is defined as the union of selected users' transactions and the union of inserted attacking transactions respectively.*

Proof. Let $SS_{1\to j}$ denote the set of successfully executed transactions swapping τ_1 for τ_j in the final batch, while $SS_{j\to 1}$ denote the set of successfully executed transactions that bring token τ_j and get τ_1. Recall that under a market equilibrium, the market clearance condition holds. It implies that for any token τ_j where $j \in [2 : n]$, we have

$$\sum_{\text{Batch}^i \in SS_{1\to j}} x_{ij} = \sum_{\text{Batch}^i \in SS_{j\to 1}} w_{ij}. \tag{2}$$

Let p_1 and p_j be the equilibrium prices of τ_1 and τ_j, respectively. The requirement that everyone spends their entire profit under the equilibrium implies that

$$\sum_{\text{Batch}^i \in SS_{1\to j}} x_{ij} p_j = \sum_{\text{Batch}^i \in SS_{1\to j}} w_{i1} p_1, \qquad \sum_{\text{Batch}^i \in SS_{j\to 1}} x_{i1} p_1 = \sum_{\text{Batch}^i \in SS_{j\to 1}} w_{ij} p_j. \tag{3}$$

Combining Equation (2) and Equation (3), we have

$$\sum_{\text{Batch}^i \in SS_{1\to j}} w_{i1} = \sum_{\text{Batch}^i \in SS_{j\to 1}} x_{i1}. \tag{4}$$

It means that under equilibrium, the consumed token τ_1 in the direction $\tau_j \to \tau_1$ are all from the transactions in the opposite direction; and vice versa for token τ_j according to Equation (2). In other words, each pair of tokens (τ_1, τ_j) in the Fisher market can be viewed as a sub-market where transactions between them are self-sufficient under equilibrium. Thus, the mediator is able to independently attack the user transactions in each sub-market, namely, to consider selecting which user transactions and inserting which attacking transactions. □

Now we are able to focus on each individual pair of tokens (τ_1, τ_j) for every $j \in [2 : n]$. The next lemma further simplifies the strategy space for optimal attacks therein.

Lemma 3.12. *For each pair (τ_1, τ_j) where $j \in [2 : n]$, there is an optimal attack which inserts transactions in at most one direction. Furthermore, all transactions in this direction are from the mediator.*

Proof. Without loss of generality, suppose the mediator inserts a transaction in each direction between τ_1 and τ_j where $j \in [2 : n]$, denoted by $\text{Batch}^{m+1} = (\tau_1 \to \tau_j, \delta_{\tau_1} = w_{m+1,1}, r_{m+1} = p_1/p_j)$ and $\text{Batch}^{m+2} = (\tau_j \to \tau_1, \delta_{\tau_j} = w_{m+2,j}, r_{m+2} = p_j/p_1)$. If $w_{m+1,1} \cdot p_1/p_j = w_{m+2,j}$ where p_1 and p_j are the tokens' equilibrium prices, Batch^{m+1} and Batch^{m+2} supply each other, bringing a utility of 0. Then, removing both transactions has no impact on the mediator's utility. Otherwise, let

$$\text{Batch}^{m+3} = \begin{cases} (\tau_1 \to \tau_j, \delta_{\tau_1} = w_{m+1,1} - w_{m+2,j} \cdot p_j/p_1, r_{m+3} = p_1/p_j), & w_{m+1,1} \cdot p_1/p_j > w_{m+2,j}; \\ (\tau_j \to \tau_1, \delta_{\tau_j} = w_{m+2,j} - w_{m+1,1} \cdot p_1/p_j, r_{m+3} = p_j/p_1), & w_{m+1,1} \cdot p_1/p_j < w_{m+2,j}. \end{cases} \tag{5}$$

Following the observation that Batch^{m+1} and Batch^{m+2} will be partially exchanged with each other, replacing them with Batch^{m+3} has no impact on the batch execution and mediator's utility. As a result, only one direction has the attack transaction. In fact, if $1 \to j$ is the direction, then the prices must satisfy $p_1/p_j \geq p_1^*/p_j^*$, and if $j \to 1$ is the direction, then the prices must satisfy $p_1/p_j \leq p_1^*/p_j^*$. Otherwise removing this attack transaction (and all user transactions between τ_1 and τ_j) from the batch increases the mediator's utility.

Then we show that all transaction in this direction are from the mediator. Under equilibrium, if there are user transactions successfully executed in a this direction, replacing them with attacking transactions of the same content is still an equilibrium but brings more profit for the mediator.

This concludes the proof. □

Now we are ready to describe our main algorithm and state our main theorem.

Theorem 3.13. *Given a set of users' transactions $\{\text{Batch}^i\}_{i\in[m]}$ such that they form a Fisher market (i.e., trades occur exclusively between τ_1 and τ_j for all $j \in [2 : n]$, with no trades taking place between τ_k and τ_j for any $k, j \in [2 : n]$), algorithm 1 finds a strategy that can obtain optimal MEV in time $\tilde{O}(m)$, where the notation $\tilde{O}(\cdot)$ hides polylogrithemic factors.*

ALGORITHM 1: Optimal MEV Strategy for BATCH Transactions that Form a Fisher Market

Input: A set of users' transactions $\{\text{BATCH}^i\}_{i\in[m]}$ that forms a Fisher market and a set of exogenous prices $\{p_i^*\}_{i\in[n]}$.

Output: A strategy for the mediator that obtains optimal profits.

Without loss of generality, assume that every transaction trades between τ_1 and τ_j for $j \in [2:n]$, *i.e.*, τ_1 is the special token.

for *each j from* 2 *to n* **do**

```
// Work on the direction τ1 → τj.
```

Let $S_{1\to j} \subseteq [m]$ be the set of indices i such that BATCH^i is in the direction $\tau_1 \to \tau_j$.

Sort transactions in $S_{1\to j}$ in an ascending order w.r.t. their exchange rate thresholds (break tie arbitrarily). Let π_1 be such an order and denote the k-th transaction in the order as $\text{BATCH}^{\pi_1(k)} = (\tau_1 \to \tau_j, \delta_{\tau_1}^{\pi_1(k)}, r^{\pi_1(k)})$.

Let $k_1 \in \arg\max_{k_1\in[|S_{1\to j}|]} \left\{\left(\sum_{k\in[k_1]} \delta_{\tau_1}^{\pi_1(k)}\right) \cdot \left(p_1^* - r^{\pi_1(k_1)} \cdot p_j^*\right)\right\}$.

Let PROFIT_1 be the value corresponding to k_1.

```
// Work on the direction τj → τ1.
```

Let $S_{j\to 1} \subseteq [m]$ be the set of indices i such that BATCH^i is in the direction $\tau_j \to \tau_1$.

Sort transactions in $S_{j\to 1}$ in an ascending order w.r.t. their exchange rate thresholds (break tie arbitrarily). Let π_2 be such an order and denote the k-th transaction in the order as $\text{BATCH}^{\pi_2(k)} = (\tau_j \to \tau_1, \delta_{\tau_j}^{\pi_2(k)}, r^{\pi_2(k)})$.

Let $k_2 \in \arg\max_{k_2\in[|S_{j\to 1}|]} \left\{\left(\sum_{k\in[k_2]} \delta_{\tau_j}^{\pi_2(k)}\right) \cdot \left(p_j^* - r^{\pi_2(k_2)} \cdot p_1^*\right)\right\}$.

Let PROFIT_2 be the value corresponding to k_2.

if $\textit{PROFIT}_1 \le 0$ & $\textit{PROFIT}_2 \le 0$ **then**

Do nothing.

end

else if $\textit{PROFIT}_1 \ge \textit{PROFIT}_2$ **then**

Include all $\{\text{BATCH}^{\pi_1(k)}\}_{k\in[k_1]}$ transactions;

Insert one mediator's BATCH $\left(\tau_j \to \tau_1, \left(\sum_{k\in[k_1]} \delta_{\tau_1}^{\pi_1(k)}\right) \cdot r^{\pi_1(k_1)}, \frac{1}{r^{\pi_1(k_1)}}\right)$.

end

else

Include all $\{\text{BATCH}^{\pi_2(k)}\}_{k\in[k_2]}$ transactions;

Insert one mediator's BATCH $\left(\tau_1 \to \tau_j, \left(\sum_{k\in[k_2]} \delta_{\tau_j}^{\pi_2(k)}\right) \cdot r^{\pi_2(k_2)}, \frac{1}{r^{\pi_2(k_2)}}\right)$.

end

end

PROOF. Algorithm 1 takes a batch of user transactions that forms a Fisher market and the exogenous prices as input and outputs a strategy to make the mediator obtain the optimal profits. It independently processes each pair (τ_1, τ_j) for all $j \in [2:n]$. For each pair, it decides which user transactions to select and how to insert its own attacking transaction (including the direction, the amount of endowment, and the exchange rate threshold).

As the first step, by Lemma 3.11, it suffices for algorithm 1 to work separately on pairs τ_1 and τ_j for every $j \in [2:n]$. Fix a $j \in [2:n]$ below. By Lemma 3.12, it suffices to consider strategies that only insert attacking transactions in one direction. So algorithm 1 tries both directions between τ_1 and τ_j and pick the best one. Take the direction $\tau_1 \to \tau_j$ as an example. Again, by Lemma 3.12, we can throw away all users' transactions that are for $\tau_1 \to \tau_j$ direction. Now the question reduces to the following: Given a set of users' transactions $\{\textsc{Batch}^i\}_{i \in S_{1\to j}}$ such that every $\textsc{Batch}^i$ is in the direction $\tau_1 \longrightarrow \tau_j$, what is the optimal way for the mediator to select a subset of users' transactions and insert its own transactions in the opposite direction, *i.e.*, the $\tau_j \longrightarrow \tau_1$ direction?

Obviously, we want to control the final prices to satisfy $p_1/p_j \le p_1^*/p_j^*$ (otherwise the mediator will loss profit), thus we should delete (and ignore) users' transactions $\textsc{Batch}^i$ such that $r^i > p_1^*/p_j^*$. Suppose that the final exchange rate is $p_1/p_j = r \le p_1^*/p_j^*$, then for each user's transaction $\textsc{Batch}^i$ such that $r^i \le r$, we can obtain profit $\delta^i_{\tau_1} \cdot (p_1^* - r \cdot p_j^*) \ge 0$. Thus we should include all users' transactions $\textsc{Batch}^i$ such that $r^i \le r$.

Now the correct way to attack them seems ready to come out: sort all user transactions in an ascending order π with respect to their exchange rate thresholds. Then the algorithm enumerates each involved threshold and calculates the corresponding profit:

$$\textsc{Profit}(k) = \left(\sum_{k' \in [k]} \delta_{\tau_1}^{\pi(k')} \right) \cdot (p_1^* - r^{\pi(k)} \cdot p_j^*), \tag{6}$$

where $r^{\pi(k)}$ is exchange rate threshold of the k-th user transaction and $\delta_{\tau_1}^{\pi(k')}$ is the endowment of the k'-th transaction in the order π. This $\textsc{Profit}(k)$ is obtained by setting the threshold of the k-th transaction as the exchange rate in this direction, selecting all user transactions with thresholds no larger than that, and inserting an attacking transaction in the opposite direction to provide the exact amount of token τ_j they need, which is $r^{\pi(k)} \cdot \left(\sum_{k' \in [k]} \delta_{\tau_1}^{\pi(k')} \right)$. Let $\textsc{Profit}_1$ and k_1 be the maximal profit and corresponding index for the direction $\tau_1 \longrightarrow \tau_j$. Symmetrically, the algorithm works on the direction $\tau_j \longrightarrow \tau_1$ and obtains the $\textsc{Profit}_2$ and k_2. If both $\textsc{Profit}_1$ and $\textsc{Profit}_2$ are no larger than 0, we just ignore all user transactions between π_1 and π_j and insert no transaction between them, as neither direction is profitable. Otherwise, choosing the strategy with the highest profit.

It is easy to verify that all transactions in the batch will be successfully executed. This finishes the proof.

□

3.5 NP-hardness for Optimal MEV under Arrow-Debreu Market

In this section, we show the computational hardness of finding an optimal MEV strategy when the user transactions form a general Arrow-Debreu market.

Theorem 3.14. *It is NP-hard to compute an optimal strategy to the Batch-MEV problem.*

Proof. We reduce the NP-hard Maximum Acyclic Subgraph Problem to our MEV optimization problem. Recall that an instance of the Maximum Acyclic Subgraph problem contains a directed graph and asks to find an acyclic subgraph with a maximum number of edges.

Suppose we are given an arbitrary graph $G = (V, E)$ where V is a set of $|V| = n$ vertices and E is a set of $|E| = m$ directed edges (there are no multiple edges with the same source and target nodes). Note that we can decide if a graph is acyclic in polynomial time, and if the graph is acyclic, we can simply output $|E|$ as the answer of the given instance of the Maximum Acyclic Subgraph problem. Thus, without loss of generality, we assume that the graph is not acyclic, which implies the maximum acyclic subgraph of G has $q < m$ edges.

Then, we construct the MEV optimization problem as follows. For each vertex $v_i \in V$, we construct a token τ_i and set its exogenous price p_i^* as 1. For each edge $(i, j) \in E$, we construct a user transaction $\textsc{Batch} = \left(\tau_i \longrightarrow \tau_j, m^m, \frac{1}{m^m}\right)$. We refer the constructed instance to $\{\textsc{Batch}^{(i,j)}\}_{(i,j)\in E}$.

The correctness of this reduction follows from the following two lemmas.

Lemma 3.15. *If the maximum acyclic subgraph of G has q edges, then the mediator's optimal profit under the instance $\left\{\text{Batch}^{(i,j)}\right\}_{(i,j)\in E}$ lies in $\left[\left(m^m - m^{m-1}\right)\cdot q, m^m \cdot q\right]$.*

Proof. Given Lemma 3.8, the upper bound proof is easy. Consider any mediator's strategy $(S, \left\{\text{Batch}^i\right\}_{i\in[m+1:m+k]})$ *such that the economy graph of* $\left\{\text{Batch}^i\right\}_{i\in S}$ *is acyclic*, so we have $|S| \leq q$. The requirement of acyclic graph is without loss of generality due to Lemma 3.8. For each token in the instance, its demand under a market equilibrium equals the supply. The mediator's profit is the value of received tokens in the final allocation minus the value of its initial endowments, which is upper bounded by the value of users' initial endowments, *i.e.*, $1 \cdot m^m \cdot q$.

The lower bound proof is more involved and we focus on it next.

Assume that $G' = (V', E')$ is a directed acyclic subgraph of G that has q edges. We construct a strategy for the mediator that can obtain at least $(m^m - m^{m-1}) \cdot q$ profit. The construction is based on *levels* of nodes in G', which we define next. We start with all nodes that have no incoming edges and label them as level 1; then we (virtually) remove the level-1 node(s) and all edges taking them as the source node; next, we label all nodes that have no incoming edges as level 2 and repeatedly process the remaining subgraph with an increasing level (*i.e.*, level 3, 4, $\cdots$) until all nodes are labeled. In this way, all selected user transactions start from a token with a lower level and end at a token with a higher level.

Now we are able to describe the mediator's strategy: First, the mediator selects all users' transactions that correspond to edges in E'. Then for each selected user transaction $\text{Batch} = \left(\tau_i \longrightarrow \tau_j, m^m, \frac{1}{m^m}\right)$ in G', the mediator inserts an attacking transaction Batch' in its opposite direction, specifically, $\text{Batch}' = \left(\tau_j \longrightarrow \tau_i, m^{m+l_i-l_j}, m^{l_j-l_i}\right)$ where l_i is the level of node i (*i.e.*, token τ_i). The two steps derive an attacking strategy. Next, we prove that these selected user transactions and newly inserted attacking transactions are executed under equilibrium, and bring a profit no less than $\left(m^m - m^{m-1}\right) \cdot q$.

Let $\mathbf{p}$ be a price vector of tokens in G', where token τ_i's price $p_i = m^{l_i}$. It is easy to verify that $\mathbf{p}$ is the price equilibrium where the unique allocation is as follows: Each selected user transaction $\text{Batch} = \left(\tau_i \longrightarrow \tau_j, m^m, \frac{1}{m^m}\right)$ is executed at the exchange rate $m^{l_i-l_j}$ which is larger than its requirement $\frac{1}{m^m}$, and receives $m^{m+l_i-l_j}$ amount of token τ_j; Each attacking transaction in the opposite direction $\text{Batch}' = \left(\tau_j \longrightarrow \tau_i, m^{m+l_i-l_j}, m^{l_j-l_i}\right)$ receives m^m amount of token τ_i, bringing a profit of $m^m - m^{m+l_i-l_j} \geq m^m - m^{m-1}$. Combining that we insert q attacking transactions in total, the mediator's profit is at least by $\left(m^m - m^{m-1}\right) \cdot q$. □

Lemma 3.16. *If the mediator's optimal profit is in $\left[\left(m^m - m^{m-1}\right)\cdot q, m^m \cdot q\right]$, then the number of edges in the maximum acyclic subgraph of G is q.*

Proof. The proof in fact follows from the disjointedness of $\left[\left(m^m - m^{m-1}\right)\cdot q, (m^m - 1) \cdot q\right]$ for different q. Specifically, we show the upper bound for $q-1$ interval is strictly smaller than the lower bound of q interval, which follows from the following simple calculation.

$$q < m \quad \Rightarrow \quad mq - q > mq - m \quad \Rightarrow \quad (m-1)qm^{m-1} > m(q-1)m^{m-1}$$
$$\Rightarrow \quad \left(m^m - m^{m-1}\right)\cdot q > m^m \cdot (q-1).$$

□

Overall, deciding which interval the mediator's optimal profit lies in is equivalent to finding the size of maximum acyclic subgraph of G. So finding an optimal strategy is as hard as solving the Maximum Acyclic Subgraph Problem, which is NP-hard.

This finishes the proof. □

3.6 UGC-hard for 1/2-approximate MEV under Arrow-Debreu Markets

The main goal of this section is to strengthen the hardness result above by showing computational hardness to computing a nearly 1/2-approximate MEV under Arrow-Debreu markets. Notably, the 1/2 threshold implies that the mediator must forgo half of the potential MEV profit. Moreover, we conjecture that even achieving a 1% approximation remains computationally hard. However, proving this would require additional techniques, which we discuss further in the final section.

Theorem 3.17. *Let $\epsilon \in (0, 1/2)$ be any constant. Assuming unique game conjecture, it is NP-hard to compute a $(1/2 + \epsilon)$-approximation to the Batch-MEV problem.*

Proof. We still reduce the Maximum Acyclic Subgraph problem to our MEV Optimization problem. Suppose we are given an arbitrary graph $G = (V, E)$ where V is a set of $|V| = n$ vertices and E is a set of $|E| = m$ directed edges. Again, we assume that G is not acyclic. We recall the hardness-of-approximation result for the Maximum Acyclic Subgraph problem.

Theorem 3.18 ([14]). *Let $\epsilon_1 \in (0, 1/2)$ be any constant. Assuming the Unique Game Conjecture, it is NP-hard to compute a $(1/2 + \epsilon_1)$-approximation to the Maximum Acyclic Subgraph problem.*

We construct the MEV optimization problem in the same way as in the proof of Theorem 3.14: For each vertex $v_i \in V$, we construct a token τ_i and set its exogenous price p_i^* as 1; for each edge $(i, j) \in E$, we construct a user transaction $\text{Batch} = \left(\tau_i \to \tau_j, m^m, \frac{1}{m^m}\right)$. We refer the constructed instance to $\left\{\text{Batch}^{(i,j)}\right\}_{(i,j)\in E}$.

Assume that m is sufficiently large. Let ϵ_1 be a constant such that $1/2 + \epsilon_1 < (1/2 + \epsilon)(1 - \frac{1}{m})$. Theorem 3.17 follows from the claim that if one can find a $(1/2+\epsilon)$-approximation to the Batch-MEV problem, then we can extract a $(1/2+\epsilon_1)$-approximation to the Maximum Acyclic Subgraph problem in polynomial time. To this end, we show the following lemma.

Lemma 3.19. *For any strategy $(S, \left\{\text{Batch}^i\right\}_{i\in[m+1:m+k]})$, if the mediator can obtain profit $\mathcal{U}$, then we can extract an acyclic subgraph of G which has at least $\mathcal{U}/m^m$ edges.*

Proof. Note that we can assume the given strategy $(S, \left\{\text{Batch}^i\right\}_{i\in[m+1:m+k]})$ satisfies that the selected user transactions form an acyclic cycle by Lemma 3.8. If not, we can use the strategy given in Lemma 3.8 to purify it and guarantee that under the new strategy, the profit is at least $\mathcal{U}$. Furthermore, this can clearly be done in polynomial time. So we assume that S forms an acyclic cycle. By Lemma 3.15, we know that $\mathcal{U} \leq |S| \cdot m^m$. □

Let q be the number of edges in a maximum acyclic subgraph of G. Then by Lemma 3.15, we know the optimal MEV is at least $(m^m - m^{m-1}) \cdot q$. If we have a strategy that can obtain $(1/2 + \epsilon)$-approximate MEV, then such a strategy can obtain $(1/2 + \epsilon)(m^m - m^{m-1}) \cdot q$ profits. Combining with Lemma 3.19, we know that we can extract an acyclic subgraph with at least $(1/2 + \epsilon)q \cdot \frac{m^m - m^{m-1}}{m^m} = (1/2 + \epsilon)q \cdot (1 - \frac{1}{m}) \geq (1/2 + \epsilon_1) \cdot q$ edges.

This finishes the proof. □

4 Hard to Approximate MEV in Constant Function Market Makers

In this section, we study the computation of approximating MEV in CFMMs. It is well known that the outcomes of transactions executed on CFMMs are sensitive to transaction-ordering attacks, which leaves space for a mediator (*e.g.*, miner, validator) to extract profits. We start by formalizing the CFMM and the execution of transactions.

4.1 CFMMs Formalization and Execution of Transactions

We consider a CFMM A for two types of tokens $\mathcal{X}, \mathcal{Y} \in \{\tau_1, \cdots, \tau_n\}$, where we use $s = (x, y)$ as a state of A that represents the current reserves of tokens $\mathcal{X}$ and $\mathcal{Y}$. The trading invariant can be modeled by a constant function on two variables $F(x, y)$. Throughout this section, we will only consider the Swap-like transactions, and there is no deposit or redemption of liquidity, so sometimes we equivalently use $s = (x, y) \in F$ to represent a reserving point on the curve $F(\cdot, \cdot)$. We assume two natural properties about these constant functions: (1) for any two points $(x, y), (x', y') \in F$, we have $x > x' \Leftrightarrow y < y'$, namely, when the reserve of $\mathcal{X}$ increases, the reserve of $\mathcal{Y}$ decreases and vice versa; (2) $F(x, y)$ is differentiable and the marginal exchange rate $\left|\frac{\partial F/\partial x}{\partial F/\partial y}\right|$ is decreasing with respect to x. We note that most CFMMs satisfy these two properties, including Uniswap v2 and v3. Finally, we denote the fraction of the swap fee in AMM A by $f \in [0, 1)$.

From the two properties above, we know that for any x, there is exactly one y such that $(x, y) \in F$ and vice versa. So for simplicity of notations, we use $F_y(x)$ to denote that y such that $(x, y) \in F$ and similarly define $F_x(y)$.

Suppose we are given a set of Swap transactions $\{\text{Swap}^1, \cdots, \text{Swap}^m\}$ on AMM A, where each $\text{Swap}^i = (\delta_{\mathcal{X}}, r)$ or $\text{Swap}^i = (\delta_{\mathcal{Y}}, r)$. Without any ambiguity, we omit the swapping direction in this section and use the shortened notation $\text{Swap}^i(\delta_{\mathcal{X}}, r)$ to represent that the i-th transaction Swap^i would like to sell $\delta_{\mathcal{X}}$ amount of token $\mathcal{X}$ to obtain at least $\delta_{\mathcal{X}} \cdot r$ amount of $\mathcal{Y}$, *i.e.*, with r as the lowest acceptable exchange rate. The meaning of $\text{Swap}^i(\delta_{\mathcal{Y}}, r)$ is analogous.

Pick an arbitrary permutation $\pi : [m] \to [m]$ as the execution order of all these m transactions and let $s_0 = (x_0, y_0)$ be the initial state of A before processing these transactions. The execution works as follows: Consider the i-th round and $\text{Swap}^{\pi(i)} = (\delta_{\mathcal{X}}, r)$. Let $\Delta \coloneqq F_y(x_{i-1}) - F_y(x_{i-1} + (1-f)\delta_{\mathcal{X}})$. If the condition $\Delta \geq \delta_{\mathcal{X}} \cdot r$ holds, then this swap is successfully executed and $s_i = \left(x_{i-1} + (1-f)\delta_{\mathcal{X}}, F_y(x_{i-1} + (1-f)\delta_{\mathcal{X}})\right)$; otherwise, the transaction fails and $s_i = s_{i-1}$. The case where $\text{Swap}^{\pi(i)} = (\delta_{\mathcal{Y}}, r)$ can be defined similarly.

4.2 MEV Optimization Problem

In this subsection, we formalize the mediator's strategies and define its strategy space. Intuitively, a mediator can potentially delete some users' transactions, insert its own transactions, and pick an arbitrary order of execution of the selected transactions.

Definition 4.1 (Strategy Space). Given a set of users' transactions $\left\{\text{Swap}^i\right\}_{i \in [m]}$ and an initial state $s_0 = (x_0, y_0)$, a mediator could create k its own transactions $\left\{\text{Swap}^i\right\}_{i \in [m+1:m+k]}$, select a subset of users' transactions $S \subseteq [m]$, and pick an execution order (a permutation) over all these transactions $\pi : [|S \cup [m+1 : m+k]|] \to S \cup [m+1 : m+k]$.

Definition 4.2 (Utility Function). Mediator's profit $U(\{\text{Swap}^i\}_{i \in [m+1:m+k]}, S, \pi)$ is defined as

$$\sum_{i \in [|S|+k], \pi(i) \in [m+1:m+k]} \frac{x_{i-1} - x_i}{1 - f \cdot \mathbb{1}_{\{x_i > x_{i-1}\}}} \cdot p_x^* + \frac{y_{i-1} - y_i}{1 - f \cdot \mathbb{1}_{\{y_i > y_{i-1}\}}} \cdot p_y^*, \tag{7}$$

where p_x^* and p_y^* are exogenous (or say, off-chain) prices of $\mathcal{X}$ and $\mathcal{Y}$ respectively.

This definition generalizes the idea of sandwich attack and can capture a wide range of order manipulation attacks. Next we formalize the computational problem associated with it.

Definition 4.3 (CFMM-MEV). The CFMM-MEV problem refers to the following computational problem: Given user transactions $\left\{\text{Swap}^i\right\}_{i \in [m]}$ and an initial state $s_0 = (x_0, y_0)$, compute a strategy $(\{\text{Swap}^i\}_{i \in [m+1:m+k]}, S, \pi)$ that achieves the optimal revenue. For $c \in (0, 1)$, we say a strategy is a c-approximation to CFMM-MEV if the strategy obtains at least c-fraction of the optimal MEV revenue.

4.3 Computational Hardness

Recall that in literature many excellent works studied the same or similar attack, with specific focus on empirical approaches [10, 24, 33, 34], or special cases (*e.g.*, no swap fee [3] or attacking one user's transaction [15]), but no polynomial time algorithm for the general setting is known. This is not a coincidence, as what we are going to show in this section.

In particular, as our main theorem in this section, we show that computing even a 0.01%-approximation to CFMM-MEV is NP-hard when $f > 0$ is any constant (say $f = 0.3\%$), which is what happens in the real blockchain world. This indicates that if one wants to attract optimal MEV, it is necessary to design heuristic algorithms, but not hope for a theoretically efficient even approximation algorithm.

THEOREM 4.4. *Let $\epsilon, f \in (0, 1)$ be any universal constants. It is NP-hard to compute an ϵ-approximation to the CFMM-MEV problem, even with the constant function $F(x, y) = xy$ (i.e., Uniswap v2).*

Due to space constraints, we postpone the proof to the appendix.

5 Discussion and Open Problems

We primarily focus on MEV issues in Batch Auctions, as they have received relatively less attention. Additionally, the approximation of CFMM-MEV has been essentially resolved.

Stronger Hardness of Approximation for Batch Auctions. As discussed earlier, we conjecture that even achieving a 1%-approximation for the BATCH-MEV problem remains computationally hard under standard complexity assumptions, such as P≠NP and the Unique Games Conjecture. However, our current reduction is based on the Maximum Acyclic Subgraph problem, which admits a very simple 1/2-approximation algorithm (see [14]). This suggests that new techniques are necessary to establish stronger hardness of approximation results. Notably, Lemma 3.8 highlights the attacker's need to identify an acyclic subgraph within the economy graph. While this provides a useful structural insight, further properties of Batch Auctions could be leveraged to reinforce the hardness results.

Market Equilibrium Manipulation. We provided a polynomial-time algorithm to attack a set of transactions that form a Fisher market. It is also possible to study the same market equilibrium manipulation problem for the general linear markets (where the utility functions may have more than two non-zero coefficients). In particular, we are curious if our ideas could be extended to general linear Fisher markets to obtain a polynomial-time optimal-attack algorithm. If so, this could be a new point of view about the fundamental differences between Fisher markets and Arrow-Debreu markets.

Empirical MEV on BATCH Auctions. The theoretical results of computational complexity are fundamental, but our interest of course extends to practical environments MEV, as batch auctions are happening in blockchain everyday with incredible volumes. The possibilities of novel MEV behaviors in batch auctions have not been discussed by the community before, and it is totally possible to design heuristic algorithms to extract MEV. This is beyond the scope of the current work, but we are expecting emerging works related to MEV in batch auctions in the near future.

Acknowledgments

Xi Chen would like to thank the support from NSF IIS-1845444. Elynn Chen's research is supported in part by the NSF Award DMS-2412577. Yuhao Li is supported by NSF grants IIS-1838154, CCF-2106429, and CCF-2107187. The authors are grateful to anonymous reviewers for helping them improve the paper and to Sen Yang for his generous help with the empirical studies.

References

[1] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. 2021. Uniswap v3 core. *Tech. rep., Uniswap, Tech. Rep.* (2021). https://uniswap.org/whitepaper-v3.pdf

[2] Kenneth J Arrow and Gerard Debreu. 1954. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society* (1954), 265–290.

[3] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. 2022. Maximizing Extractable Value from Automated Market Makers. In *Financial Cryptography (Lecture Notes in Computer Science, Vol. 13411)*. Springer, 3–19.

[4] Eric Budish, Peter Cramton, and John Shim. 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130, 4 (2015), 1547–1621.

[5] Andrea Canidio and Robin Fritsch. 2023. Arbitrageurs' profits, LVR, and sandwich attacks: batch trading as an AMM design response. *arXiv preprint arXiv:2307.02074* (2023).

[6] Andrea Canidio and Robin Fritsch. 2023. Batching trades on automated market makers. In *5th Conference on Advances in Financial Technologies (AFT 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.

[7] TH Chan, Ke Wu, and Elaine Shi. 2024. Mechanism Design for Automated Market Makers. *arXiv preprint arXiv:2402.09357* (2024).

[8] Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. 2017. The Complexity of Non-Monotone Markets. *J. ACM* 64, 3 (2017), 20:1–20:56. https://doi.org/10.1145/3064810

[9] Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. 2006. Leontief economies encode nonzero sum two-player games. In *SODA*, Vol. 6. 659–667.

[10] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*. IEEE, 910–927. https://doi.org/10.1109/SP40000.2020.00040

[11] Flashbots. February 8th 2025. Cowswap Solver Metrics. https://dune.com/flashbots/cowswap-solver-metrics

[12] David Gale. 1976. The linear exchange model. *Journal of Mathematical economics* 3, 2 (1976), 205–209.

[13] Tivas Gupta, Mallesh M. Pai, and Max Resnick. 2023. The Centralizing Effects of Private Order Flow on Proposer-Builder Separation. In *AFT (LIPIcs, Vol. 282)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 20:1–20:15.

[14] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. 2008. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 573–582.

[15] Lioba Heimbach and Roger Wattenhofer. 2022. Eliminating Sandwich Attacks with the Help of Game Theory. In *ASIA CCS '22: ACM Asia Conference on Computer and Communications Security, Nagasaki, Japan, 30 May 2022 - 3 June 2022*, Yuji Suga, Kouichi Sakurai, Xuhua Ding, and Kazue Sako (Eds.). ACM, 153–167. https://doi.org/10.1145/3488932.3517390

[16] Subhash Khot. 2002. On the power of unique 2-prover 1-round games. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*. 767–775.

[17] Kshitij Kulkarni, Theo Diamandis, and Tarun Chitra. 2023. Routing MEV in Constant Function Market Makers. In *International Conference on Web and Internet Economics*. Springer, 456–473.

[18] Yuhao Li, Mengqian Zhang, Jichen Li, Elynn Chen, Xi Chen, and Xiaotie Deng. 2023. MEV Makes Everyone Happy under Greedy Sequencing Rule. In *Proceedings of the 2023 Workshop on Decentralized Finance and Security, DeFi 2023, Copenhagen, Denmark, 30 November 2023*, Kaihua Qin and Fan Zhang (Eds.). ACM, 9–15. https://doi.org/10.1145/3605768.3623543

[19] Zihao Li, Jianfeng Li, Zheyuan He, Xiapu Luo, Ting Wang, Xiaoze Ni, Wenwu Yang, Xi Chen, and Ting Chen. 2023. Demystifying DeFi MEV Activities in Flashbots Bundle. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS 2023, Copenhagen, Denmark, November 26-30, 2023*, Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda (Eds.). ACM, 165–179. https://doi.org/10.1145/3576915.3616590

[20] Robert R Maxfield. 1997. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics* 27, 1 (1997), 23–51.

[21] Lionel McKenzie. 1954. On equilibrium in Graham's model of world trade and other competitive systems. *Econometrica: Journal of the Econometric Society* (1954), 147–161.

[22] Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046* (2022).

[23] Alexandre Obadia, Alejo Salles, Lakshman Sankar, Tarun Chitra, Vaibhav Chellani, and Philip Daian. 2021. Unity is Strength: A Formalization of Cross-Domain Maximal Extractable Value. *CoRR* abs/2112.01472 (2021). arXiv:2112.01472 https://arxiv.org/abs/2112.01472

[24] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying Blockchain Extractable Value: How dark is the forest?. In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*. IEEE, 198–214. https://doi.org/10.1109/SP46214.2022.9833734

[25] Geoffrey Ramseyer, Ashish Goel, and David Mazières. 2023. SPEEDEX: A Scalable, Parallelizable, and Economically Efficient Decentralized EXchange. In *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2023, Boston, MA, April 17-19, 2023*, Mahesh Balakrishnan and Manya Ghobadi (Eds.). USENIX Association, 849–875. https://www.usenix.org/conference/nsdi23/presentation/ramseyer

[26] Geoffrey Ramseyer, Mohak Goyal, Ashish Goel, and David Mazières. 2024. Augmenting batch exchanges with constant function market makers. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 986–1016.

[27] Jan Christoph Schlegel, Mateusz Kwasnicki, and Akaki Mamageishvili. 2023. Axioms for Constant Function Market Makers. In *EC*. ACM, 1079.

[28] Cowswap Team. Feb 8th 2025. Cow Protocol Overview. https://docs.cow.fi/

[29] Christof Ferreira Torres, Ramiro Camino, and Radu State. 2021. Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael D. Bailey and Rachel Greenstadt (Eds.). USENIX Association, 1343–1359. https://www.usenix.org/conference/usenixsecurity21/presentation/torres

[30] Ye Wang, Yan Chen, Haotian Wu, Liyi Zhou, Shuiguang Deng, and Roger Wattenhofer. 2022. Cyclic Arbitrage in Decentralized Exchanges. In *Companion of The Web Conference 2022, Virtual Event / Lyon, France, April 25 - 29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 12–19. https://doi.org/10.1145/3487553.3524201

[31] Sen Yang, Kartik Nayak, and Fan Zhang. 2024. Decentralization of Ethereum's Builder Market. *arXiv preprint arXiv:2405.01329* (2024).

[32] Mengqian Zhang, Sen Yang, and Fan Zhang. 2024. RediSwap: MEV Redistribution Mechanism for CFMMs. *arXiv preprint arXiv:2410.18434* (2024).

[33] Liyi Zhou, Kaihua Qin, Antoine Cully, Benjamin Livshits, and Arthur Gervais. 2021. On the Just-In-Time Discovery of Profit-Generating Transactions in DeFi Protocols. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 919–936. https://doi.org/10.1109/SP40001.2021.00113

[34] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc Viet Le, and Arthur Gervais. 2021. High-Frequency Trading on Decentralized On-Chain Exchanges. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 428–445. https://doi.org/10.1109/SP40001.2021.00027

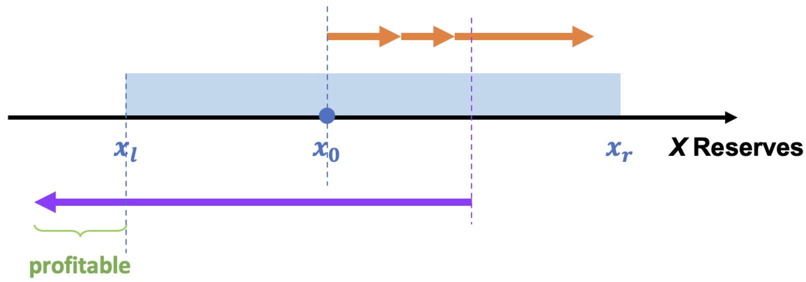

Fig. 1. An illustration of the reduction in the proof of Theorem 4.4. The dark line means x-axis which represents the reserved amount of $\mathcal{X}$ tokens. Light blue shadow means the "arbitrage-free" interval $[x_\ell, x_r]$, where any mediator's transaction cannot obtain profits if the execution state is in this interval due to the trade fees. The orange arrows represent the set of $\mathcal{X} \to \mathcal{Y}$ users' transactions; they satisfy that starting from the initial state, after executing all these transactions, the state remains in the arbitrage-free interval (so that the mediator can never back-run them for profits). The purple arrow represents the only user's $\mathcal{Y} \to \mathcal{X}$ transaction (which is the source where the mediator could potentially obtain profits by back-running), and the purple dotted line represents the state that exactly satisfies its slippage tolerance. The left green part is where the mediator can potentially back-run and obtain profits. We design the purple transaction such that the following holds: If there is a subset of transactions, after executing them, the state of the pool can reach precisely the purple dotted line, then the mediator can obtain non-trivial profits indicated by the green area; otherwise, the mediator cannot obtain any profits. The reduction uses the idea that the orange transactions can encode any instance of the Partition problem. The formal reasoning is shown in the proof of Theorem 4.4.

A Proof of Theorem 4.4

We first include a simple observation that can help us simplify the notations.

Observation 1 (No-Deleting in AMMs). Given any set of users' transactions $\left\{\textsc{Swap}^i\right\}_{i\in[m]}$ and an initial state $s_0 = (x_0, y_0)$, it is without loss of generality to assume that the strategy satisfies $S = [m]$, *i.e.*, the mediator will always select the full set of transactions.

Proof. The proof is easy. For mediator's transactions, it can always create the transactions that it needs. For the users' transactions, if the mediator didn't select some users' transactions in a strategy, it can also equivalently put them at the very end of the sequence, and it will not affect the mediator's profits. □

From Observation 1, we may omit the parameter S for simplicity of notations.

Theorem 4.4. *Let $\epsilon, f \in (0, 1)$ be any universal constants. It is NP-hard to compute an ϵ-approximation to the CFMM-MEV problem, even with the constant function $F(x, y) = xy$ (i.e., Uniswap v2).*

Proof. We reduce the NP-hard Partition problem to our problem. Recall that an instance of the Partition problem contains m positive integers $\{c_1, \cdots, c_m\}$ and asks if it can be partitioned into two subsets $S_1, S_2 \subseteq [m]$ with $S_1 \cap S_2 = \emptyset$ and $S_1 \cup S_2 = [m]$ such that the sum of numbers in S_1 equals that in S_2, *i.e.*, $\sum_{i\in S_1} c_i = \sum_{i\in S_2} c_i = \frac{1}{2}\sum_{i\in[m]} c_i$.

Suppose that we are given an arbitrary set of m positive integers $\{c_1, \cdots, c_m\}$ such that the sum of all c_i's is an even number (otherwise, the answer to the Partition problem is obviously "no").

To start the reduction, we will construct a CFMM, a set of uses' transactions $\left\{\textsc{Swap}^i\right\}$, and an initial state. Concretely, it suffices for us to use the product function $F(x, y) = xy$ (other CFMMs that satisfy our two properties defined in Section 4 will also work). Let the external prices $p_x^* = p_y^* = 1$, and the initial state $s_0 = (x_0, y_0)$ be such that $x_0 = y_0$. Note that $\frac{\partial F/\partial x}{\partial F/\partial y} = p_x^*/p_y^*$. We will have $m+1$ users' transactions; m of them will be $\textsc{Swap}(\mathcal{X} \to \mathcal{Y})$, namely, they are trying to sell $\mathcal{X}$ to obtain $\mathcal{Y}$, and the last one will be $\textsc{Swap}(\mathcal{Y} \to \mathcal{X})$.

Let x_r be such that $\frac{\partial F_y(x_r)}{\partial x} = \frac{(1-f)p_x^*}{p_y^*}$ and x_ℓ be such that $\frac{\partial F_y(x_\ell)}{\partial x} = \frac{p_x^*}{(1-f)p_y^*}$. Note that these two points correspond to the two endpoints of the "non-arbitrage" interval.

Let K be the smallest positive integer such that $\sum_{i\in[n]} c_i/K < x_r - x_0$. Then we let $\{d_1, \cdots, d_m\} = \{c_1/K, \cdots, c_m/K\}$ and $t = \frac{1}{2}\sum_{i\in[m]} d_i$. Our problem is still to decide whether there exists $S_1 \subseteq [m]$ such that $\sum_{i\in S_1} d_i = t$. Obviously this doesn't change the original problem, but makes it easier for us to finish the reduction.

Let $x^* = x_0 + t$. We first construct the user transaction $\textsc{Swap}^{m+1} = (\mathcal{Y} \to \mathcal{X}, \delta_y)$, namely, the only $\textsc{Swap}(\mathcal{Y} \to \mathcal{X})$ transaction. Both parameters δ_y and the exchange ratio threshold are very crucial, and they are defined as follows.

We let

$$\delta_y = (F_y(x_\ell) - F_y(x^*) + \eta)/(1-f)$$

with a sufficiently small number η (to be specified later) and the exchange ratio threshold to be

$$r = \frac{x^* - F_x((1-f)\delta_y + F_y(x^*))}{\delta_y}.$$

Intuitively, these two parameters work as follows: ($\diamond$) the exchanged ratio threshold guarantees that the transaction $\textsc{Swap}^{m+1}$ can be executed at some state $s = (x, y)$ if and only if $x \geq x^*$; and ($\triangle$) the parameter δ_y guarantees that after the transaction $\textsc{Swap}^{m+1}$ is successfully executed, the new state $s' = (x', y')$ satisfies that $y' - F_y(x_\ell) \leq \eta$.

We then construct m users' transactions $\left\{\textsc{Swap}^i\right\}_{i\in[m]}$, where each $\textsc{Swap}^i = (\mathcal{X} \to \mathcal{Y}, \delta_{\mathcal{X}} = d_i/(1-f))$ (we'll specify the thresholds of exchange ratios of these transactions later; they will also be crucial). Intuitively, this construction means if we only execute users' transactions, then we can reach the state $s^* = (x^*, F_y(x^*))$ *if and only if* there exists $S_1 \subseteq [m]$ such that $\sum_{i\in S_1} d_i = t$, which corresponds to the answer to the given instance of the Partition problem.

For each transaction $\textsc{Swap}^i$ for $i \in [m]$, we set the threshold of exchange ratio to be

$$\frac{F_y(x^* - d_i) - F_y(x^*)}{d_i/(1-f)}.$$

Intuitively, this means if a user's transaction $\textsc{Swap}^i$ is executed at state $s = (x, y)$, then the $\textsc{Swap}(\delta_{\mathcal{X}})$ is successfully executed if and only if $x \leq x^* - d_i$. Equivalently, after successfully executing any user's transaction, the state $s' = (x', y')$ satisfies $x' \leq x^*$. ($\star$)

This finishes the construction of an instance of the CFMM-MEV problem. An intuitive illustration of the construction is presented in Figure 1. Next, we show that if a mediator can find an ϵ-approximation to the CFMM-MEV problem, then it can solve the given Partition instance.

To this end, we define the mediator's (ideal) maximal MEV as follows (recall that x_ℓ is such that $\frac{\partial F_y(x_\ell)}{\partial x} = \frac{p_y^*}{(1-f)p_x^*}$):

$$\texttt{IDL} = \eta \cdot p_y^* - \frac{x_\ell - F_x(F_y(x_\ell) + \eta)}{1-f} \cdot p_x^*.$$

This value corresponds to the arbitrage value starting from the state $(\cdot, F_y(x_\ell) + \eta)$ to the state $(\cdot, F_y(x_\ell))$. (We omit the x part for convenience since y can determine the state of the pool.)

We set η to be sufficiently small such that $\texttt{IDL} < \frac{f}{(1-f)K}$. The correctness of our reduction follows from the next lemma.

Lemma A.1. *If there exists S_1 such that $\sum_{i\in S_1} d_i = t$, then the mediator's profit is at least* `IDL`*; otherwise, the mediator's best profit is 0.*

The reason that this lemma implies the correctness of the reduction is as follows. If an ϵ-approximation has profits larger than 0, then it must be the case that there exists S_1 such that $\sum_{i\in S_1} d_i = t$; if an ϵ-approximation has profits 0, it must be the case that there does not exist S_1 such that $\sum_{i\in S_1} d_i = t$.

We provide the proof of Lemma A.1.

Proof of Lemma A.1. One direction is relatively easy: If there exists S_1 such that $\sum_{i\in S_1} d_i = t$, then the mediator can obtain `IDL` profits. The strategy is as follows:

(1) Execute users' transactions Swap^i for all $i \in S_1$ (under arbitrary order);
(2) Execute the user's transaction Swap^{m+1};
(3) Execute one mediator's transaction $\text{Swap}(\mathcal{X} \to \mathcal{Y}, \delta_x = (x_\ell - F_x(F_y(x_\ell)+\eta))/(1-f), r = 0)$.

It is easy to verify that mediator's profits equal `IDL`.

The other direction is more involved. We will show that if there doesn't exist S_1 such that $\sum_{i\in S_1} d_i = t$, then for any mediator's strategy, we have $U(\{\text{Swap}^i\}_{i\in[m+2:m+1+k]}, \pi) \le 0$.

Fix arbitrary strategy $(\{\text{Swap}^i\}_{i\in[m+2:m+1+k]}, \pi)$. Without loss of generality, we assume that all transactions $\left\{\text{Swap}^{\pi(i)}\right\}_{i\in[m']}$ for some $m' \in [m+1+k]$ are successfully executed (otherwise we can put these transactions at the back and mediator's profits stay the same). We define a potential function of state $\psi : \{s\}_{s\in F} \to \mathbb{R}$ as follows (recall that $x^* = x_0 + t$):

$$\psi(x,y) = \begin{cases} (x - x_r)\cdot p_x^* + \frac{y - F_y(x_r)}{1-f}\cdot p_y^*, & x > x_r; \\ \frac{x-x_\ell}{1-f}\cdot p_x^* + \left(y - F_y(x_\ell)\right)\cdot p_y^*, & x < x_\ell; \\ 0, & x \in [x_\ell, x_r]. \end{cases}$$

Let s_i be the state that is after executing transaction $\text{Swap}^{\pi(i)}$ and U_i denote the mediator's profit after executing i-th transaction. Let i^* be the index with $\pi(i^*) = m+1$, *i.e.*, the user's transaction $\text{Swap}^{m+1}(\mathcal{Y} \to \mathcal{X})$. We will inductively show that after executing i-th transaction, the profit of the mediator

$$U_i + \psi(s_i) \le U_{i-1} + \psi(s_{i-1})$$

for all $i \in [i^* - 1]$. Note that at the beginning of s_0, we have $U_0 = 0$ and $\psi(s_0) = 0$. Now let's move to the induction step and consider any $i \in [i^* - 1]$.

Case 1: $\text{Swap}^{\pi(i)}$ is a user's transaction. By the observation mentioned above, we know that $\text{Swap}^{\pi(i)}$ is successfully executed. Due to the thresholds of exchange ratios we constructed, we know that $x_i \le x^*$ (recall $(\star)$ above). Combining with $x^* \le x_r$, this means $\psi(s_i) \le \psi(s_{i-1})$. Note also that executing a user's transaction doesn't affect mediators' profit. So we have $U_i = U_{i-1}$. To conclude, we have $U_i + \psi(s_i) \le U_{i-1} + \psi(s_{i-1})$.

Case 2: $\text{Swap}^{\pi(i)}$ is a mediator's transaction. In this case, $U_i + \psi(s_i) \le U_{i-1} + \psi(s_{i-1})$ directly follows from the definition of ψ function.

This finishes the induction. Thus, we conclude that $U_{i^*-1} + \psi(s_{i^*-1}) \le U_0 + \psi(s_0) = 0$. Note that $\psi(s) \ge 0$ for all $s \in F$, so we can conclude that $U_{i^*-1} \le 0$.

Importantly, if all transactions $\left\{\text{Swap}^{\pi(1)}, \cdots, \text{Swap}^{\pi(i^*-1)}\right\}$ are users' transactions, then we have $x_0 \le x_1 \le x_2 \le \cdots \le x_{i^*-1} \le x^* - 1/K$. This follows from our condition from the very beginning that there doesn't exist $S_1 \subset [m]$ such that $\sum_{i\in S_1} d_i = t$. If this is the case, then $\text{Swap}^{\pi(i^*)}$ will not be able to be executed successfully due to $(\diamond)$ and mediator will not get any profit. Intuitively, $\text{Swap}^{\pi(i^*)}$ is the only transaction where mediator can extract some profits. So to make $\text{Swap}^{\pi(i^*)}$ successfully executed, it has to be the case $x_{i^*-1} \ge x^*$, which means there are some mediator's transaction in $\left\{\text{Swap}^{\pi(i)}\right\}_{i\in[i^*-1]}$. In particular, the cumulative amount of $\mathcal{X}$ token that the mediator puts in the pool is at least $1/K$. Thus the mediator has the cost for the swap fee at least $\frac{f}{(1-f)K}$. So under the nontrivial case that Swap^{m+1} is successfully executed (*i.e.*, $x_{i^*-1} \geq x^*$), we have $U_{i^*-1} \leq -\frac{f}{(1-f)K}$. Since the i^*-th transaction is user's transaction, we have $U_{i^*} = U_{i^*-1} \leq -\frac{f}{(1-f)K}$.

Note that the parameter η is set such that $\texttt{IDL} < \frac{f}{(1-f)K}$. By ($\triangle$), we know that $\psi_{i^*} \leq \texttt{IDL}$. Thus, we conclude that $U_{i^*} + \psi_{i^*} \leq 0$. It is easy to use the same induction above by two cases to show that $U_{m+1+k} + \psi_{m+1+k} \leq U_{i^*} + \psi_{i^*} \leq 0$. Combining with the fact that $\psi(s) \geq 0$, we conclude that $U_{m+1+k} \leq 0$.

This finishes the proof. □

□