

Maximizing Extractable Value from Automated Market Makers

Massimo Bartoletti¹, James Hsin-yu Chiang², Alberto Lluçh Lafuente²

¹ Università degli Studi di Cagliari, Cagliari, Italy

² Technical University of Denmark, DTU Compute, Copenhagen, Denmark

Abstract. Automated Market Makers (AMMs) are decentralized applications that allow users to exchange crypto-tokens without the need for a matching exchange order. AMMs are one of the most successful DeFi use cases: indeed, major AMM platforms process a daily volume of transactions worth USD billions. Despite their popularity, AMMs are well-known to suffer from transaction-ordering issues: adversaries can influence the ordering of user transactions, and possibly front-run them with their own, to extract value from AMMs, to the detriment of users. We devise an effective procedure to construct a strategy through which an adversary can *maximize* the value extracted from user transactions.

Keywords: miner extractable value, front-running, decentralized finance

1 Introduction

Decentralized finance (DeFi) is emerging as an alternative to traditional finance, boosted by blockchains, crypto-tokens and smart contracts [17]. *Automated Market Makers (AMMs)* — one of the main DeFi applications — allow users to exchange crypto-tokens without the need to find another party wanting to participate in the exchange. Major AMM platforms like e.g. Uniswap, Curve Finance, and SushiSwap, hold dozens of billions of USD and process hundreds of millions worth of transactions daily [8,1,5].

AMMs are sensitive to *transaction-ordering attacks*, where adversaries who can influence the ordering of transactions in the blockchain exploit this power to *extract value* from user transactions [13,15,16,20]. We illustrate this kind of attacks through a minimal example. Assume a Uniswap-like AMM holding 100 units of a crypto-token τ_0 and 100 units of another token τ_1 , and assume that both tokens have the same price in the reference currency (say, USD 1,000). Now, suppose that user **A** wants to swap 20 units of τ_0 in her wallet for at least 15 units of τ_1 . This requires to append to the blockchain a transaction of the form **A** : $\text{swap}^0(20 : \tau_0, 15 : \tau_1)$, where the prefix **A** indicates the wallet involved in the transaction, swap is the called AMM function, and the superscript 0 indicates the swap direction, i.e. deposit $20 : \tau_0$ to receive back at least $15 : \tau_1$ (a superscript 1 would indicate the opposite direction). In a *constant-product* AMM platform

like Uniswap, the actual amount of τ_1 transferred to **A** must be such that the product between the AMM reserves remains constant before and after a swap.

Now, suppose that an adversary **M** (possibly a miner) observes **A**'s transaction in the txpool, and appends to the blockchain the following *sandwich*:

M : $\text{swap}^0(5.9 : \tau_0, 5.5 : \tau_1)$ **A** : $\text{swap}^0(20 : \tau_0, 15 : \tau_1)$ **M** : $\text{swap}^1(25.9 : \tau_0, 20.6 : \tau_1)$

where the last transaction is in the opposite direction, i.e. **M** sends $20.6 : \tau_1$ to receive at least $25.9 : \tau_0$. As a result, **A** only yields the *minimum* amount of $15 : \tau_1$ in return for $20 : \tau_0$. This implies that USD 5,000 have been gained by **M** and lost by **A**. This has been called *Miner Extractable Value* (MEV) [13].

Recent works study this and other kinds of attacks to AMMs [13,16,19,20]: however, all these approaches are preeminently *empirical*, as they focus on the definition of heuristics to extract value from AMMs, and on their evaluation in the wild. To the best of our knowledge, a general solution to obtain *optimal* MEV is still missing, even in the special case of constant-product AMMs.

To exemplify a case where prior approaches fail to extract optimal MEV, consider the following set of user transactions, containing a swap of τ_0 for τ_1 , a deposit of units of τ_0 and τ_1 , and a redeem of units of minted (liquidity) tokens:

{ **A** : $\text{swap}^0(40 : \tau_0, 35 : \tau_1)$, **A** : $\text{dep}(30 : \tau_0, 40 : \tau_1)$, **A** : $\text{rdm}(10 : (\tau_0, \tau_1))$ }

Here, both the swap and the dep transactions would be rejected. For instance, the constant-product invariant dictates that $40 : \tau_0$ sent by the user swap in the initial AMM state ($100 : \tau_0, 100 : \tau_1$) will return exactly $28.6 : \tau_1$; since the swap transaction requires $35 : \tau_1$, it would be discarded. The known heuristics here fail to extract any value. Even considering only the swap , the sandwich would not be profitable for **M**, since it requires the *same* direction for **M**'s and **A**'s swap (offer τ_0 to obtain τ_1), making **A**'s swap not enabled. Further, the known heuristics only operate on swap actions, neglecting user deposits and redeems. This paper proposes a layered construction to extract the *maximum value* from all user transactions, through a multi-layer sandwich that we call *Dagwood sandwich*. In our example, **M**'s strategy would be to fire the following three-layer sandwich:

M : $\text{swap}^1(11 : \tau_0, 13 : \tau_1)$ **A** : $\text{swap}^0(40 : \tau_0, 35 : \tau_1)$
M : $\text{swap}^1(42 : \tau_0, 38 : \tau_1)$ **A** : $\text{dep}(30 : \tau_0, 40 : \tau_1)$
M : $\text{swap}^0(18 : \tau_0, 21 : \tau_1)$

The first transaction is a swap in the opposite direction (i.e., pay τ_1 to get τ_0) w.r.t. the subsequent user swap , unlike in the classical sandwich heuristic. **M**'s second swap enables **A**'s deposit; the final swap is an arbitrage move [9]. The user redeem is dropped, since it would negatively contribute to **M**'s profit. By firing the transaction sequence above, **M** can extract approx. USD 5,700 from **A**, improving over swap -only attacks, that would only extract USD 5,000.

Contributions To the best of our knowledge, this work is the first to formalise the *MEV game* for AMMs (Section 3), and the first to effectively construct

optimal solutions which attack all types of transactions supported by constant-product AMMs (Section 4). We discuss in Section 6 the applicability of our technique in the wild. The proofs of our statements are in Appendix A.

2 Automated Market Makers

We assume a set \mathbb{T}_0 of *atomic token types* (ranged over by τ, τ', \dots), representing native cryptocurrencies and application-specific tokens. We denote by $\mathbb{T}_1 = \mathbb{T}_0 \times \mathbb{T}_0$ the set of *minted token types*, representing shares in AMMs. In our model, tokens are *fungible*, i.e. individual units of the same type are interchangeable. In particular, amounts of tokens of the same type can be split into smaller parts, and two amounts of tokens of the same type can be joined. We use v, v', r, r' to range over nonnegative real numbers (\mathbb{R}_0^+), and we write $r : \tau$ to denote r units of token type $\tau \in \mathbb{T} = \mathbb{T}_0 \cup \mathbb{T}_1$.

We model the *wallet* of a user A as a term $A[\sigma]$, where the partial map $\sigma \in \mathbb{T} \rightarrow \mathbb{R}_0^+$ represents A 's token holdings, and write $A[_]$ if the wallet balance is clear from context. We denote with $\text{dom}(\sigma)$ the domain of σ . An **AMM** is a pair of the form $(r_0 : \tau_0, r_1 : \tau_1)$, representing the fact that the AMM is holding r_0 units of τ_0 and r_1 units of τ_1 . We denote by $\text{res}_{\tau_0, \tau_1}(\Gamma)$ the *reserves* of τ_0 and τ_1 in Γ , i.e. $\text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1)$ if $(r_0 : \tau_0, r_1 : \tau_1)$ is in Γ .

A **state** is a composition of wallets and AMMs, represented as a term:

$$A_1[\sigma_1] \mid \dots \mid A_n[\sigma_n] \mid (r_1 : \tau_1, r'_1 : \tau'_1) \mid \dots \mid (r_k : \tau_k, r'_k : \tau'_k)$$

where: (i) all A_i are distinct, (ii) the token types in an AMM are *distinct*, and (iii) distinct AMMs cannot hold exactly the same token types. Note that two AMMs can have a common token type τ , as in $(r_1 : \tau_1, r : \tau) \mid (r' : \tau, r_2 : \tau_2)$, thus enabling indirect trades between token pairs not directly provided by any AMM. We use Γ, Γ', \dots to range over states. For a base term Q (either wallet or AMM), we write $Q \in \Gamma$ when $\Gamma = Q \mid \Gamma'$, for some Γ' , where we assume that two states are equivalent when they contain the same base terms.

We define the **supply** of a token type τ in a state Γ as the sum of the balances of τ in all the wallets and the AMMs occurring in Γ . Formally:

$$\text{sply}_\tau(A[\sigma]) = \begin{cases} \sigma(\tau) & \text{if } \tau \in \text{dom}(\sigma) \\ 0 & \text{otherwise} \end{cases} \quad \text{sply}_\tau(r_0 : \tau_0, r_1 : \tau_1) = \begin{cases} r_i & \text{if } \tau = \tau_i \\ 0 & \text{otherwise} \end{cases}$$

and the supply of τ in $\Gamma \mid \Gamma'$ is the summation $\text{sply}_\tau(\Gamma) + \text{sply}_\tau(\Gamma')$.

We model the interaction between users and AMMs as a transition system between states. A transition $\Gamma \xrightarrow{\mathbb{T}} \Gamma'$ represents the evolution of the state Γ into Γ' upon the execution of the **transaction** \mathbb{T} . The possible transactions are:

- $A : \text{dep}(v_0 : \tau_0, v_1 : \tau_1)$, which allows A to **deposit** $v_0 : \tau_0$ and $v_1 : \tau_1$ to an AMM, receiving in return units of the minted token (τ_0, τ_1) .
- $A : \text{swap}^d(v_0 : \tau_0, v_1 : \tau_1)$ with $d \in \{0, 1\}$, which allows A to **swap** tokens, i.e. transfer $v_d : \tau_d$ to an AMM, and receive in return *at least* $v_{1-d} : \tau_{1-d}$.

- $\mathbf{A} : \text{rdm}(v : \tau)$, which allows to \mathbf{A} *redeem* v units of minted token $\tau = (\tau_0, \tau_1)$ from an AMM, receiving in return units of the atomic tokens τ_0 and τ_1 .

We now formalise the one-step relation $\dashv\rightarrow$ through rewriting rules, inspired by [9]. We use the standard notation $\sigma\{v/x\}$ to update a partial map σ at point x : namely, $\sigma\{v/x\}(x) = v$, while $\sigma\{v/x\}(y) = \sigma(y)$ for $y \neq x$. For a partial map $\sigma \in \mathbb{T} \rightarrow \mathbb{R}_0^+$, a token type $\tau \in \mathbb{T}$ and a partial operation $\circ \in \mathbb{R}_0^+ \times \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$, we define the partial map $\sigma \circ v : \tau$ (updating τ 's balance in σ by v) as follows:

$$\sigma \circ v : \tau = \begin{cases} \sigma\{\sigma(\tau) \circ v/\tau\} & \text{if } \tau \in \text{dom } \sigma \text{ and } \sigma(\tau) \circ v \in \mathbb{R}_0^+ \\ \sigma\{v/\tau\} & \text{if } \tau \notin \text{dom } \sigma \end{cases}$$

Deposit Any user can create an AMM for a token pair (τ_0, τ_1) , provided that such an AMM is not already present in the state. This is achieved by the transaction $\mathbf{A} : \text{dep}(v_0 : \tau_0, v_1 : \tau_1)$, through which \mathbf{A} transfers $v_0 : \tau_0$ and $v_1 : \tau_1$ to the new AMM. In return, \mathbf{A} receives an amount of units of a new token type (τ_0, τ_1) , which is minted by the AMM. We formalise this behaviour by the rule:

$$\frac{\sigma(\tau_i) \geq v_i > 0 \ (i \in \{0, 1\}) \quad \tau_0 \neq \tau_1 \quad \tau_0, \tau_1 \in \mathbb{T}_0 \quad (- : \tau_0, - : \tau_1), (- : \tau_1, - : \tau_0) \notin \Gamma}{\mathbf{A}[\sigma] \mid \Gamma \xrightarrow{\mathbf{A}:\text{dep}(v_0:\tau_0, v_1:\tau_1)} \mathbf{A}[\sigma - v_0 : \tau_0 - v_1 : \tau_1 + v_0 : (\tau_0, \tau_1)] \mid (v_0 : \tau_0, v_1 : \tau_1) \mid \Gamma} \text{[DEP0]}$$

Once an AMM is created, any user can deposit tokens into it, as long as doing so preserves the ratio of the token holdings in the AMM. When a user deposits $v_0 : \tau_0$ and $v_1 : \tau_1$ to an existing AMM, it receives in return an amount of minted tokens of type (τ_0, τ_1) . This amount is the ratio between the deposited amount v_0 and the redeem rate of (τ_0, τ_1) in the current state Γ . This redeem rate is the ratio between the amount r_0 of τ_0 stored in the AMM, and the total supply $\text{sply}_{(\tau_0, \tau_1)}(\Gamma)$ of the minted token in the state.

$$\frac{\sigma(\tau_i) \geq v_i > 0 \ (i \in \{0, 1\}) \quad r_1 v_0 = r_0 v_1 \quad v = \frac{v_0}{r_0} \cdot \text{sply}_{(\tau_0, \tau_1)}(\Gamma)}{\Gamma = \mathbf{A}[\sigma] \mid (r_0 : \tau_0, r_1 : \tau_1) \mid \Gamma' \xrightarrow{\mathbf{A}:\text{dep}(v_0:\tau_0, v_1:\tau_1)} \mathbf{A}[\sigma - v_0 : \tau_0 - v_1 : \tau_1 + v : (\tau_0, \tau_1)] \mid (r_0 + v_0 : \tau_0, r_1 + v_1 : \tau_1) \mid \Gamma'} \text{[DEP]}$$

The premise $r_1 v_0 = r_0 v_1$ ensures that the ratio between the reserves of τ_0 and τ_1 in the AMM is preserved, i.e. $r_1 + v_1 / r_0 + v_0 = r_1 / r_0$.

Swap Any user \mathbf{A} can swap units of τ_0 in her wallet for units of τ_1 in an AMM $(r_0 : \tau_0, r_1 : \tau_1)$, or *vice versa* swap units of τ_1 in the wallet for units of τ_0 in the AMM. This is achieved by the transaction $\mathbf{A} : \text{swap}^d(v_0 : \tau_0, v_1 : \tau_1)$, where $d \in \{0, 1\}$ is the *swap direction*. If $d = 0$ (“left” swap), then v_0 is the amount of τ_0 transferred from \mathbf{A} 's wallet to the AMM, while v_1 is a *lower bound* on the amount of τ_1 that \mathbf{A} will receive in return. Conversely, if $d = 1$ (“right” swap), then v_1 is the amount of τ_1 transferred from \mathbf{A} 's wallet, and v_0 is a lower bound on the received amount of τ_0 . The actual amount v of received units of τ_{1-d} must satisfy the *constant-product invariant* [18], as in Uniswap [7], SushiSwap [6] and other common AMMs implementations:

$$r_0 \cdot r_1 = (r_d + v_d) \cdot (r_{1-d} - v)$$

Formally, for $d \in \{0, 1\}$ we define:

$$\frac{\sigma(\tau_d) \geq v_d > 0 \quad v = \frac{\tau_{1-d} \cdot v_d}{r_d + v_d} \quad 0 < v_{1-d} \leq v}{\text{A}[\sigma] \mid (r_0 : \tau_0, r_1 : \tau_1) \mid \Gamma \xrightarrow{\text{A:swap}^d(v_0 : \tau_0, v_1 : \tau_1)} \text{A}[\sigma - v_d : \tau_d + v : \tau_{1-d}] \mid (r_0 : \tau_0, r_1 : \tau_1) + v_d : \tau_d - v : \tau_{1-d} \mid \Gamma} \text{[SWAP]}$$

where we define the update of the units of τ in an AMM, for $\circ \in \{+, -\}$, as:

$$(r_0 : \tau_0, r_1 : \tau_1) \circ v : \tau = \begin{cases} (r_0 \circ v : \tau_0, r_1 : \tau_1) & \text{if } \tau = \tau_0 \text{ and } r_0 \circ v \in \mathbb{R}_0^+ \\ (r_0 : \tau_0, r_1 \circ v : \tau_1) & \text{if } \tau = \tau_1 \text{ and } r_1 \circ v \in \mathbb{R}_0^+ \end{cases}$$

Redeem Users can redeem units of a minted token (τ_0, τ_1) for units of the underlying atomic tokens τ_0 and τ_1 . Each unit of (τ_0, τ_1) can be redeemed for equal fractions of τ_0 and τ_1 remaining in the AMM:

$$\frac{\sigma(\tau_0, \tau_1) \geq v > 0 \quad v_0 = v \frac{\tau_0}{\text{sply}(\tau_0, \tau_1)(\Gamma)} \quad v_1 = v \frac{\tau_1}{\text{sply}(\tau_0, \tau_1)(\Gamma)}}{\Gamma = \text{A}[\sigma] \mid (r_0 : \tau_0, r_1 : \tau_1) \mid \Gamma' \xrightarrow{\text{A:rdm}(v : (\tau_0, \tau_1))} \text{A}[\sigma + v_0 : \tau_0 + v_1 : \tau_1 - v : (\tau_0, \tau_1)] \mid (r_0 - v_0 : \tau_0, r_1 - v_1 : \tau_1) \mid \Gamma'} \text{[RDM]}$$

A key property of the transition system is *determinism*, i.e. if $\Gamma \xrightarrow{\mathbb{T}} \Gamma'$ and $\Gamma \xrightarrow{\mathbb{T}} \Gamma''$, then the states Γ' and Γ'' are equivalent. We denote with $\text{type}(\mathbb{T})$ the type of \mathbb{T} (i.e., **dep**, **swap**, **rdm**), and with $\text{usr}(\mathbb{T})$ the user issuing \mathbb{T} . For a sequence of transactions $\lambda = \mathbb{T}_1 \cdots \mathbb{T}_n$, we write $\Gamma \xrightarrow{\lambda} \Gamma'$ whenever there exist intermediate states $\Gamma_1, \dots, \Gamma_{n-1}$ such that $\Gamma \xrightarrow{\mathbb{T}_1} \Gamma_1 \xrightarrow{\mathbb{T}_2} \dots \xrightarrow{\mathbb{T}_{n-1}} \Gamma_{n-1} \xrightarrow{\mathbb{T}_n} \Gamma'$. When this happens, we say that λ is *enabled* in Γ , or just $\Gamma \xrightarrow{\lambda}$. A state Γ is *reachable* if there exist some Γ_0 only containing wallets with atomic tokens and some λ such that $\Gamma_0 \xrightarrow{\lambda} \Gamma$.

3 The MEV game

The model in the previous section defines how the state of AMMs and wallets evolves upon a sequence of transactions, but it does not specify how this sequence is formed. We specify this as a single-player, single-round game where the only player is an adversary **M** who attempts to maximize its MEV. Accordingly, we call this the **MEV game**. The *initial state* of the game is given by a reachable state Γ (not including **M**'s wallet) and by a finite multiset \mathcal{X} of user transactions, representing the pool of pending transactions (also called **txpool**). The *moves* of **M** are pairs (σ, λ) , where σ is **M**'s initial balance, and λ is a sequence formed by (part of) the transactions in \mathcal{X} , and by any number of **M**'s transactions. We require that the sequence λ in a move is enabled in Γ . The MEV game assumes the following (see Section 6 for a discussion thereof):

1. Users balances in Γ are sufficiently high to not interfere with the validity of any specific ordering of actions in \mathcal{X} .
2. The balance σ of **M** does not include minted tokens.

3. The length of the sequence λ is unbounded.
4. Prices of atomic tokens are fixed throughout the game execution.

Besides the above, some further assumptions are implied by our AMM model:

5. AMMs only hold atomic tokens (this is a consequence of $[\text{DEP0}]$).
6. Swap actions do not require fees (this is a consequence of $[\text{SWAP}]$).
7. There are no transaction fees.
8. Interval constraints on received token amounts are modelled in swaps only.

A *solution* to the game is a move that maximizes M 's *gain*, i.e. the change in M 's net worth after performing the sequence λ from Γ . Intuitively, the net worth of a user is the overall *value* of tokens in her wallet. To define it, we need to associate a *price* to each token. We assume that the prices of atomic tokens are given by an oracle $P \in \mathbb{T}_0 \rightarrow \mathbb{R}_0^+$: naturally, the MEV game solution will need to be recomputed should the price of atomic tokens be updated. The price $P_\Gamma(\tau_0, \tau_1)$ of a minted token (τ_0, τ_1) in a state Γ is defined as follows:

$$P_\Gamma(\tau_0, \tau_1) = \frac{r_0 \cdot P(\tau_0) + r_1 \cdot P(\tau_1)}{\text{sply}_{(\tau_0, \tau_1)}(\Gamma)} \quad \text{if } \text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1) \quad (1)$$

Minted tokens are priced such that the net worth of a user is preserved when she deposits or redeems minted tokens in her wallet. We assume that the reserves in an AMM are never reduced to zero in an execution, in order to preserve equality of minted token prices between two states with equal reserves, thereby facilitating proofs and analysis. While our semantics of AMMs allows reserves to be emptied, we note that this does not occur in practice, as it would halt the operation of the respective AMM pair. We define the *net worth* of a user A in a state Γ such that $A[\sigma] \in \Gamma$ as follows:

$$W_A(\Gamma) = \sum_{\tau \in \text{dom}(\sigma)} \sigma(\tau) \cdot P_\Gamma(\tau) \quad (2)$$

and we denote by $G_A(\Gamma, \lambda)$ the *gain* of user A upon performing a sequence of transactions λ enabled in state Γ (if λ is not enabled, the gain is zero):

$$G_A(\Gamma, \lambda) = W_A(\Gamma') - W_A(\Gamma) \quad \text{if } \Gamma \xrightarrow{\lambda} \Gamma' \quad (3)$$

A *rational player* is a player which, for all initial states (Γ, \mathfrak{X}) of the game, always chooses a move (σ, λ) that maximizes the function $G_M(M[x] \mid \Gamma, y)$ on variables x and y . We define the *miner extractable value* in (Γ, \mathfrak{X}) as the gain obtained by a rational player by applying such a solution (σ, λ) , i.e.:

$$\text{MEV}(\Gamma, \mathfrak{X}) = G_M(M[\sigma] \mid \Gamma, \lambda)$$

Lemma 1 states that firing transactions preserves the *global* net worth, i.e. the gains of some users are balanced by equal overall losses of other users.

Lemma 1. $\sum_A G_A(\Gamma, \mathbb{T}) = 0$.

By using a simple inductive argument, we can extend Lemma 1 to sequences of transactions: if $\Gamma \xrightarrow{\lambda} \Gamma'$, then the summation of the gains $G_A(\Gamma, \lambda)$ over all users (including M) is 0. Hence, the MEV game is zero-sum. The following lemma ensures that deposit and redeem actions do not directly affect the net worth of the user who performs them.

Lemma 2. *If $\text{type}(\mathbb{T}) \in \{\text{dep}, \text{rdm}\}$, then $G_{usr(\mathbb{T})}(\Gamma, \mathbb{T}) = 0$.*

Finally, we note that prices of a minted token in two states are equal if the reserve ratio in the two states are as well.

Lemma 3. *Let $\Gamma \xrightarrow{\lambda} \Gamma'$, and let $\text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1)$, $\text{res}_{\tau_0, \tau_1}(\Gamma') = (r'_0, r'_1)$. Then, $P_\Gamma(\tau_0, \tau_1) = P_{\Gamma'}(\tau_0, \tau_1)$ if and only if $r_0/r_1 = r'_0/r'_1$.*

4 Solving the MEV game

By Lemma 1, a move which minimizes the gain of all users but M must maximize M 's gain, and therefore is a solution to the MEV game. More formally, we have:

Corollary 1. *$G_M(\Gamma, \lambda)$ is maximized iff $G_A(\Gamma, \lambda)$ is minimized for all $A \neq M$.*

The net worth W_A of a user A can be decomposed in two parts: W_A^0 , which accounts for the atomic tokens, and W_A^1 , which accounts for the minted tokens:

$$W_A^0(\Gamma) = \sum_{\tau \in \mathbb{T}_0} \sigma_A(\tau) \cdot P(\tau) \quad W_A^1(\Gamma) = \sum_{\tau \in \mathbb{T}_1} \sigma_A(\tau) \cdot P_\Gamma(\tau) \quad (4)$$

This provides M with two levers to reduce the users' gain: token balances, and the price of minted tokens. To use the first lever, M needs to exploit user actions in the txpool \mathcal{X} of the MEV game. For the second lever, since the prices of atomic tokens ($\tau \in \mathbb{T}_0$) are fixed, M can only influence the price of minted tokens ($\tau \in \mathbb{T}_1$). This can be achieved by performing actions on the respective AMMs.

In the rest of the section we devise an *optimal* strategy to exploit these two levers. Intuitively, our strategy constructs a multi-layer **Dagwood Sandwich**³, containing an **inner layer** for each exploitable user action in \mathcal{X} , which M *front-runs* by a swap transaction to enable it (if necessary), and a **final layer** of swaps by M to minimize the prices of all minted tokens.

The construction of the final layer of the Dagwood sandwich is shown in §4.1, while the construction of the inner layers is presented in §4.2.

4.1 Price minimization

Lemma 4 below states that, in any state, M can minimize the price of a minted token by using a single swap, at most. In particular, this minimization can always be performed in the final layer of the Dagwood sandwich.

³ We name it after Dagwood Bumstead, a comic strip character who is often illustrated while producing enormous multi-layer sandwiches.

Lemma 4. *There exists a function P^{min} such that if $\mathbf{M}[\sigma] \mid \Gamma \rightarrow^* \mathbf{M}[\sigma'] \mid \Gamma'$ then: (i) $P_{\Gamma'}(\tau_0, \tau_1) \geq P_{\Gamma}^{min}(\tau_0, \tau_1)$; (ii) there exist σ'' and λ consisting at most of a swap by \mathbf{M} such that $\mathbf{M}[\sigma''] \mid \Gamma' \xrightarrow{\lambda} \mathbf{M}[\cdot] \mid \Gamma''$ and $P_{\Gamma''}(\tau_0, \tau_1) = P_{\Gamma}^{min}(\tau_0, \tau_1)$.*

In order to construct the **swap** transaction which minimizes the price of a minted token (τ_0, τ_1) in Γ , we need some auxiliary definitions. For each swap direction $d \in \{0, 1\}$, we define the **canonical swap values** as:

$$w_d^d(\tau_0, \tau_1, \Gamma) = \sqrt{\frac{P(\tau_{1-d})}{P(\tau_d)} \cdot r_0 \cdot r_1} - r_d \quad w_{1-d}^d(\tau_0, \tau_1, \Gamma) = \frac{r_{1-d} \cdot w_d^d(\tau_0, \tau_1, \Gamma)}{r_d + w_d^d(\tau_0, \tau_1, \Gamma)}$$

Intuitively, w_d^d is the amount of tokens *deposited* in a swap of direction d : it is defined such that, after the swap, the AMM reaches an equilibrium, where the ratio of the AMM reserves is equal to the (inverse) ratio of the token prices. Instead, w_{1-d}^d is the amount of tokens *received* after the swap, i.e. it is the unique value for which the swap invariant is satisfied.

If both $w_0^0(\tau_0, \tau_1, \Gamma) \leq 0$ and $w_1^1(\tau_0, \tau_1, \Gamma) \leq 0$, then the price of the minted token (τ_0, τ_1) is already minimized. Otherwise, if $w_d^d(\tau_0, \tau_1, \Gamma) > 0$ for some d (and there may exist at most one d for which this holds), then we define the **price minimization transaction** $\mathbf{X}^d(\tau_0, \tau_1, \Gamma)$ as:

$$\mathbf{M} : \text{swap}^d(w_d^d(\tau_0, \tau_1, \Gamma) : \tau_0, w_{1-d}^d(\tau_0, \tau_1, \Gamma) : \tau_1) \quad (5)$$

Theorem 1 constructs the final layer of the Dagwood sandwich. We show that this layer is the solution of the MEV game on an empty txpool. This is because if \mathbf{M} cannot leverage user transactions, the solution is just to minimize the price of all minted tokens. The solution is obtained by sequencing price minimization transactions on all AMMs. Since the price of a minted token is a function of the reserves of the corresponding AMM, this can be done in any order.

Theorem 1. *Let $\Gamma = \parallel_{i \in I} (r_{i,0} : \tau_{i,0}, r_{i,1} : \tau_{i,1}) \mid \Gamma_w$, where Γ_w only contains wallets. For all $j \in I$ and $d \in \{0, 1\}$, let $v_j^d = w_d^d(\tau_{j,0}, \tau_{j,1}, \Gamma)$, and let:*

$$\sigma_j = \begin{cases} v_j^d : \tau_{j,d} & \text{if } v_j^d > 0 \\ 0 & \text{if } v_j^0, v_j^1 \leq 0 \end{cases} \quad \lambda_j = \begin{cases} \mathbf{X}^d(\tau_{j,0}, \tau_{j,1}, \Gamma) & \text{if } v_j^d > 0 \\ \varepsilon & \text{if } v_j^0, v_j^1 \leq 0 \end{cases}$$

Then, $(\sigma_1 \cdots \sigma_n, \lambda_1 \cdots \lambda_n)$ is a solution to the game (Γ, \mathbf{X}) for an empty \mathbf{X} .

4.2 Constructing the inner layers

Consider a solution (σ, λ) to the game $(\mathbf{A}[\sigma_{\mathbf{A}}] \mid \Gamma, \mathbf{X})$, and let:

$$\mathbf{M}[\sigma] \mid \mathbf{A}[\sigma_{\mathbf{A}}] \mid \Gamma \xrightarrow{\lambda} \mathbf{M}[\sigma'] \mid \mathbf{A}[\sigma'_{\mathbf{A}}] \mid \Gamma'$$

By decomposing the net worth as in (4), we find that \mathbf{A} 's gain for λ is:

$$\begin{aligned} G_{\mathbf{A}}(\mathbf{M}[\sigma] \mid \mathbf{A}[\sigma_{\mathbf{A}}] \mid \Gamma, \lambda) &= W_{\mathbf{A}}^0(\Gamma') - W_{\mathbf{A}}^0(\Gamma) + W_{\mathbf{A}}^1(\Gamma') - W_{\mathbf{A}}^1(\Gamma) \\ &= \sum_{\tau \in \mathbb{T}_0} (\sigma'_{\mathbf{A}}(\tau) - \sigma_{\mathbf{A}}(\tau)) \cdot P(\tau) + \sum_{\tau \in \mathbb{T}_1} (\sigma'_{\mathbf{A}}(\tau) \cdot P_{\Gamma'}(\tau) - \sigma_{\mathbf{A}}(\tau) \cdot P_{\Gamma}(\tau)) \end{aligned}$$

Since λ is a solution, by Lemma 4 we can replace $P_{\Gamma'}(\tau)$ with $P_{\Gamma}^{min}(\tau)$:

$$= \sum_{\tau \in \mathbb{T}_0} (\sigma'_A(\tau) - \sigma_A(\tau)) \cdot P(\tau) + \sum_{\tau \in \mathbb{T}_1} (\sigma'_A(\tau) \cdot P_{\Gamma}^{min}(\tau) - \sigma_A(\tau) \cdot P_{\Gamma}(\tau)) \quad (6)$$

Note that all token prices in (6) are already defined in state Γ . Thus, A 's gain can be minimized by considering only the effect on the token balance σ'_A , which we can rewrite as $\sigma_A + \Delta_0 + \Delta_1 + \dots$ where Δ_i is the effect on the balance by the i 'th transaction in λ . We will show that Δ_i is *fixed* for any user transaction: the position of the i 'th transaction in λ does not affect its optimality.

The following theorem states that solutions to the MEV game can be constructed incrementally, by layering the local solutions for each individual transaction in the txpool. Intuitively, we choose a transaction \mathbb{T} from \mathcal{X} , we solve the game for $(\Gamma, [\mathbb{T}])$, we compute the state Γ' obtained by executing this solution, and we inductively solve the game in the (Γ', \mathcal{X}') , where \mathcal{X}' is \mathcal{X} minus \mathbb{T} .

Theorem 2. *With respect to the MEV game in (Γ, \mathcal{X}) :*

1. *If \mathcal{X} is empty, the solution is the final layer constructed for $(\Gamma, [])$ in §4.1.*
2. *Otherwise, if $\mathcal{X} = [\mathbb{T}] + \mathcal{X}'$, let (σ, λ) be the inner layer constructed for $(\Gamma, [\mathbb{T}])$, let $\mathbb{M}[\sigma] \mid \Gamma \xrightarrow{\lambda} \mathbb{M}[\cdot] \mid \Gamma'$, and let (σ', λ') be the solution for (Γ', \mathcal{X}') . Then, the solution to (Γ, \mathcal{X}) is $(\sigma + \sigma', \lambda \lambda')$.*

We now describe how to define the inner layers of the Dagwood sandwich, i.e. the base case of the inductive construction given by Theorem 2. Each inner layer includes a user transaction from the txpool, possibly front-run by \mathbb{M} such that executing the layer leads the user's net worth to a local minimum. We define below the construction of these inner layers for each transaction type.

Swap inner layer Swap actions only affect the balance of *atomic* tokens. To minimize the gain of A after a swap, \mathbb{M} must make A receive exactly the *minimum* amount of requested tokens. The effect of the swap on A 's *atomic net worth* is:

$$W_A^0(\Gamma') - W_A^0(\Gamma) = -v_d \cdot P(\tau_d) + v_{1-d} \cdot P(\tau_{1-d}) \quad \text{if } \Gamma \xrightarrow{\mathbb{A}:\text{swap}^d(v_0:\tau_0, v_1:\tau_1)} \Gamma'$$

If the change in A 's atomic net worth is negative, A 's transaction is included in the layer. Although this transaction minimizes A 's atomic net worth, it simultaneously affects the price of the minted token (τ_0, τ_1) . This is not an issue, since the final layer of the Dagwood sandwich minimizes the prices of *all* minted tokens. Thus, the change of minted token prices due to the swap inner layer will *not* affect the user gain in the full Dagwood sandwich, as evident from (6). Note that the amount of tokens exchanged in a swap is chosen by the user, so the actual position of the layer in the Dagwood sandwich is immaterial (Theorem 2).

We now define the transaction used by \mathbb{M} to front-run A 's swap, ensuring that A receives the least amount of tokens from the swap. For $\Gamma = (r_0 : \tau_0, r_1 : \tau_1) \mid \dots$

and $\mathbf{T} = \mathbf{A} : \text{swap}^{d_A}(v_0 : \tau_0, v_1 : \tau_1)$, let the *swap front-run reserves* be:

$$\text{SFr}_{d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \frac{\left| \sqrt{v_0^2 \cdot v_1^2 + 4 \cdot v_0 \cdot v_1 \cdot r_0 \cdot r_1} \right| - v_0 \cdot v_1}{2 \cdot v_{1-d_A}}$$

$$\text{SFr}_{1-d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \frac{r_0 \cdot r_1}{\text{SFr}_{d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T})}$$

These values define the reserves of (τ_0, τ_1) in the state Γ' reached from $\mathbf{M}[\sigma] \mid \Gamma$ with \mathbf{M} 's transaction. Intuitively, if the swap front-run reserves do *not* coincide with the reserves r_0, r_1 in Γ , then \mathbf{M} 's transaction is needed to enable \mathbf{A} 's swap. We define the *swap front-run direction* d_M as:

$$d_M = \begin{cases} d_A & \text{if } \text{SFr}_{d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) > r_{d_A} \\ 1 - d_A & \text{if } \text{SFr}_{1-d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) > r_{1-d_A} \end{cases}$$

We define the *swap front-run values* (i.e., the parameters of \mathbf{M} 's swap) as:

$$\text{SF}w_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \begin{cases} \text{SFr}_{d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) - r_{d_A} & \text{if } d_M = d_A \\ r_{d_A} - \text{SFr}_{d_A}(\tau_0, \tau_1, \Gamma, \mathbf{T}) & \text{if } d_M = 1 - d_A \end{cases} \quad (7)$$

$$\text{SF}w_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \begin{cases} r_{1-d_M} - \text{SFr}_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) & \text{if } d_M = d_A \\ \text{SFr}_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) - r_{1-d_M} & \text{if } d_M = 1 - d_M \end{cases}$$

We combine these values to craft the *swap front-run transaction*:

$$\text{SF}\mathbf{X}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \mathbf{M} : \text{swap}^{d_M}(\text{SF}w_0(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_0, \text{SF}w_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_1)$$

The inner layer is included in the Dagwood sandwich if it reduces \mathbf{A} 's net worth, i.e. if $-v_d \cdot P(\tau_d) + v_{1-d} \cdot P(\tau_{1-d}) < 0$. The swap front-run transaction is omitted if the reserves in Γ coincide with the swap front-run reserves. The balance of \mathbf{M} in the (local) game solution is $\text{SF}w_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_{d_M}$.

Example 1. We recast our first example in §1 as a MEV game, assuming a txpool $\mathcal{X} = \{\mathbf{A} : \text{swap}^0(40 : \tau_0, 35 : \tau_1)\}$. The initial state is $\Gamma = (100 : \tau_0, 100 : \tau_1) \mid \Gamma_w$, where Γ_w is made of user wallets, among which $\mathbf{A}[40 : \tau_0]$, and $P(\tau_0) = P(\tau_1) = 1,000$. We construct the Dagwood sandwich. Since \mathbf{A} 's swap yields a reduction in \mathbf{A} 's atomic net worth, $35 \cdot P(\tau_1) - 40 \cdot P(\tau_0) = -5,000$, then \mathbf{A} 's transaction is included in the inner layer. To check if \mathbf{A} 's swap must be front-run by \mathbf{M} , we first compute the swap front-run reserves:

$$\text{SFr}_0(\tau_0, \tau_1, \mathbf{T}, \Gamma) = \frac{\sqrt{40^2 \cdot 35^2 + 4 \cdot 40 \cdot 35 \cdot 100^2} - 40 \cdot 35}{2 \cdot 35} \approx 88.8$$

$$\text{SFr}_1(\tau_0, \tau_1, \mathbf{T}, \Gamma) = \frac{100^2}{89} \approx 112.7$$

Since these values differ from the reserves in the initial game state, \mathbf{M} must front-run \mathbf{A} 's transaction. The direction d_M of \mathbf{M} 's swap is 1, as $\text{SFr}_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) > r_1$. The swap front-run values (7) are given by:

$$\text{SF}w_0(\tau_0, \tau_1, \Gamma, \mathbf{T}) = 100 - 88.8 \approx 11.2 \quad \text{SF}w_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) = 112.7 - 100 \approx 12.7$$

Therefore, the swap inner layer is made of two transactions:

$$\mathbf{M} : \text{swap}^1(11.2 : \tau_0, 12.7 : \tau_1) \quad \mathbf{A} : \text{swap}^0(40 : \tau_0, 35 : \tau_1)$$

and \mathbf{M} 's balance of the (local) game solution is $12.7 : \tau_1$. To construct the final layer, we consider the state $\Gamma'' = (128.8 : \tau_0, 77.7 : \tau_1) \mid \dots$, shown in Figure 1. In Γ'' , the canonical swap values are given by:

$$w_0^1(\tau_0, \tau_1, \Gamma'') = \frac{128.8 \cdot 22.3}{77.7 + 22.3} \approx 28.7$$

$$w_1^1(\tau_0, \tau_1, \Gamma'') = \sqrt{\frac{1}{1} \cdot 128.8 \cdot 77.7 - 77.7} \approx 22.3$$

Since $w_1^1(\tau_0, \tau_1, \Gamma'') > 1$, the direction d of the price minimization swap is 1. Therefore, the final layer is made of a single swap on the pair (τ_0, τ_1) :

$$\mathbf{M} : \text{swap}^1(28.7 : \tau_0, 22.3 : \tau_1)$$

where \mathbf{M} 's required balance is $22.3 : \tau_1$. Summing up, the Dagwood sandwich is constructed by appending the final layer to the inner layer, and \mathbf{M} 's required balance is $\sigma = 12.7 : \tau_1 + 22.3 : \tau_1 = 35 : \tau_1$. The MEV obtained by \mathbf{M} through the Dagwood sandwich is $(11.2 - 12.7) \cdot 1,000 + (28.7 - 22.3) \cdot 1,000 \approx 5,000$. \square

Deposit inner layer By Lemma 2, deposits preserve the user's net worth. Thus, executing $\mathbf{T} = \mathbf{A} : \text{dep}(v_0 : \tau_0, v_1 : \tau_1)$ in Γ does not bring any gain to \mathbf{A} :

$$G_{\mathbf{A}}(\Gamma, \mathbf{T}) = -v_0 \cdot P(\tau_0) - v_1 \cdot P(\tau_1) + v \cdot P_{\Gamma}(\tau_0, \tau_1) = 0 \quad (8)$$

where v is the amount of minted tokens (τ_0, τ_1) given to \mathbf{A} upon the deposit. By Lemma 4, $P_{\Gamma}(\tau_0, \tau_1) \geq P_{\Gamma}^{\min}(\tau_0, \tau_1)$. By using this inequality in (8), we have:

$$-v_0 \cdot P(\tau_0) - v_1 \cdot P(\tau_1) + v \cdot P_{\Gamma}^{\min}(\tau_0, \tau_1) \leq 0$$

$$\iff v \cdot P_{\Gamma}^{\min}(\tau_0, \tau_1) \leq v_0 \cdot P(\tau_0) + v_1 \cdot P(\tau_1)$$

By (6) it follows that including \mathbf{T} in a game solution λ reduces \mathbf{A} 's net worth, since the decrease of \mathbf{A} 's net worth in atomic tokens is *not* always offset

$$\begin{aligned} & \mathbf{M}[35 : \tau_1] \mid \Gamma = (100 : \tau_0, 100 : \tau_1) \mid \dots \\ & \xrightarrow{\text{SFX}(\tau_0, \tau_1, \Gamma, \mathbf{T})} \mathbf{M}[11.2 : \tau_0, 22.3 : \tau_1] \mid \Gamma' = (88.8 : \tau_0, 112.7 : \tau_1) \mid \dots \\ & \xrightarrow{\mathbf{T} = \mathbf{A} : \text{swap}^0(40 : \tau_0, 35 : \tau_1)} \mathbf{M}[11.2 : \tau_0, 22.3 : \tau_1] \mid \Gamma'' = (128.8 : \tau_0, 77.7 : \tau_1) \mid \dots \\ & \xrightarrow{\times(\tau_0, \tau_1, \Gamma'')} \mathbf{M}[40 : \tau_0, 0 : \tau_1] \mid \Gamma''' = (100 : \tau_0, 100 : \tau_1) \mid \dots \end{aligned}$$

Fig. 1. A Dagwood sandwich exploiting a single user swap.

by the increase of net worth in minted tokens. Additionally, the minted token price $P_\Gamma(\tau_0, \tau_1)$ in (8) when the user deposit occurs is determined by deposit parameters v_0, v_1 alone: let $\Gamma \rightarrow^* \Gamma'$ be such that the given user deposit \mathbf{T} is *enabled* in both Γ and Γ' . By [DEP], this implies $v_0/v_1 = r_0/r_1 = r'_0/r'_1$ where $(r_0, r_1) = \text{res}_{\tau_0, \tau_1}(\Gamma)$ and $(r'_0, r'_1) = \text{res}_{\tau_0, \tau_1}(\Gamma')$. Then, by Lemma 3, $P_\Gamma(\tau_0, \tau_1) = P_{\Gamma'}(\tau_0, \tau_1)$, as the reserve ratios in Γ and Γ' are equal. Thus, the amount of minted tokens v received by the depositing user in (6) is fixed by (v_0, v_1) , and the effect of a deposit inner layer does not depend on its position along the Dagwood sandwich (Theorem 2).

Similarly to the construction of the swap inner layer, \mathbf{M} may need to front-run transaction $\mathbf{T} = \mathbf{A} : \text{dep}(v_0 : \tau_0, v_1 : \tau_1)$ to enable it. For $\Gamma = (r_0 : \tau_0, r_1 : \tau_1) \mid \dots$, we define the **deposit front-run reserves** as:

$$\text{DF}r_0(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \left| \sqrt{v_0/v_1 \cdot r_0 \cdot r_1} \right| \quad \text{DF}r_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \left| \sqrt{v_1/v_0 \cdot r_0 \cdot r_1} \right|$$

which satisfy $\text{DF}r_0(\tau_0, \tau_1, \Gamma, \mathbf{T}) \cdot v_1 = \text{DF}r_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) \cdot v_0$, as required by [DEP]. Given a swap direction d_M , we define the **deposit front-run values** as:

$$\begin{aligned} \text{DF}w_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) &= \text{DF}r_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) - r_{d_M} \\ \text{DF}w_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) &= r_{1-d_M} - \text{DF}r_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) \end{aligned}$$

If $\text{DF}w_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) > 0$ and $\text{DF}w_{1-d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) > 0$ holds for a swap direction d_M , then we define the **deposit front-run transaction** as:

$$\text{DFX}(\tau_0, \tau_1, \Gamma, \mathbf{T}) = \mathbf{M} : \text{swap}^{d_M}(\text{DF}w_0(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_0, \text{DF}w_1(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_1)$$

If the reserve ratio in the initial state does not coincide with the ratio of deposited funds, i.e. $v_0/v_1 \neq r_0/r_1$, then the deposit inner layer is $\text{DFX}(\tau_0, \tau_1, \Gamma, \mathbf{T}) \mathbf{T}$, and the balance required by \mathbf{M} is $\text{DF}w_{d_M}(\tau_0, \tau_1, \Gamma, \mathbf{T}) : \tau_{d_M}$. Otherwise, the deposit inner layer is made just by \mathbf{T} , and the required balance is zero.

Redeem inner layer By Lemma 2, redeem actions preserve the user's net worth, i.e. \mathbf{A} 's gain is zero when firing $\mathbf{T} = \mathbf{A} : \text{rdm}(v : (\tau_0, \tau_1))$ in Γ :

$$G_{\mathbf{A}}(\Gamma, \mathbf{T}) = -v \cdot P_\Gamma(\tau_0, \tau_1) + v_0 \cdot P(\tau_0) + v_1 \cdot P(\tau_1) = 0$$

Unlike for the deposit inner layer, redeem transactions *increase* the users' gain when executed in the game solution. This is apparent when substituting in the above equation $P_\Gamma(\tau_0, \tau_1) = P_\Gamma^{\text{min}}(\tau_0, \tau_1)$ (as per Lemma 4) to express the user gain *contribution* (6) of the redeem action.

$$-v \cdot P_\Gamma^{\text{min}}(\tau_0, \tau_1) + v_0 \cdot P(\tau_0) + v_1 \cdot P(\tau_1) \geq 0$$

Therefore, user redeem actions always *reduce* \mathbf{M} 's gain, and so they are *not* included in the solution. Therefore, the redeem inner layer is always empty.

$$\begin{aligned}
& \mathbf{M}[18 : \tau_0, 50.5 : \tau_1] \mid \Gamma = (100 : \tau_0, 100 : \tau_1) \mid \dots \\
& \xrightarrow{\text{SFX}(\tau_0, \tau_1, \Gamma, \mathbf{T})} \mathbf{M}[29.3 : \tau_0, 37.8 : \tau_1] \mid \Gamma' = (88.8 : \tau_0, 112.7 : \tau_1) \mid \dots \\
& \xrightarrow{\mathbf{T}=\mathbf{A}:\text{swap}^0(40:\tau_0,35:\tau_1)} \mathbf{M}[29.3 : \tau_0, 37.8 : \tau_1] \mid \Gamma'' = (128.8 : \tau_0, 77.7 : \tau_1) \mid \dots \\
& \xrightarrow{\text{DFX}(\tau_0, \tau_1, \Gamma'', \mathbf{T}')} \mathbf{M}[71.4 : \tau_0, 0 : \tau_1] \mid \Gamma''' = (86.6 : \tau_0, 115.5 : \tau_1) \mid \dots \\
& \xrightarrow{\mathbf{T}'=\mathbf{A}:\text{dep}(30:\tau_0,40:\tau_1)} \mathbf{M}[71.4 : \tau_0, 0 : \tau_1] \mid \Gamma'''' = (116.6 : \tau_0, 155.5 : \tau_1) \mid \dots \\
& \xrightarrow{\mathbf{X}(\tau_0, \tau_1, \Gamma''''')} \mathbf{M}[53.4 : \tau_0, 20.8 : \tau_1] \mid (134.6 : \tau_0, 134.6 : \tau_1) \mid \dots
\end{aligned}$$

Fig. 2. A Dagwood sandwich exploiting a user swap, deposit and redeem (dropped).

Example 2. We now recast the full example in Section 1 as a MEV game, considering all three user transactions in the txpool:

$$\mathcal{X} = \{ \mathbf{A} : \text{swap}^0(40 : \tau_0, 35 : \tau_1), \mathbf{A} : \text{dep}(30 : \tau_1, 40 : \tau_1), \mathbf{A} : \text{rdm}(10 : (\tau_0, \tau_1)) \}$$

The game solution is shown in Figure 2: note that we can reuse the swap inner layer from Example 1, since the initial state and user swap action are identical. Thus, we continue by constructing the deposit inner layer for user deposit \mathbf{T}' in state $\Gamma'' = (128.8 : \tau_0, 77.7 : \tau_1)$. Here, the deposit front-run reserves are:

$$\begin{aligned}
\text{DFr}_0(\tau_0, \tau_1, \Gamma'', \mathbf{T}') &= \left| \sqrt{30/40 \cdot 128.8 \cdot 77.7} \right| = 86.6 \\
\text{DFr}_1(\tau_0, \tau_1, \Gamma'', \mathbf{T}') &= \left| \sqrt{40/30 \cdot 128.8 \cdot 77.7} \right| = 115.5
\end{aligned}$$

Since the ratio of the deposit front-run reserves does not coincide with the reserve ratio in Γ'' ($86.6/115.5 \neq 128.8/77.7$), the deposit front-run by \mathbf{M} is necessary to enable the user deposit action. By choosing a swap direction $d_{\mathbf{M}} = 1$, we obtain the positive deposit front-run values, which confirm the choice of the direction:

$$\text{DFw}_0(\tau_0, \tau_1, \Gamma'', \mathbf{T}') = 128.8 - 86.6 \approx 42.2 \quad \text{DFw}_1(\tau_0, \tau_1, \Gamma'', \mathbf{T}') = 115.5 - 77.7 \approx 37.8$$

Therefore, \mathbf{M} 's deposit front-run transaction is:

$$\text{DFX}(\tau_0, \tau_1, \Gamma'', \mathbf{T}') = \mathbf{M} : \text{swap}^1(42.2 : \tau_0, 37.8 : \tau_1)$$

which requires a balance $\sigma(\tau_1) \geq 37.8$. The deposit inner layer is obtained by prepending this transaction to \mathbf{A} 's deposit. The redeem inner layer is empty, as shown before. By (5), the final layer to minimize the price of minted tokens is:

$$\mathbf{M} : \text{swap}^1(18.0 : \tau_0, 20.8 : \tau_1)$$

Summing up, the full Dagwood sandwich (see also Figure 2) is:

$$\text{SFX}(\tau_0, \tau_1, \Gamma, \mathbf{T}) \quad \mathbf{T} \quad \text{DFX}(\tau_0, \tau_1, \Gamma'', \mathbf{T}') \quad \mathbf{T}' \quad \mathbf{X}(\tau_0, \tau_1, \Gamma''''')$$

which requires an initial balance $\sigma = \{18.0 : \tau_0, 12.7 + 37.8 : \tau_1\}$ by M . By inspection of the Dagwood sandwich execution in Figure 2, it can be seen that the miner has obtained a gain of approximately 5,700. \square

5 Related work

Daian *et al.* [13] study the effect of transaction reordering obtained through *priority gas auctions*. These are games between users who compete to include a bundle of transactions in the next block, bidding on transaction fees to incentivize miners to include their own bundle. Notably, [13] finds empirical evidence of the fact that the gain derived from transaction reorderings in decentralized exchanges (DEX) exceeds the gain given by block rewards and transaction fees in Ethereum. The same work also proposes a game model of priority gas auctions, showing a Nash equilibrium for players to take turns bidding, compatibly with behavior observed in the wild on Ethereum. Our mining game differs from that in [13], since we assume a greedy adversary wanting to maximize its gain at the expense of all the other users, exploiting arbitrages on AMMs.

Zhou *et al.* [20] provide a theoretical framework to study the front-running on AMMs. Two sandwich heuristics are studied: the *front-run & back-run swap* sandwich, and the novel *front-run redeem & back-run swap and deposit*. The swap semantics used in [20] is simplified, compared to ours, since no minimum amount of received tokens is enforced by the AMM, users only perform swaps and hold no minted tokens (depositing and swapping agents are decoupled). Further, extractable value from arbitrage is considered separately. In comparison, we emphasize that we propose a solution to attack all main user action types offered by leading AMMs, thereby extracting value from user submitted swaps and deposits. Our model also accurately model minted tokens: their value is dynamically affected by miner and user swaps during the execution of the attack. Thus, our game solution extracts the maximum value in a more concrete setting, considering the victim transactions of both aforementioned attacks in [20], and leaving no arbitrage opportunities unexploited.

More general ordering and injection of transactions by a rational agent is generally referred to as *front-running*. Eskandari *et al.* [15] provide a taxonomy for various front-running attacks in blockchain applications and networks. This taxonomy is expanded in [16] with liquidations, sandwich attacks and arbitrage actions between DEX.

Some works investigate the problem of detecting front-running attacks on public blockchains. For example, in [16], Qin *et al.* introduce front-running detection heuristics which are deployed to empirically study the presence of such attacks on public DeFi applications. On the other hand, various fair ordering schemes have been proposed to mitigate front-running or exploitation of miner-extractable value. However, simple commit-and-reveal schemes still leak information such as account balances. Breidenbach *et al.* [11] propose “submarine commitments”, which rely on k-anonymity to prevent any leaks from user commitments. Baum *et al.* [10] introduce a order-book based DEX which delegates

the matching of orders to an out-sourced, off-chain multi-party computation committee. Private user orders are not revealed to other participants, such that no front-running can occur in each privately-computed order matching round. Ciampi *et al.* [12] introduce a market maker protocol in which the strictly sequential trade history between an off-chain market maker and traders are verifiable as a hash-chain. Any subsequent reordering by the AMM is publicly provable: collateral from the market maker incentivizes honest, fair-ordering behaviour. Such work aims to provide alternative, front-running resistant designs with AMM-like functionality. In contrast, our work is intended to formalize the behaviour of current, mainstream AMMs in the presence of a rational adversary.

The DeFi community is developing tools to enable agents to extract value from smart contracts: e.g., flashbots [2] is a project aiming to develop Ethereum implementations which support transaction bundles: Rather than front-running individual transactions by adjusting their fees, an agent can communicate a sequence or *bundle* of transactions to the miner, asking its inclusion in the next block. Our game solutions could be implemented to solve for such bundles.

6 Conclusions

We have addressed the problem of adversaries extracting value from AMMs interactions to the detriment of users. We have constructed an *optimal* strategy that adversaries can use to extract value from AMMs, focussing on the widespread class of constant-product AMMs. Our results apply to any adversary with the power to reorder, drop or insert transactions: besides miners, this includes *roll-up aggregators*, like e.g. Optimism and StarkWare [3,4]. Notably, our work shows that it is possible to extract value from *all* types of AMM transactions, while previous works focus on extracting value from token swaps, only.

In practice, value is also extracted from AMMs by colluding mining and non-mining agents: for the Ethereum blockchain, agents can send *transaction bundles* [2] to mining pools for block inclusion, in return for a fee. Our technique naturally applies to this setting, where the actions of the miner are simply replaced by actions by the agent submitting the transaction bundle.

We now discuss the simplifying assumptions (1-8) listed in Section 3. (1) User balances do not limit the order in which transactions in the *txpool* can be executed. In practice, in some cases it would be possible to perform a sequence of actions by exploiting the funds received from previous actions. We leave ordering constraints imposed by limited wallet balances for future work. (2) The adversary holds no minted tokens prior to executing the game solution. Yet, the adversary can exploit an (unbounded) initial balance of atomic tokens to acquire minted tokens as part of the game solution by performing deposits. The optimality of the *Dagwood sandwich* illustrates that this is not necessary. (3) The size of the *Dagwood sandwich* is unbounded. In practice, a typical block of transactions will include other transactions besides those directed to AMMs, and so the adversary can find enough space for its sandwiches by dropping non-AMM transactions. During times of block-congestion, a constraint on the length of the *Dagwood*

sandwich will apply: we conjecture that solving such an optimization is NP-hard, and regard this as an relevant question for future work. (4) Prices of atomic tokens are fixed for the duration of the game: the Dagwood sandwich will need to be recomputed should prices change. (5) AMMs only hold atomic tokens. This is common in practice, but we note that extending the mining game to account for arbitrary nesting of minted tokens by AMM pairs is an interesting direction of future research. (6) No AMM swap fees and (7) no transaction fees are modelled: the adversary’s gain resulting from the Dagwood sandwich is an upper bound to profitability as fees tend to zero. Yet, fees affect this gain, so they should be taken into account to construct an optimal strategy. Furthermore, transaction fees may make it convenient for a miner to include user redeem transactions in the sandwich, while these are never exploited by our strategy. (8) Besides fees, we abstract from the intervals that users can express to constrain the amount of tokens received upon deposits and redemptions (we only model these constraints for swaps). This is left for future work.

In this paper we have considered AMMs which implement the constant-product swap invariant, like e.g. Uniswap and SushiSwap. A relevant research question is how to solve the MEV game under different swap invariants, e.g. those used by Curve Finance and SushiSwap. Uniform frameworks which address this problem have been proposed in [14,9] where swap invariants are abstracted as functions subject to a given set of constraints.

References

1. Curve statistics (2020), <https://www.curve.fi/dailystats>
2. Flashbots (2021), <https://github.com/flashbots/pm>
3. Optimism website (2021), <https://optimism.io/>
4. Starkware website (2021), <https://starkware.co/>
5. SushiSwap statistics (2021), <https://analytics.sushi.com/>
6. SushiSwap token pair implementation (2021), <https://github.com/sushiswap/sushiswap/blob/94ea7712daaa13155dfab9786aacf69e24390147/contracts/uniswapv2/UniswapV2Pair.sol>
7. Uniswap token pair implementation (2021), <https://github.com/Uniswap/uniswap-v2-core/blob/4dd59067c76dea4a0e8e4bfdda41877a6b16dedc/contracts/UniswapV2Pair.sol>
8. Uniswap V2 statistics (2021), <https://v2.info.uniswap.org/>
9. Bartoletti, M., Chiang, J.H., Lluch-Lafuente, A.: A theory of Automated Market Makers in DeFi. In: Coordination Models and Languages (COORDINATION). LNCS, vol. 12717, pp. 168–187. Springer (2021). https://doi.org/10.1007/978-3-030-78142-2_11
10. Baum, C., David, B., Frederiksen, T.K.: P2DEX: privacy-preserving decentralized cryptocurrency exchange. In: Applied Cryptography and Network Security (ACNS). LNCS, vol. 12726, pp. 163–194. Springer (2021). https://doi.org/10.1007/978-3-030-78372-3_7
11. Breidenbach, L., Daian, P., Tramèr, F., Juels, A.: Enter the Hydra: Towards principled bug bounties and exploit-resistant smart contracts. In: USENIX Security Symposium. pp. 1335–1352. USENIX Association (2019)

12. Ciampi, M., Ishaq, M., Magdon-Ismail, M., Ostrovsky, R., Zikas, V.: FairMM: A fast and frontrunning-resistant crypto market-maker. Cryptology ePrint Archive, Report 2021/609 (2021), <https://eprint.iacr.org/2021/609>
13. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: IEEE Symp. on Security and Privacy. pp. 910–927. IEEE (2020). <https://doi.org/10.1109/SP40000.2020.00040>
14. Engel, D., Herlihy, M.: Composing Networks of Automated Market Makers. In: Advances in Financial Technologies (AFT). p. 15–28. ACM (2021). <https://doi.org/10.1145/3479722.3480987>
15. Eskandari, S., Moosavi, S., Clark, J.: SoK: Transparent Dishonesty: Front-Running Attacks on Blockchain. In: Financial Cryptography. pp. 170–189. Springer (2020). https://doi.org/10.1007/978-3-030-43725-1_13
16. Qin, K., Zhou, L., Gervais, A.: Quantifying blockchain extractable value: How dark is the forest? (2021), <https://arxiv.org/abs/2101.05511>
17. Werner, S.M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., Knottenbelt, W.J.: SoK: Decentralized finance (DeFi). CoRR **abs/2101.08778** (2021)
18. Zhang, Y., Chen, X., Park, D.: Formal specification of constant product market maker model & implementation (2018), <https://github.com/runtimeverification/verified-smart-contracts/blob/uniswap/uniswap/x-y-k.pdf>
19. Zhou, L., Qin, K., Cully, A., Livshits, B., Gervais, A.: On the just-in-time discovery of profit-generating transactions in DeFi protocols. In: IEEE Symp. on Security and Privacy. pp. 919–936. IEEE (2021). <https://doi.org/10.1109/SP40001.2021.00113>
20. Zhou, L., Qin, K., Torres, C.F., Le, D.V., Gervais, A.: High-Frequency Trading on Decentralized On-Chain Exchanges. In: IEEE Symp. on Security and Privacy. pp. 428–445. IEEE (2021). <https://doi.org/10.1109/SP40001.2021.00027>

A Proofs

Lemma 1. $\sum_{\mathbb{A}} G_{\mathbb{A}}(\Gamma, \mathbb{T}) = 0$.

Proof. Follows from Lemma 3 (preservation of net worth) in [9].

Lemma 2. If $\text{type}(\mathbb{T}) \in \{\text{dep}, \text{rdm}\}$, then $G_{usr(\mathbb{T})}(\Gamma, \mathbb{T}) = 0$.

Proof. Follows from Lemma 3 (preservation of net worth) in [9].

Lemma 3. Let $\Gamma \xrightarrow{\lambda} \Gamma'$, and let $\text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1)$, $\text{res}_{\tau_0, \tau_1}(\Gamma') = (r'_0, r'_1)$. Then, $P_{\Gamma}(\tau_0, \tau_1) = P_{\Gamma'}(\tau_0, \tau_1)$ if and only if $r_0/r_1 = r'_0/r'_1$.

Proof. Let the *projected minted token price* of (τ_0, τ_1) at reserve ratio $R > 0$ in state Γ be defined as:

$$P_{\Gamma}^R(\tau_0, \tau_1) = \frac{r'_0}{\text{sply}_{\Gamma}(\tau_0, \tau_1)} \cdot P(\tau_0) + \frac{r'_1}{\text{sply}_{\Gamma}(\tau_0, \tau_1)} \cdot P(\tau_1)$$

where for the *projected reserves* (r'_0, r'_1) , both $r'_0 \cdot r'_1 = r_0 \cdot r_1$ and $R = r'_0/r'_1$ hold. Thus, the projected minted token price can be rewritten entirely in terms of token reserves and supply in Γ and *projected ratio* R :

$$P_{\Gamma}^R(\tau_0, \tau_1) = \frac{\sqrt{r_0 \cdot r_1 \cdot R}}{\text{sply}_{\Gamma}(\tau_0, \tau_1)} \cdot P(\tau_0) + \frac{\sqrt{r_0 \cdot r_1 / R}}{\text{sply}_{\Gamma}(\tau_0, \tau_1)} \cdot P(\tau_1) \quad (9)$$

We note that from (9) and (1) it follows that

$$P_{\Gamma}^R(\tau_0, \tau_1) = P_{\Gamma}(\tau_0, \tau_1) \quad \text{if} \quad \begin{array}{l} \text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1) \\ R = r_0/r_1 \end{array} \quad (10)$$

Alternatively, the *projected minted token price* in a given state Γ can be interpreted as the minted token price in Γ' of execution $\mathbb{M}[\sigma] \mid \Gamma \xrightarrow{\mathbb{T}} \mathbb{M}[_]\mid \Gamma'$ where \mathbb{T} is a miner swap action and the reserve ratio $r'_0/r'_1 = R$ holds in Γ' but not in Γ . By definition then, there exists σ and swap \mathbb{T} for any reachable state Γ and $R > 0$, such that $\mathbb{M}[\sigma] \mid \Gamma \xrightarrow{\mathbb{T}} \mathbb{M}[_]\mid \Gamma'$ and $P_{\Gamma}^R(\tau_0, \tau_1) = P_{\Gamma'}(\tau_0, \tau_1)$ if $R \neq r_0/r_1$.

We prove Lemma 3 by showing that for any R , the *projected minted token price* of a pair remains *constant* for any execution. Thus, if in two states Γ, Γ' along an execution the AMM pair reserve ratios both equal $R = r_0/r_1 = r'_0/r'_1$, prices must also be equal, thereby proving the lemma.

$$P_{\Gamma}^R(\tau_0, \tau_1) = P_{\Gamma'}^R(\tau_0, \tau_1) = P_{\Gamma}(\tau_0, \tau_1) = P_{\Gamma'}(\tau_0, \tau_1) \quad (11)$$

We prove that the *projected minted token price* remains constant for any execution by induction.

Base case: empty For an empty step, the projected minted token price remains constant (trivially).

Induction step: deposit/redeem For a deposit or redeem execution $\Gamma_n \rightarrow^\top \Gamma_{n+1}$ the following must hold for $c > 0$ by definition of [DEP] and [RDM]

$$(c \cdot r_0^n, c \cdot r_1^n) = (r_0^{n+1}, r_1^{n+1}) \quad c \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) = \text{sply}_{\Gamma_{n+1}}(\tau_0, \tau_1)$$

Thus, we can write the *projected minted token price* in Γ_{n+1} in terms of reserves and token supply in Γ_n , such that the equality is apparent.

$$\begin{aligned} P_{\Gamma_{n+1}}^R(\tau_0, \tau_1) &= \frac{\sqrt{c^2 \cdot r_0^n \cdot r_1^n \cdot R}}{c \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_0) + \frac{\sqrt{c^2 \cdot r_0^n \cdot r_1^n / R}}{c \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_1) \\ &= \frac{\sqrt{r_0^n \cdot r_1^n \cdot R}}{\text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_0) + \frac{\sqrt{r_0^n \cdot r_1^n / R}}{\text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_1) = P_{\Gamma_n}^R(\tau_0, \tau_1) \end{aligned}$$

Induction step: swap For a swap execution $\Gamma_n \rightarrow^\top \Gamma_{n+1}$ both the supply of minted tokens and the reserve product is maintained by definition of SWAP

$$r_0^n \cdot r_1^n = r_0^{n+1} \cdot r_1^{n+1} \quad \text{sply}_{\Gamma_n}(\tau_0, \tau_1) = \text{sply}_{\Gamma_{n+1}}(\tau_0, \tau_1)$$

Again, we can express the *projected minted token price* in Γ_{n+1} in terms of reserves and token supply in Γ_n to illustrate the equality.

$$P_{\Gamma_{n+1}}^R(\tau_0, \tau_1) = \frac{\sqrt{r_0^n \cdot r_1^n \cdot R}}{\text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_0) + \frac{\sqrt{r_0^n \cdot r_1^n / R}}{\text{sply}_{\Gamma_n}(\tau_0, \tau_1)} \cdot P(\tau_1) = P_{\Gamma_n}^R(\tau_0, \tau_1)$$

Thus, we have shown that the projected minted token price remains constant for all executions. Therefore, (11) holds, proving the lemma. \square

Lemma 4. *There exists a function P^{min} such that if $\mathbf{M}[\sigma] \mid \Gamma \rightarrow^* \mathbf{M}[\sigma'] \mid \Gamma'$ then: (i) $P_{\Gamma'}^R(\tau_0, \tau_1) \geq P_{\Gamma}^{min}(\tau_0, \tau_1)$; (ii) there exist σ'' and λ consisting at most of a swap by \mathbf{M} such that $\mathbf{M}[\sigma''] \mid \Gamma' \xrightarrow{\lambda} \mathbf{M}[\cdot] \mid \Gamma''$ and $P_{\Gamma''}^R(\tau_0, \tau_1) = P_{\Gamma}^{min}(\tau_0, \tau_1)$.*

Proof. lma:price-minimum The proof reuses the definition of the *projected minted token price* (9) defined in the proof of Lemma 3: there, we showed that the *projected minted token price* for any given reserve ratio $R > 0$ remains constant for all executions. Thus, by definition (9), the projected minted token price in Γ for all $R > 0$ is the minted token price *range* which can be achieved by executing a swap in any reachable state Γ .

To find $P_{\Gamma}^{min}(\tau_0, \tau_1)$, we first determine the R for which $P_{\Gamma}^R(\tau_0, \tau_1)$ is minimized in any reachable state Γ .

$$\frac{\partial}{\partial R} P_{\Gamma}^R(\tau_0, \tau_1) = \frac{\sqrt{r_0^n \cdot r_1^n}}{2 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{R}} \cdot P(\tau_0) - \frac{\sqrt{r_0^n \cdot r_1^n}}{2 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{R} \cdot R} \cdot P(\tau_1)$$

Setting the expression above as equal to zero and then solving for $R = R^{min}$ we obtain

$$R^{min} = \frac{r_0}{r_1} = \frac{P(\tau_1)}{P(\tau_0)}$$

Further, we have determined the projected minted token price *minimum* since the second derivative is positive

$$\begin{aligned}
\frac{\partial^2}{\partial R^2} P_\Gamma^R(\tau_0, \tau_1) &= -\frac{\sqrt{r_0^n \cdot r_1^n}}{4 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{R} \cdot R} + \frac{3 \cdot \sqrt{r_0^n \cdot r_1^n}}{4 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{R} \cdot R^2} \\
&= -\frac{\sqrt{r_0^n \cdot r_1^n}}{4 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{\frac{P(\tau_1)}{P(\tau_0)}} \cdot \frac{P(\tau_1)}{P(\tau_0)^2}} + \frac{3 \cdot \sqrt{r_0^n \cdot r_1^n}}{4 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{\frac{P(\tau_1)}{P(\tau_0)}} \cdot \frac{P(\tau_1)}{P(\tau_0)^2}} \\
&= \frac{2 \cdot \sqrt{r_0^n \cdot r_1^n}}{4 \cdot \text{sply}_{\Gamma_n}(\tau_0, \tau_1) \cdot \sqrt{\frac{P(\tau_1)}{P(\tau_0)}} \cdot \frac{P(\tau_1)}{P(\tau_0)^2}} > 0
\end{aligned}$$

Thus, the function $P_\Gamma^{\min}(\tau_0, \tau_1)$ is given as

$$P_\Gamma^{\min}(\tau_0, \tau_1) = P_\Gamma^{P(\tau_1)/P(\tau_0)}(\tau_0, \tau_1)$$

By definition of the *project minted token price*, a swap exists such that the projected price for reserve ratio R is achieved in the resulting state if the reserve ratio in Γ is not equal to R . Otherwise the reserve ratio must equal R , and thus the empty step achieves the projected price trivially. We have shown that $P_\Gamma^R(\tau_0, \tau_1) \geq P_\Gamma^{\min}(\tau_0, \tau_1)$ for any $R > 0$, thereby proving the lemma. \square

Theorem 1. Let $\Gamma = \parallel_{i \in I} (r_{i,0} : \tau_{i,0}, r_{i,1} : \tau_{i,1}) \mid \Gamma_w$, where Γ_w only contains wallets. For all $j \in I$ and $d \in \{0, 1\}$, let $v_j^d = w_d^d(\tau_{j,0}, \tau_{j,1}, \Gamma)$, and let:

$$\sigma_j = \begin{cases} v_j^d : \tau_{j,d} & \text{if } v_j^d > 0 \\ 0 & \text{if } v_j^0, v_j^1 \leq 0 \end{cases} \quad \lambda_j = \begin{cases} \mathbb{X}^d(\tau_{j,0}, \tau_{j,1}, \Gamma) & \text{if } v_j^d > 0 \\ \varepsilon & \text{if } v_j^0, v_j^1 \leq 0 \end{cases}$$

Then, $(\sigma_1 \cdots \sigma_n, \lambda_1 \cdots \lambda_n)$ is a solution to the game (Γ, \mathbb{X}) for an empty \mathbb{X} .

Proof. Theorem 1 states that the solution to (Γ, \mathbb{X}) can be greedily constructed from canonical swaps for each AMM pair in Γ , thereby minimizing the prices of all minted tokens and net worth of users whilst maximizing the gain for the miner.

We prove the lemma by showing that the *price minimization swap* (5) for a pair (τ_0, τ_1) minimizes the respective minted token price. Since all AMM actions affect single pair reserves only, the miner can minimize the minted token price in any order, thereby proving the lemma.

To prove that the *price minimization swap* minimizes the minted token price of a pair, we show that it updates the pair reserve ratio to $r_0/r_1 = P(\tau_1)/P(\tau_0)$, which, as shown in the proof of Lemma 4, minimizes the price for all executions.

Case: $d = 0$ We assume the canonical swap direction to be $d = 0$. By definition of the canonical swap values at page 8, we have:

$$w_0^0(\tau_0, \tau_1, \Gamma) = \sqrt{\frac{P(\tau_1)}{P(\tau_0)}} \cdot r_0 \cdot r_1 - r_0$$

$$w_1^0(\tau_0, \tau_1, \Gamma) = \frac{r_1 \cdot w_0^0(\tau_0, \tau_1, \Gamma)}{r_0 + w_0^0(\tau_0, \tau_1, \Gamma)} = \frac{r_1 \cdot \sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1} - r_0 \cdot r_1}{\sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1}}$$

Further, the reserve product invariant must hold before and after the price minimization swap in direction $d = 0$. We show that this holds:

$$(r_0 + w_0^0(\tau_0, \tau_1, \Gamma)) \cdot (r_1 - w_1^0(\tau_0, \tau_1, \Gamma)) = \sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1} \cdot \frac{r_0 \cdot r_1}{\sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1}} = r_0 \cdot r_1$$

Finally, we can show that the resulting reserve ratio following the price minimization swap is indeed $P(\tau_1)/P(\tau_0)$, thereby minimizing the minted token price (see proof of Lemma 4).

$$\frac{r_0 + w_0^0(\tau_0, \tau_1, \Gamma)}{r_1 - w_1^0(\tau_0, \tau_1, \Gamma)} = \frac{\sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1}}{\frac{r_0 \cdot r_1}{\sqrt{\frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1}}} = \frac{P(\tau_1)}{P(\tau_0)} \cdot r_0 \cdot r_1}{r_0 \cdot r_1} = \frac{P(\tau_1)}{P(\tau_0)}$$

Case: $d = 1$ Follows similarly and is omitted for brevity. \square

Theorem 2. *With respect to the MEV game in (Γ, \mathcal{X}) :*

1. If \mathcal{X} is empty, the solution is the final layer constructed for (Γ, \emptyset) in §4.1.
2. Otherwise, if $\mathcal{X} = [\mathbb{T}] + \mathcal{X}'$, let (σ, λ) be the inner layer constructed for $(\Gamma, [\mathbb{T}])$, let $\mathbf{M}[\sigma] \mid \Gamma \xrightarrow{\lambda} \mathbf{M}[\cdot] \mid \Gamma'$, and let (σ', λ') be the solution for (Γ', \mathcal{X}') . Then, the solution to (Γ, \mathcal{X}) is $(\sigma + \sigma', \lambda \lambda')$.

Proof. We restate the user gain (6) from the execution of a game solution following Lemma 4.

$$\begin{aligned} & G_{\mathbf{A}}(\mathbf{M}[\sigma] \mid \mathbf{A}[\sigma_{\mathbf{A}}] \mid \Gamma, \lambda) \\ &= \sum_{\tau \in \mathbb{T}_0} \sigma'_{\mathbf{A}}(\tau) \cdot P(\tau) - \sigma_{\mathbf{A}}(\tau) \cdot P(\tau) + \sum_{\tau \in \mathbb{T}_1} \sigma'_{\mathbf{A}}(\tau) \cdot P_{\Gamma}^{\min}(\tau) - \sigma_{\mathbf{A}}(\tau) \cdot P_{\Gamma}(\tau) \end{aligned}$$

Here, the prices are either of atomic ($P_{\Gamma}(\tau)$), or minted tokens ($P_{\Gamma}(\tau)$ and $P_{\Gamma}^{\min}(\tau)$), all determined in the initial state Γ . Thus, the exploitation of individual user actions by the miner is decided on the action's effect the user *token balance* only.

We prove Theorem 2 by showing that the "inner layer" for each user action type are *optimal* when constructed in any order from the submitted user actions in \mathcal{X} .

Swap-inner-layer Firstly, we show that the *swap front-run* by the miner will always minimize the amount of tokens *received* by the user. Let $\mathbb{T} = \mathbf{A} : \text{swap}^0(v_0 : \tau_0, v_1 : \tau_1)$ where $d_{\mathbf{A}} = 0$ and

$$\mathbf{M}[\cdot] \mid \Gamma \xrightarrow{\text{SFX}(\tau_0, \tau_0, \Gamma, \mathbb{T})} \mathbf{M}[\cdot] \mid \Gamma' \xrightarrow{\mathbb{T}} \mathbf{M}[\cdot] \mid \Gamma''$$

If the execution of user swap \mathbb{T} results in the *minimal* received output amount v_1 for \mathbf{A} , then for $\text{res}_{\tau_0, \tau_1}(\Gamma) = (r_0, r_1)$, $\text{res}_{\tau_0, \tau_1}(\Gamma') = (r'_0, r'_1)$ and $\text{res}_{\tau_0, \tau_1}(\Gamma'') = (r'_0 + v_0, r'_1 - v_1)$ the reserve product invariant must hold by definition of $[\text{SWAP}]$.

$$(r'_0 + v_0) \cdot (r'_1 - v_1) = r'_0 \cdot r'_1 = r_0 \cdot r_1$$

Solving for r'_0 , we can rewrite as:

$$\begin{aligned} (r'_0 + v_0) \cdot \left(\frac{r_0 \cdot r_1}{r'_0} - v_1\right) &= r_0 \cdot r_1 \\ r_0 \cdot r_1 - v_0 \cdot r'_0 + \frac{v_0 \cdot r_0 \cdot r_1}{r'_0} - v_0 \cdot v_1 &= r_0 \cdot r_1 \\ v_1 \cdot r_0'^2 + v_0 \cdot v_1 \cdot r'_0 - v_0 \cdot r_0 \cdot r_1 &= 0 \end{aligned}$$

The determinant to the quadratic equation is

$$D = v_0^2 \cdot v_1^2 + 4 \cdot v_0 \cdot v_1 \cdot r_0 \cdot r_1$$

Thus we can solve for *positive* reserves r'_0, r'_1 in state Γ' expressed in terms of the swap parameters (v_0, v_1) and reserves r_0, r_1 in initial state Γ , which coincide with definitions of the *swap front-run reserves* for $d_A = 0$ (the case $d_A = 1$ is omitted for brevity).

$$r'_0 = \frac{-v_0 \cdot v_1 + \sqrt{v_0^2 \cdot v_1^2 + 4 \cdot v_0 \cdot v_1 \cdot r_0 \cdot r_1}}{2 \cdot v_1} \quad r'_1 = \frac{r_0 \cdot r_1}{r'_0}$$

If $(r_0, r_1) = (r'_0, r'_1)$, then clearly no swap front-run is required. Otherwise, the direction of the swap front-run depends on the value of r'_0, r'_1 . For $r'_0 > r_0$ and $r'_1 < r_0$, the swap-front run direction $d_M = 0$ is implied. For $r'_0 > r_0$ and $r'_1 < r_0$, $d_M = 1$. The *swap front-run values* (7) follow from the difference between initial and *swap front-run* reserves.

Since the swap front-run always enables the user swap such that the the minimum output amount is returned, this implies that the *effect* on the user token balance when executing the solution (6) is solely determined by user swap parameters (v_0, v_1) : it is not affected by its position in the full solution, enabling the greedy construction of the swap-inner-layer in Theorem 2.

The optimality of the swap-inner-layer can be easily shown: For our assumed user swap direction $d_A = 0$, if $-v_0 \cdot P(\tau_0) + v_1 \cdot P(\tau_1) < 0$ holds, then the contribution to the user gain (6) must be negative, and furthermore, since by definition of $[\text{SWAP}]$, v_1 is the minimum amount the user can receive, the swap-inner-layer must be optimal.

If $-v_0 \cdot P(\tau_0) + v_1 \cdot P(\tau_1) \geq 0$, then the swap-inner-layer will be $(0, \varepsilon)$, since there the user swap can never reduce the user gain in any game solution. We omit the case $d_A = 1$ for brevity.

Deposit-inner-layer The optimality of the deposit-inner-layer follows from Section 4.2.

Redeem-inner-layer The optimality of the redeem-inner-layer $(0, \varepsilon)$ follows from Section 4.2. \square