

Defensive Rebalancing in Networks of Automated Market Makers

Sam Devorsetz 

Computer Science Department, Brown University, Providence, RI, USA

Maurice Herlihy  

Computer Science Department, Brown University, Providence, RI, USA

Abstract

This paper introduces and analyzes *defensive rebalancing*, a novel mechanism for preventing arbitrageurs from extracting surplus from networks of constant-function market makers (CFMMs).

A *rebalancing* transfers assets directly from one CFMM's pool to another's, bypassing the CFMMs' standard trading protocols. In any *arbitrage-prone* configuration, we prove there exists a rebalancing to an *arbitrage-free* configuration that strictly increases the liquidity of at least one participating CFMM without decreasing the liquidity of any other participating CFMM. Moreover, we prove that a configuration is arbitrage-free if and only if it is *Pareto efficient* under rebalancing.

We prove that for any log-concave trading function, including the ubiquitous constant-product market maker, finding an optimal arbitrage-free rebalancing that maximizes the aggregate liquidity of the participating CFMMs while preserving each participant's liquidity can be formulated as a convex optimization problem with a unique, computationally tractable solution.

We extend this framework to *mixed rebalancing*, where a subset of participating CFMMs use a combination of direct transfers and standard trades to transition to an arbitrage-free configuration while capturing arbitrage opportunities from non-participating CFMMs and external price-setting oracles such as centralized exchanges.

Our results provide a rigorous foundation for future CFMM protocols that coordinate defensive rebalancing to internalize arbitrage surplus within networks of CFMMs.

2012 ACM Subject Classification Information systems → Distributed storage

Keywords and phrases Automated Market Makers, Arbitrage, Convex Optimization

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.12

Funding *Maurice Herlihy*: Supported by a grant from Nunchi corporation.

1 Introduction

A *constant-function automated market maker* [3] (CFMM) $\mathcal{A}$ is an automaton that trades one token for another. It owns two pools of tokens, let's call them dollar tokens (\$) and euro tokens (€). $\mathcal{A}$ holds x_1 dollars in its dollar pool, and x_2 euros in its euro pool. Token exchange rates are determined by a *trading function* $F(x_1, x_2)$, satisfying some common-sense technical conditions given below. Initially $F(x_1, x_2) = k$, for some $k > 0$, called $\mathcal{A}$'s *liquidity*. A trader who transfers $\delta_\$$ dollars to $\mathcal{A}$ receives $\delta_€$ euros in return, where $F(x_1 + \delta_\$, x_2 - \delta_€) = k$. That is, a *trade* changes $\mathcal{A}$'s pool sizes but leaves its liquidity unchanged.

Consider the following *circular trade* scenario. There are three CFMMs, each trading one token type for another. $\mathcal{A}$ trades dollars for euros, $\mathcal{B}$ trades euros for pounds (£), and $\mathcal{C}$ trades pounds for dollars. All three have trading function $F(x_1, x_2) := x_1 \cdot x_2 = 3$.

- $\mathcal{A}$ has €1 and \$3,
- $\mathcal{B}$ has £1 and €3, and
- $\mathcal{C}$ has \$1 and £3.



© Sam Devorsetz and Maurice Herlihy;

licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 12; pp. 12:1–12:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

12:2 Defensive Rebalancing in Networks of AMMs

Alice (an *arbitrageur*) sees an opportunity to make a quick risk-free profit. She sells \$1 to $\mathcal{C}$ for $\pounds \frac{3}{2}$, which she sells to $\mathcal{B}$ for $\text{€} \frac{9}{5}$, which she sells to $\mathcal{A}$ for $\$ \frac{27}{14}$, collecting a profit of $\$ \frac{13}{14}$. Her profit comes at the expense of the *liquidity providers* (LPs), the investors who loaned their tokens to $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$.

Suppose instead that before Alice strikes, the LPs for $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$, recognizing their danger, agree to transfer $\text{€}1$ directly from $\mathcal{B}$'s euro pool to $\mathcal{A}$'s euro pool, $\pounds 1$ from $\mathcal{C}$'s pound pool to $\mathcal{B}$'s pound pool, and $\$1$ from $\mathcal{A}$'s dollar pool to $\mathcal{C}$'s dollar pool. These are direct pool-to-pool transfers, bypassing the CFMMs' usual trading protocols.

The state resulting from this *rebalancing* has three interesting properties. First, arbitrage across $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ has become impossible: no circular trade results in a profit. Second, rebalancing has improved the liquidity levels of *all three* CFMMs (each went from 3 to 4), without reducing the value of any CFMM's assets (in a sense formalized below). Third, no subsequent rebalancing can increase one CFMM's liquidity without reducing another's. The first condition improves social welfare by harmonizing exchange rates; the second improves individual welfare because higher liquidity implies lower price slippage, making that CFMM more attractive to traders. The third condition means individual welfare cannot be further improved.

This paper shows that the first and third conditions are equivalent, and the second follows as a consequence of the others. As explained in Section 2, a *configuration* is the global state of a network of CFMMs. A configuration is *arbitrage-free* if no agent can make a profit simply by executing a sequence of trades. As noted, a *rebalancing* is a set of direct transfers from one CFMM pool directly to another. A configuration c is *Pareto efficient* (under rebalancing) if any rebalancing starting from c that increases one CFMM's liquidity must reduce another's. In Section 3, we show that *the arbitrage-free configurations are precisely the Pareto-efficient configurations*. In short, the Pareto-efficient configurations are exactly those that eliminate arbitrage opportunities while preserving or improving every participating CFMM's liquidity.

This equivalence has several implications.

- From any arbitrage-prone configuration, there exist rebalancings that simultaneously eliminate arbitrage exposure and strictly increase some CFMMs' liquidities without decreasing the others'.
- One could provide “rebalancing as a service” where a set of participating CFMMs periodically check whether their configuration is arbitrage prone. If so, they agree on a rebalancing that simultaneously protects them from arbitrage and renders their liquidity levels Pareto efficient.
- Section 4 shows that for any log-concave trading function, the problem of finding a rebalancing leading to an arbitrage-free/Pareto-efficient configuration can be formulated as a convex optimization problem, computationally tractable using widely available solvers. Moreover, the equivalence between arbitrage-free and Pareto-efficient configurations simplifies the optimization program's formulation.
- Section 5 shows that this convex optimization program can be readily extended to systems where some CFMMs refuse to allow direct pool-to-pool transfers, and some act like powerful centralized exchanges, setting global market prices for certain tokens.

Our analysis assumes that a proposed rebalancing is executed before the corresponding arbitrage opportunity has been exhausted by external traders. We focus on the existence and properties of such rebalancings rather than on the mechanisms required to schedule or prioritize them.

2 Model

A function $F : \mathbb{R}^k \rightarrow \mathbb{R}$ is *strictly convex* if for all $t \in (0, 1)$ and distinct x_1, x_2 ,

$$F(tx_1 + (1-t)x_2) < tF(x_1) + (1-t)F(x_2).$$

A positive-valued function F is *log-concave* if for all $t \in (0, 1)$ and distinct x_1, x_2 ,

$$\log F(tx_1 + (1-t)x_2) > t \log F(x_1) + (1-t) \log F(x_2).$$

We use “=” for equality and “:=” for definitions.

2.1 Automated Market Makers

Agents own tokens. A token is any valuable item that can be transferred electronically. An *automated market maker* (AMM) is an automaton that trades between two token types, T_1, T_2 . An AMM manages a *pool* of each token. An AMM’s *state* is a pair (x_1, x_2) , meaning the AMM has x_1 units of T_1 in its first pool, and x_2 units of T_2 in its second pool.¹ A *configuration* c is an ordered collection of AMM states.

We consider *constant-function market makers* [3] (CFMMs) in which each trade preserves an invariant of the form $F(x_1, x_2) = k$, where $k > 0$, and $F : \mathbb{R}^2 \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing in each argument, and log-concave. The value k is called the CFMM’s *liquidity*.² Perhaps the most popular CFMM is the *constant-product market maker* (CPMM) of the form $F(x_1, x_2) := x_1 \cdot x_2 = k$. Uniswap [1], Balancer [14], Bullish [7], CoinEx [10], and Skate [19] are examples of commercial CPMMs.

A *liquidity provider* (LP) is an agent who invests in a CFMM by loaning tokens to its liquidity pools. LPs expect to profit from trading fees, and typically receive governance rights over the CFMM.

For an AMM holding $x_\$$ dollars and x_ϵ euros, a *trade* is a pair $(\delta_\$, \delta_\epsilon)$ such that

$$F(x_\$ + \delta_\$, x_\epsilon - \delta_\epsilon) = F(x_\$, x_\epsilon).$$

Thus, if a trader sends $\delta_\$$ dollars to the AMM, it receives δ_ϵ euros in return, where δ_ϵ is chosen so that the liquidity remains unchanged. Formally, the AMM’s possible states under trading correspond to a *level set* of F :

$$\{(x_1, x_2) \mid F(x_1, x_2) = k\},$$

where the liquidity k remains constant for the duration of the trade.

For a CFMM with trading function $F(x_1, x_2)$, the *spot price* of token T_1 in terms of T_2 in state (x_1, x_2) is:

$$-\frac{\partial F / \partial x_1}{\partial F / \partial x_2}(x_1, x_2).$$

Trades change spot prices, an effect known as *slippage*. For example, a CPMM $x_1 x_2 = 2$ in state $(1, 2)$ has a dollar spot price of 2 euros, but a trader who deposits 2 euros changes the state to $(\frac{1}{2}, 4)$, and receives only $\frac{1}{2}$ dollar. The trade increased the spot price as dollars become more scarce compared to euros.

¹ Our results generalize to AMMs that trade more than two types of tokens, but we focus on two-pool AMMs for simplicity.

² This term is sometimes called the CFMM’s *invariant* or *liquidity constant*. In our model, however, it is neither invariant nor constant over the course of repeated rebalancings.

12:4 Defensive Rebalancing in Networks of AMMs

Consider two CFMMs $\mathcal{A}, \mathcal{B}$, both trading dollars against euros. If $\mathcal{A}$ and $\mathcal{B}$ have different spot prices, then they are vulnerable to *arbitrage*. Suppose euros are cheaper in $\mathcal{B}$ than in $\mathcal{A}$. An *arbitrageur* agent spends d dollars to buy euros from $\mathcal{B}$, then sells those euros to $\mathcal{A}$ where they are more valuable, receiving $d' > d$ dollars in return. The agent has made a risk-free profit at the expense of the CFMMs' LPs. Arbitrage can of course involve multiple AMMs.

► **Definition 1.** A configuration is *arbitrage-prone* if an agent starting with a basket of tokens $(b_1, \dots, b_m)$ can execute a sequence of trades and end up with a basket $(b'_1, \dots, b'_m)$ where $b'_i \geq b_i$, and at least one inequality is strict. A configuration that is not arbitrage-prone is *arbitrage-free*.

A *rebalancing* transfers tokens directly between liquidity pools belonging to different CFMMs. Unlike trades, rebalancing typically changes both AMMs' liquidities and their spot prices. A rebalancing is self-funding in the sense that it does not change the total amounts of each token held in participants' pools. A rebalancing *improves liquidity* if every AMM's post-rebalancing liquidity is at least its original liquidity, and at least one such inequality is strict.

Note that trades and rebalancings are distinct operations. A trade takes place between a party and a CFMM, and always leaves the AMM's liquidity unchanged. A rebalancing takes place among multiple CFMMs, and typically alters their liquidities. Trades are a standard part of classical AMMs, while rebalancings are (as far as we know) analyzed for the first time in this paper.

Formally, a rebalancing is defined as follows. Starting from a configuration

$$c := (x_{1,1}, x_{1,2}), \dots, (x_{n,1}, x_{n,2}),$$

a rebalancing carries c to a new configuration

$$c' := (x'_{1,1}, x'_{1,2}), \dots, (x'_{n,1}, x'_{n,2}).$$

Tokens can be transferred only between pools holding the same token type. We encode permitted transfers in the index set:

$$\mathcal{T} := \{(i, j, k, \ell) \mid i < k \text{ and the pools corresponding to } x_{i,j}, x_{k,\ell} \text{ have the same token type}\}.$$

A rebalancing is defined by a collection:

$$\{\delta_{i,j,k,\ell} \mid (i, j, k, \ell) \in \mathcal{T}\}.$$

Each $\delta_{i,j,k,\ell}$ represents a transfer of $|\delta_{i,j,k,\ell}|$ tokens. If $\delta_{i,j,k,\ell} > 0$, then $|\delta_{i,j,k,\ell}|$ tokens are transferred from the pool corresponding to $x_{i,j}$ to the pool corresponding to $x_{k,\ell}$; if $\delta_{i,j,k,\ell} < 0$, the transfer is in the opposite direction. More precisely, the rebalancing sends each pool $x_{i,j}$ to $x'_{i,j}$, where:

$$x'_{i,j} = x_{i,j} + \sum_{\substack{k,\ell: \\ (k,\ell,i,j) \in \mathcal{T}}} \delta_{k,\ell,i,j} - \sum_{\substack{k,\ell: \\ (i,j,k,\ell) \in \mathcal{T}}} \delta_{i,j,k,\ell}.$$

A rebalancing is *feasible* if all final pool balances are positive. All rebalancings considered here are assumed to be feasible.

► **Definition 2.** A configuration c is *Pareto efficient* if no rebalancing starting from c can strictly increase the liquidity of any CFMM without strictly reducing the liquidity of another CFMM.

Protecting a network of CFMMs from arbitrage is a public good that benefits the network as a whole, although those benefits may be distributed unevenly among the individual CFMMs. Because a given arbitrage-prone configuration typically admits a continuum of Pareto-efficient rebalancings, there is considerable freedom in selecting a target configuration. For example, in Section 4, we consider a rebalancing that selects the Pareto-efficient configuration maximizing the product of the squares of the CFMMs' resulting liquidities. Each participating CFMM may communicate additional constraints to the rebalancing agent (e.g., "I must retain at least 100 euros"). The agent, in turn, is free to exclude any CFMM whose individual constraints make rebalancing infeasible. Ultimately, however, participation is voluntary: every CFMM retains veto power and may decline to participate in any proposed rebalancing for any reason.

We assume each sequence of trades or rebalancings executes *atomically*, so that no other agent's actions are interleaved with it.

Our results characterize the allocations that are achievable under rebalancing. They do not address how a rebalancing transaction obtains priority over competing arbitrage transactions on a blockchain.

As discussed in Section 6, rebalancing is more powerful than trading alone for protection against arbitrage. Standard CFMM trades do not improve liquidity. Moreover, no self-funding sequence of trades can, in general, harmonize prices: the agent executing those trades generally cannot eliminate arbitrage while simultaneously breaking even.

3 Arbitrage and Pareto Efficiency

► **Theorem 3.** *Arbitrage-prone configurations are not Pareto efficient under rebalancing.*

Proof. Let

$$c := (x_{1,1}, x_{1,2}), \dots, (x_{n,1}, x_{n,2})$$

be an arbitrage-prone configuration. We construct a rebalancing that leads to a configuration

$$c' := (x'_{1,1}, x'_{1,2}), \dots, (x'_{n,1}, x'_{n,2})$$

where $F_i(x_{i,1}, x_{i,2}) \leq F_i(x'_{i,1}, x'_{i,2})$ for $i \in \{1, \dots, n\}$, and at least one inequality is strict, contradicting the hypothesis that c is Pareto efficient.

Because c is arbitrage prone by hypothesis, an arbitrageur starting with a basket of tokens $(b_1, \dots, b_m)$ can execute a sequence of trades, ending with a basket $(b'_1, \dots, b'_m)$, where each $b'_j \geq b_j$, for $j \in \{1, \dots, m\}$, and at least one inequality is strict.

To construct a rebalancing, borrow a basket $(b_1, \dots, b_m)$, transfer tokens among the pools in a way that mimics that arbitrageur's sequence of trades, producing $(b'_1, \dots, b'_m)$, where each $b'_j \geq b_j$, and at least one inequality is strict. Repay the loan, leaving the basket $(b'_1 - b_1, \dots, b'_m - b_m)$. Recall that each of these trades leaves each CFMM's liquidity unchanged.

For each token type T_k for which $b'_k - b_k > 0$, choose any CFMM that trades token type T_k , and transfer $b'_k - b_k$ tokens of type T_k to that pool. This transfer strictly increases that CFMM's liquidity. Repeating this procedure for every token type with $b'_k > b_k$ produces a rebalancing in which every CFMM's liquidity is at least its original value and at least one CFMM's liquidity is strictly greater, contradicting the hypothesis that c is Pareto efficient. ◀

► **Definition 4.** *A valuation for a network of CFMMs trading token types $T_1, \dots, T_m$ is a vector $V = (V_1, \dots, V_m) \in \mathbb{R}_{>0}^m$, where V_i is the value assigned to one unit of token T_i .*

12:6 Defensive Rebalancing in Networks of AMMs

► **Definition 5.** A valuation $V = (V_1, \dots, V_m)$ is a global valuation for a configuration if, for every CFMM that trades token types T_p and T_q ,

$$\frac{V_p}{V_q} = P_{pq},$$

where P_{pq} is the spot price of T_p measured in units of T_q .

The proof of the following lemma appears in the appendix.

► **Lemma 6.** Any arbitrage-free configuration has a global valuation.

The remainder of this section uses the global valuation to assign a linear value to each CFMM. We show that increasing a CFMM's liquidity necessarily increases its value, while the total value of all CFMMs is preserved under rebalancing.

In an arbitrage-free configuration with global valuation V , define the *value* of the CFMM state $(x_{i,1}, x_{i,2})$, which trades token types T_j and T_k , by

$$\Phi_V(x_{i,1}, x_{i,2}) := V_j x_{i,1} + V_k x_{i,2}.$$

The tangent line to the level set

$$\{(X_1, X_2) \mid F_i(X_1, X_2) = F_i(x_{i,1}, x_{i,2})\}$$

at point $(x_{i,1}, x_{i,2})$ is:

$$L_i := \{(X_1, X_2) \mid \Phi_V(X_1, X_2) = \Phi_V(x_{i,1}, x_{i,2})\}. \quad (1)$$

The tangent L_i has slope $-\frac{V_j}{V_k}$, and divides $\mathbb{R}_{>0}^2$ into two half-planes. Points in the half-plane below L_i have lower value than $\Phi_V(x_{i,1}, x_{i,2})$ and points above have higher values.

► **Lemma 7.** $\Phi_V(x'_{i,1}, x'_{i,2}) < \Phi_V(x_{i,1}, x_{i,2})$ if and only if $(x'_{i,1}, x'_{i,2})$ lies in the half-plane below L_i .

Proof. From Equation (1). ◀

Every point with higher liquidity lies above L_i .

► **Lemma 8.** The region $\{(X_1, X_2) \mid F_i(X_1, X_2) \geq F_i(x_{i,1}, x_{i,2})\}$ lies entirely in the half-plane above L_i .

Proof. Since the upper level set of a log-concave function is convex, the tangent line L_i is a supporting line. ◀

► **Corollary 9.** Any rebalancing that increases the liquidity $F_i(x_{i,1}, x_{i,2})$ must also increase the value $\Phi_V(x_{i,1}, x_{i,2})$.

► **Corollary 10.** Any rebalancing that decreases the value $\Phi_V(x_{i,1}, x_{i,2})$ must also decrease the liquidity $F_i(x_{i,1}, x_{i,2})$.

► **Theorem 11.** Any arbitrage-free configuration is Pareto efficient under rebalancing.

Proof. We will show that from an arbitrage-free configuration, any rebalancing that increases one CFMM's liquidity must decrease another's.

Consider a rebalancing that carries each $(x_{i,1}, x_{i,2})$ to $(x'_{i,1}, x'_{i,2})$, where $F_i(x'_{i,1}, x'_{i,2}) > F_i(x_{i,1}, x_{i,2})$. By Corollary 9, $\mathcal{A}_i$'s value increases: $\Phi_V(x'_{i,1}, x'_{i,2}) > \Phi_V(x_{i,1}, x_{i,2})$.

Because rebalancing neither creates nor destroys tokens, the sum of the CFMMs' values is invariant:

$$\sum_{i=1}^n \Phi_V(x_{i,1}, x_{i,2}) = \sum_{i=1}^n \Phi_V(x'_{i,1}, x'_{i,2}).$$

Therefore there exists some j such that

$$\Phi_V(x'_{j,1}, x'_{j,2}) < \Phi_V(x_{j,1}, x_{j,2}).$$

By Corollary 10, $F_j(x'_{j,1}, x'_{j,2}) < F_j(x_{j,1}, x_{j,2})$, implying that the rebalancing strictly lowered at least one CFMM's liquidity. $\blacktriangleleft$

Theorems 3 and 11 immediately imply the main result:

► **Theorem 12.** *A configuration is arbitrage free if and only if it is Pareto efficient.*

4 Rebalancing as Convex Optimization

Parameters:	$n \in \mathbb{N}$ (Number of CFMMs) $x_{i,j} \quad i \in \{1, \dots, n\}, j \in \{1, 2\}$ (Initial pool sizes) $\mathcal{T} = \{(i, j, k, \ell) \mid \text{transfer allowed between pools } x_{i,j}, x_{k,\ell}\}$
Decision variables:	$x'_{i,j} \quad i \in \{1, \dots, n\}, j \in \{1, 2\}$ (Final pool sizes) $\delta_{i,j,k,\ell}$ Amount (+/-) transferred along $(i, j, k, \ell) \in \mathcal{T}$
Objective:	$\max \quad L(x'_{1,1}, x'_{1,2}, \dots, x'_{n,1}, x'_{n,2}) := \sum_{i=1}^n \log F_i(x'_{i,1}, x'_{i,2})$
Subject to:	$x'_{i,j} \geq 0 \quad \text{for } i \in \{1, \dots, n\}, j \in \{1, 2\} \quad (\text{non-negativity})$ $\log F_i(x_{i,1}, x_{i,2}) \leq \log F_i(x'_{i,1}, x'_{i,2}), \quad i = 1, \dots, n \quad (\text{liquidity non-reduction})$ $x'_{i,j} = x_{i,j} + \sum_{(k,\ell,i,j) \in \mathcal{T}} \delta_{k,\ell,i,j} - \sum_{(i,j,k,\ell) \in \mathcal{T}} \delta_{i,j,k,\ell} \quad (\text{conservation})$

■ **Figure 1** Liquidity Rebalancing Optimization.

In this section, we show how to find a unique rebalancing to protect a configuration from arbitrage. We require the solution to be unique because mutually mistrustful LPs must independently verify the result. Any ambiguity would lead to missed rebalancing opportunities.

Specifically, we show that finding an arbitrage-free rebalancing can be cast as a convex optimization problem. Moreover, we can exploit Theorem 12 to simplify the optimization program's constraints by explicitly seeking a Pareto-efficient final configuration instead of a (possibly more complicated) arbitrage-free final configuration.

4.1 Parameters

Rebalancings can transfer tokens only between pools of the same type. We capture which transfers are permissible through the index set:

$$\mathcal{T} := \{(i, j, k, \ell) \mid \text{transfers allowed between pools } x_{i,j}, x_{k,\ell}\}.$$

The other parameters are the starting pool sizes:

$$(x_{1,1}, x_{1,2}), \dots, (x_{n,1}, x_{n,2}).$$

4.2 Decision Variables

There are transfer amounts

$$\delta_{i,j,k,\ell}, \quad (i,j,k,\ell) \in \mathcal{T}.$$

If $\delta_{i,j,k,\ell} > 0$, then $|\delta_{i,j,k,\ell}|$ tokens are transferred from the pool corresponding to $x_{i,j}$ to the pool corresponding to $x_{k,\ell}$. If $\delta_{i,j,k,\ell} < 0$, the transfer is in the opposite direction. If $\delta_{i,j,k,\ell} = 0$, no transfer occurs.

The remaining decision variables are the new pool sizes:

$$(x'_{1,1}, x'_{1,2}), \dots, (x'_{n,1}, x'_{n,2}).$$

4.3 Constraints

The *non-negativity* constraints require new pool sizes to be non-negative:

$$x'_{i,j} \geq 0 \quad \text{for } i \in \{1, \dots, n\}, j \in \{1, 2\}.$$

This system of linear inequalities defines a convex set.

The *conservation* constraints define the relationship between the pool sizes and transfers. For $i \in \{1, \dots, n\}, j \in \{1, 2\}$,

$$x'_{i,j} = x_{i,j} + \sum_{\substack{k,\ell: \\ (k,\ell,i,j) \in \mathcal{T}}} \delta_{k,\ell,i,j} - \sum_{\substack{k,\ell: \\ (i,j,k,\ell) \in \mathcal{T}}} \delta_{i,j,k,\ell}.$$

This system of linear equalities defines a convex set.

The *liquidity non-reduction* constraints require that rebalancing not decrease any liquidity levels:

$$F_i(x_{i,1}, x_{i,2}) \leq F_i(x'_{i,1}, x'_{i,2}), \quad i \in \{1, \dots, n\}.$$

This constraint, however, is generally not convex, but it can be made convex through a log transform:

$$\log F_i(x_{i,1}, x_{i,2}) \leq \log F_i(x'_{i,1}, x'_{i,2}), \quad i \in \{1, \dots, n\}.$$

Since the left-hand side is constant and the right-hand side is concave, the feasible region defined by this inequality is convex.

4.4 Objective Function

The objective function can express goals that span individual CFMMs, thereby promoting social welfare.

► **Theorem 13.** *Let $L : \mathbb{R}^n \rightarrow \mathbb{R}$ be strictly increasing in each argument. Any configuration that maximizes $L(F_1(x'_{1,1}, x'_{1,2}), \dots, F_n(x'_{n,1}, x'_{n,2}))$ over feasible rebalancings is arbitrage-free.*

Proof. Suppose c maximizes $L(F_1(x_{1,1}, x_{1,2}), \dots, F_n(x_{n,1}, x_{n,2}))$ but is not arbitrage-free. By Theorem 12, c is not Pareto efficient, so there is a rebalancing that takes each $(x_{i,1}, x_{i,2})$ in c to $(x'_{i,1}, x'_{i,2})$, where each $F_i(x_{i,1}, x_{i,2}) \leq F_i(x'_{i,1}, x'_{i,2})$, and at least one inequality is strict. It follows that $L(F_1(x_{1,1}, x_{1,2}), \dots, F_n(x_{n,1}, x_{n,2})) < L(F_1(x'_{1,1}, x'_{1,2}), \dots, F_n(x'_{n,1}, x'_{n,2}))$, contradicting the hypothesis. ◀

► **Corollary 14.** *Any configuration that maximizes a weighted sum or weighted product of liquidities*

$$\max \sum_{i=1}^n w_i F_i(x'_{i,1}, x'_{i,2}) \quad \text{or} \quad \max \prod_{i=1}^n F_i(x'_{i,1}, x'_{i,2})^{w_i}$$

for $w_i > 0$ is arbitrage-free.

For the objective function, we choose to maximize the product of the CFMMs' liquidities. Maximizing this objective is not directly a convex optimization problem, but it can be made convex through a log transform:

$$\max \quad L(x'_{1,1}, x'_{1,2}, \dots, x'_{n,1}, x'_{n,2}) := \sum_{i=1}^n \log F_i(x'_{i,1}, x'_{i,2}).$$

Because each $\log F_i(\cdot, \cdot)$ is concave on its domain by hypothesis, $L(\cdot, \cdot)$ is concave as the sum of concave functions. Figure 1 shows the complete convex optimization problem.

► **Theorem 15.** *The optimization program in Figure 1 is a convex optimization problem.*

Proof. The claim follows because we have shown that (1) the objective function to be maximized is concave, and (2) the feasible region defined by the constraints is convex. ◀

► **Theorem 16.** *The optimization program in Figure 1 has a unique solution.*

Proof. Because each $\log F_i$ is strictly concave, the objective function is strictly concave. A strictly concave function has at most one maximizer over a convex feasible set. ◀

As noted earlier, if CFMMs have additional convex constraints, they can be incorporated into the optimization program, though care must be taken that these additional constraints do not render the problem infeasible.

5 Mixed Populations

So far we have considered rebalancing among a population of CFMMs, all of whom are willing to rebalance with one another to mitigate a common arbitrage vulnerability.

Some LPs might be unwilling to allow a rebalancing agent to transfer tokens directly in and out of their CFMMs' pools. We distinguish between *active* CFMMs who are willing to allow direct pool-to-pool transfers, and *passive* CFMMs, who are not. We assume there is at least one active CFMM. Passive CFMMs cannot be ignored, because their spot prices might expose the others to arbitrage losses, but passive CFMMs can interact with a rebalancing agent only through standard trades, not direct pool-to-pool transfers.

A *price oracle* is a market maker that models an external market that trades at a fixed reference price. For example, consider a centralized exchange with (effectively) inexhaustible liquidity and no slippage that is willing to sell arbitrary amounts of a token at a constant price. This section shows how rebalancing can be extended to accommodate both passive CFMMs and price oracle market makers.

5.1 Active vs. Passive CFMMs

Recall the circular trade scenario where $\mathcal{A}, \mathcal{B}, \mathcal{C}$ have the constant-product trading function $F(x_{i,1}, x_{i,2}) := x_{i,1} \cdot x_{i,2} = 3$, in a configuration where $\mathcal{A}$ has €1 and \$3, $\mathcal{B}$ has £1 and €3, and $\mathcal{C}$ has \$1 and £3. Because this configuration is arbitrage-prone, the LPs for $\mathcal{A}$ decide to rebalance, but this time without cooperation from the LPs for $\mathcal{B}, \mathcal{C}$. The rebalancing agent does the following:

- Transfer \$0.5 out of $\mathcal{A}$,
- Trade \$0.5 to $\mathcal{C}$ for £1.0,
- Trade £1.0 to $\mathcal{B}$ for €1.5, and
- Transfer €1.5 into $\mathcal{A}$.

Note that each trade is compatible with that CFMM's trading function. The resulting configuration is arbitrage-free. Because $\mathcal{B}, \mathcal{C}$ performed only trades, their liquidity levels remain unchanged, but the liquidity for $\mathcal{A}$ has increased from 3 to 6.25. By actively rebalancing, $\mathcal{A}$ has effectively executed the threatened arbitrage and allocated the profit to itself. In general, that profit would be spread across the active CFMMs.

A rebalancing is *restricted* if it does not change passive CFMMs' liquidities. A restricted rebalancing can be implemented by executing both direct transfers between active CFMMs and trades with any CFMMs, while ensuring the process is self-funding by making sure that every token ends up in some CFMM's pool.

As before, consider CFMMs $\mathcal{A}_1, \dots, \mathcal{A}_n$, with trading functions $F_1(x, y), \dots, F_n(x, y)$. The *active* CFMMs $\mathcal{A}_1, \dots, \mathcal{A}_p$ support arbitrary transfers in and out of their pools, for $p > 0$, while the *passive* CFMMs $\mathcal{A}_{p+1}, \dots, \mathcal{A}_n$ support only standard trades, that is, transfers that leave their liquidities unchanged.

The network of CFMMs forms a *CFMM graph*, where the CFMMs are the vertices, and there is an edge linking each pair of CFMMs $\mathcal{A}, \mathcal{B}$ that trade the same token type. Without loss of generality, we assume this graph is *connected*, meaning that between any pair of vertices $\mathcal{A}, \mathcal{B}$, there is a sequence $\mathcal{A}_1, \dots, \mathcal{A}_k$ called a *path*, where consecutive CFMMs are connected by an edge, $\mathcal{A} = \mathcal{A}_1$, and $\mathcal{B} = \mathcal{A}_k$.

► **Lemma 17.** *Let T be a token type traded by some CFMM $\mathcal{A}_1$. An agent in possession of $\epsilon > 0$ tokens of type T can increase the liquidity of some active CFMM while leaving all other liquidity levels unchanged.*

Proof. By induction on ℓ , the length of the shortest path $\mathcal{A}_1, \dots, \mathcal{A}_\ell$ linking $\mathcal{A}_1$ to an active CFMM $\mathcal{A}_\ell$ in the CFMM graph.

For the base case, $\ell = 1$, so $\mathcal{A}_1$ is active. The agent simply transfers the ϵ tokens into $\mathcal{A}_1$'s pool for T , increasing $\mathcal{A}_1$'s liquidity while leaving the others unchanged.

For the induction step, assume the result for paths shorter than $\ell > 1$. The agent first trades at $\mathcal{A}_1$ the ϵ tokens of type T for $\epsilon' > 0$ tokens of type T' , leaving $\mathcal{A}_1$'s liquidity unchanged. The agent now holds $\epsilon' > 0$ tokens of type T' traded by $\mathcal{A}_2$, a passive CFMM linked to an active CFMM by a path shorter than ℓ . The claim follows from the induction hypothesis. ◀

► **Theorem 18.** *Arbitrage-prone configurations are not Pareto efficient under restricted rebalancing.*

Proof. Let c be an arbitrage-prone configuration.

As in the proof of Theorem 3, an arbitrageur starting with a basket of tokens $(b_1, \dots, b_m)$ can execute a sequence of trades starting from c , ending with the basket $(b'_1, \dots, b'_m)$, where each $b'_j \geq b_j$, for $j \in \{1, \dots, m\}$, and at least one inequality is strict.

For each token type T_k for which $b'_k - b_k > 0$, Lemma 17 implies that the agent can transfer those surplus tokens to some active CFMM trading T_k , increasing that CFMM's liquidity while leaving all other liquidity levels unchanged. Repeating this procedure for every profitable token type yields a rebalancing in which every CFMM's liquidity is at least its original value and at least one CFMM's liquidity is strictly greater, contradicting the hypothesis that c is Pareto efficient. ◀

► **Theorem 19.** *Any arbitrage-free configuration is Pareto efficient under restricted rebalancing.*

Proof. The claim follows immediately from Theorem 11. ◀

The *restricted rebalancing convex optimization program* requires only one change to the optimization program of Figure 1: instead of maximizing the product of *all* CFMM liquidities, the restricted program maximizes the product of the active CFMM liquidities *only* (replacing n with p):

$$\max \quad L(x'_{1,1}, x'_{1,2}, \dots, x'_{p,1}, x'_{p,2}) := \sum_{i=1}^p \log F_i(x'_{i,1}, x'_{i,2}). \quad (2)$$

The passive CFMMs' liquidities are bound only by the liquidity non-reduction constraint:

$$F_i(x_{i,1}, x_{i,2}) \leq F_i(x'_{i,1}, x'_{i,2}), \quad i \in \{p+1, \dots, n\}.$$

Let $c' = (x'_{1,1}, x'_{1,2}), \dots, (x'_{n,1}, x'_{n,2})$ be the optimal configuration produced by solving the restricted rebalancing program. Note that because the program's objective function is strictly concave, and the feasible region is convex, the solution c' is unique.

We claim that in c' , the passive CFMMs' non-reduction constraints hold with equality.

► **Lemma 20.**

$$F_i(x_{i,1}, x_{i,2}) = F_i(x'_{i,1}, x'_{i,2}), \quad i \in \{p+1, \dots, n\}.$$

Proof. Assume by way of contradiction that in c' , $F_i(x'_{i,1}, x'_{i,2}) > F_i(x_{i,1}, x_{i,2})$ for some passive $\mathcal{A}_i$. Because $F_i(\cdot, \cdot)$ is strictly increasing in both arguments, we can transfer $\epsilon > 0$ tokens out of pool $x'_{i,1}$ without violating the program's constraints: $F_i(x'_{i,1} - \epsilon, x'_{i,2}) \geq F_i(x_{i,1}, x_{i,2})$. By Lemma 17, we can use these tokens to increase the liquidity of some active CFMM while leaving all other liquidity levels unchanged. This action increases the objective function's value, contradicting the hypothesis that c' is optimal. ◀

► **Lemma 21.** *Configuration c' is arbitrage-free.*

Proof. By Theorem 12, it suffices to show that c' is Pareto efficient. Assume by way of contradiction that there is a rebalancing taking c' to a configuration $c^\dagger = (x^\dagger_{1,1}, x^\dagger_{1,2}), \dots, (x^\dagger_{n,1}, x^\dagger_{n,2})$ where for $i \in \{1, \dots, n\}$, $F_i(x^\dagger_{i,1}, x^\dagger_{i,2}) \geq F_i(x'_{i,1}, x'_{i,2})$, and at least one inequality is strict.

If $c^\dagger$ strictly increases an active CFMM's liquidity, then $c^\dagger$ has a higher objective function value than c' , while continuing to satisfy the restricted program's other constraints, contradicting the hypothesis that c' is Pareto efficient.

If, on the other hand, no active CFMM's liquidity increases, then every strict increase must occur at a passive CFMM. By Lemma 20, however, every passive CFMM already satisfies its liquidity non-reduction constraint with equality in c' . Any strict increase in a passive CFMM's liquidity would continue to satisfy the restricted program's constraints, and by Lemma 17 the additional liquidity could be transferred to an active CFMM without decreasing any other liquidity levels, contradicting the optimality of c' . ◀

12:12 Defensive Rebalancing in Networks of AMMs

In summary, the optimization program of Figure 1 extends naturally to encompass passive CFMMs simply by omitting their liquidities from the objective function.

5.2 Price Oracles

A *price oracle* models a centralized exchange that sets the “market price” for a token by trading at that price. Recall the circular trade scenario with CFMMs $\mathcal{A}, \mathcal{B}, \mathcal{C}$. As before, $\mathcal{A}$ and $\mathcal{B}$ have constant-product trading functions $F_i(x_{i,1}, x_{i,2}) := x_{i,1} \cdot x_{i,2} = 3$. $\mathcal{C}$, however, is an oracle market maker with unlimited liquidity that trades pounds and dollars at par. Initially, $\mathcal{A}$ has €1 and \$3, $\mathcal{B}$ has £1 and €3, and $\mathcal{C}$ has inexhaustible amounts of dollars and pounds. This configuration is arbitrage-prone, so the LPs for $\mathcal{A}$ alone decide to rebalance. The rebalancing agent:

- Transfers $\$(\sqrt{3} - 1) \approx \0.732051 from $\mathcal{A}$,
- Trades $\$(\sqrt{3} - 1)$ for $\pounds(\sqrt{3} - 1)$ at $\mathcal{C}$,
- Trades $\pounds(\sqrt{3} - 1)$ for $\pounds(3 - \sqrt{3}) \approx \pounds1.26795$ at $\mathcal{B}$, and
- Transfers $\pounds(3 - \sqrt{3})$ to $\mathcal{A}$.

In the end, $\mathcal{A}$ has pools $(4 - \sqrt{3}, 4 - \sqrt{3})$ and $\mathcal{B}$ has $(\sqrt{3}, \sqrt{3})$. $\mathcal{A}$ ’s liquidity has increased from 3 to $19 - 8\sqrt{3} \approx 5.14359$, while the others, being passive, remain unchanged. This final configuration is arbitrage-free, and hence Pareto-efficient.

To create a convex optimization program that encompasses one or more oracle AMMs, the key insight is that a price-setting oracle AMM is simply a passive CFMM with linear trading function $F_i(x_{i,1}, x_{i,2}) := ax_{i,1} + bx_{i,2}$, where $a, b > 0$, initialized with sufficient reserves to ensure no pool is exhausted by the rebalancing. (Alternatively, one could permit oracle AMMs to have negative reserve balances in the optimization program.)

6 Remarks

6.1 Alternatives to CFMMs

Although we have focused on CFMMs, the equivalence between arbitrage-free and Pareto-efficient configurations extends to any financial system prone to arbitrage where parties have distinct private utility functions.

For example, suppose Alice owns a basket containing one gold token and one silver token. She subjectively values gold at 1 dollar per token and silver at 2 dollars per token, and thus values her basket at 3 dollars. Bob also owns a basket containing one gold token and one silver token. He subjectively values gold at 2 dollars per token and silver at 1 dollar per token, and thus also values his basket at 3 dollars. Both parties will accept any trade that increases the perceived value of their baskets according to their own valuations.

Suppose an arbitrageur offers Alice $((1 + \epsilon)/2)$ silver tokens in exchange for her gold token. She will accept, because she values the silver she receives at $(1 + \epsilon)$ dollars, while she values the gold she gives up at only 1 dollar. Her new basket is worth $(3 + \epsilon)$ dollars according to her valuation. Similarly, the arbitrageur can offer Bob $((1 + \epsilon)/2)$ gold tokens in exchange for his silver token. Bob will also accept, and he too will value his new basket at $(3 + \epsilon)$ dollars.

Collectively, the arbitrageur receives one gold token and one silver token while paying out $((1 + \epsilon)/2)$ of each. The arbitrageur therefore earns a profit of $((1 - \epsilon)/2)$ gold tokens and $((1 - \epsilon)/2)$ silver tokens.

By comparison, a rebalancing might exchange Alice’s gold token for Bob’s silver token, leaving both parties in an arbitrage-free, Pareto-efficient configuration in which each values their basket at 4 dollars.

6.2 Alternatives to Rebalancing

As an alternative to rebalancing, LPs might instead execute ordinary trades among the CFMMs to mitigate their exposure to arbitrage. Here we use a simple example to illustrate why LPs might prefer defensive rebalancing (as introduced here) to trading. In general, however, defensive trading cannot be *self-funding*: the trading agent must either inject tokens into the system of CFMMs or withdraw them.

Consider the following example. Recall the circular trade scenario where $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are in the arbitrage-prone configuration $\mathcal{A} : (\text{€}1, \$3)$, $\mathcal{B} : (\text{£}1, \text{€}3)$, and $\mathcal{C} : (\$1, \text{£}3)$. Figure 2 shows an optimization program that seeks an arbitrage-free configuration. It uses slack variables $\sigma_{\$}, \sigma_{\text{€}}, \sigma_{\text{£}}$ to represent the amount of each token the agent adds or removes from $\mathcal{A}, \mathcal{B}, \mathcal{C}$. The objective function is to minimize the sum of the squares of the slack variables. If a self-funding solution were possible, these variables would all be zero.

This optimization problem is not convex, but it is small enough to be readily solved by off-the-shelf solvers³. In this case, minimizing the sum of the squares of the slack variables requires transferring $\sigma_{\$} = \sigma_{\text{€}} = \sigma_{\text{£}} = 4 - 2\sqrt{3} \approx 0.535898$ units of each currency from the three CFMMs to the agent, yielding a solution that is both less effective than rebalancing (because the CFMMs permanently surrender assets) and more complicated (because the agent must retain custody of the withdrawn assets).

Parameters:	$(x_{1,1}, x_{1,2}), \quad (x_{2,1}, x_{2,2}), \quad (x_{3,1}, x_{3,2})$
Decision variables:	$(x'_{1,1}, x'_{1,2}), \quad (x'_{2,1}, x'_{2,2}), \quad (x'_{3,1}, x'_{3,2})$ $\sigma_{\text{€}}, \quad \sigma_{\text{£}}, \quad \sigma_{\$}$
Objective:	$\min \quad \sigma_{\text{€}}^2 + \sigma_{\text{£}}^2 + \sigma_{\2
Subject to:	$x'_{1,1}x'_{1,2} = x_{1,1}x_{1,2} = 3 \quad (\text{Invariant for } \mathcal{A})$ $x'_{2,1}x'_{2,2} = x_{2,1}x_{2,2} = 3 \quad (\text{Invariant for } \mathcal{B})$ $x'_{3,1}x'_{3,2} = x_{3,1}x_{3,2} = 3 \quad (\text{Invariant for } \mathcal{C})$ $x_{1,1} + x_{2,2} = x'_{1,1} + x'_{2,2} + \sigma_{\text{€}}$ $x_{2,1} + x_{3,2} = x'_{2,1} + x'_{3,2} + \sigma_{\text{£}}$ $x_{3,1} + x_{1,2} = x'_{3,1} + x'_{1,2} + \sigma_{\$}$ $x'_{1,1}x'_{2,1}x'_{3,1} = x'_{1,2}x'_{2,2}x'_{3,2} \quad (\text{Arbitrage-free condition})$ $x'_{i,j} > 0 \quad \text{for } i \in \{1, 2, 3\}, j \in \{1, 2\}$

■ **Figure 2** Arbitrage protection with trades alone.

7 Related Work

The notion of a CFMM is due to Angeris *et al.* [3], who define the notion of a CFMM, analyze how CFMMs can serve as decentralized price oracles, and formulate the *optimal arbitrage problem*, namely, how to exploit an arbitrage-prone configuration most profitably. Angeris *et al.* [4] consider how to use convex optimization to detect arbitrage-prone configurations, and how best to exploit them, a problem they call *optimal routing*. Our use of optimization complements theirs: rather than exploiting arbitrage opportunities, we compute optimal rebalancings that eliminate them. Angeris *et al.* [5] investigate what happens when a CFMM displaces an external market as a price setter.

³ For this example, we used Mathematica's `NMinimize`.

Milionis *et al.* [17] propose *loss-versus-rebalancing* (LVR) as an improvement over divergence loss as a way of measuring LP costs in configurations with a central reference market. LVR measures the difference between the value of a hypothetical investment strategy that responds optimally to market price changes and the value of an LP's actual investment in a CFMM, which responds only through arbitrage. Defensive rebalancing promises to reduce LVR by allowing LPs to adjust their pool sizes dynamically in response to price changes at other markets.

Danos *et al.* [11] formulate various arbitrage problems in a network of AMMs as optimization programs. Angeris *et al.* [2] describe how a variety of problems associated with CFMMs can be formulated as optimization problems. Wang *et al.* [20] perform an empirical investigation of arbitrage transactions in Uniswap V2, reporting on the frequency of arbitrage opportunities, how often they are exploited, and their profitability. Bartoletti *et al.* [6] describe arbitrage between AMM pairs as a two-party game. Dimitri [12] analyzes how arbitrage can affect CFMM parameters such as pool sizes, prices, and liquidity. Milionis *et al.* [16] study how block time affects arbitrage profits.

As a way to avoid arbitrage altogether, Krishnamachari *et al.* [13] describe *dynamic curves*, which allow AMMs to adjust their trading curves to match oracle-reported prices.

In classical economics, the *First Welfare Theorem* states that, under suitable assumptions, every competitive equilibrium is Pareto efficient, while the *Second Welfare Theorem* states that, under additional convexity assumptions, every Pareto-efficient allocation can be supported as a competitive equilibrium after an appropriate redistribution of initial endowments [15, Chapter 16]. Theorems 3 and 11 bear a strong resemblance to the classical welfare theorems and might at first seem to be special cases. The analogy seems close: CFMMs play the role of consumers, tokens the role of commodities, liquidity the role of utility, rebalancing the role of feasible reallocation, and a common token valuation the role of a competitive price vector. Nevertheless, a direct reduction to the classical welfare theorems is not immediate. In particular, CFMMs do not explicitly optimize utility subject to exogenously specified budget constraints, and the feasible reallocations are constrained by the network of token pools. Thus, establishing Theorems 3 and 11 as formal consequences of those theorems would require a careful correspondence between the two models.

8 Conclusion

8.1 Transaction Costs and Fees

In practice, CFMMs charge *trading fees*. A CFMM $\mathcal{A}_i$ with trading function $F_i(x_1, x_2)$ retains a fraction $\gamma_i \in (0, 1]$ of each incoming trade. If a trader transfers $\delta_{\$}$ dollars to $\mathcal{A}_i$ and receives δ_{ϵ} euros in return, the pool invariants update according to $F_i(x_{i,1} + \gamma_i \delta_{\$}, x_{i,2} - \delta_{\epsilon}) = k$. (The un-discounted fraction $(1 - \gamma_i)\delta_{\$}$ of dollars is typically retained by LPs as fee revenue.) For example, in Uniswap, γ_i depends on the pool's fee tier.

Adding fees to the model requires a few changes to Figure 1. There is a new parameter: each $\mathcal{A}_i$ has a fee parameter γ_i . The conservation constraint is split into two parts:

$$\text{for } i \leq p: \quad x'_{i,j} = x_{i,j} + \sum_{\substack{k,\ell: \\ (k,\ell,i,j) \in \mathcal{T}}} \delta_{k,\ell,i,j} - \sum_{\substack{k,\ell: \\ (i,j,k,\ell) \in \mathcal{T}}} \delta_{i,j,k,\ell}, \quad (\text{active conservation})$$

$$\text{for } i > p: \quad x'_{i,j} = x_{i,j} + \sum_{\substack{k,\ell: \\ (k,\ell,i,j) \in \mathcal{T}}} \gamma_i \delta_{k,\ell,i,j} - \sum_{\substack{k,\ell: \\ (i,j,k,\ell) \in \mathcal{T}}} \delta_{i,j,k,\ell}, \quad (\text{passive conservation})$$

The conservation equations for active CFMMs are unchanged (we assume pool-to-pool transfers between active CFMMs do not incur trading fees), while the conservation constraints for passive CFMMs reflect their respective fees.

Agents also typically incur a small fixed cost with each trade. Angeris *et al.* [4] observe that incorporating such costs in the optimization program transforms a convex program into a mixed-integer convex program. Exact solutions for such programs may be computationally expensive, but efficient approximate solvers exist (see the cited reference for details).

The optimization program could be extended in a number of ways. It might impose additional constraints, such as enforcing upper or lower bounds on pool sizes. The objective function could be extended to take into account costs such as inventory risk or gas fees.

8.2 Practical Issues

Rebalancing requires active CFMMs to move tokens directly from one pool to another without changing LPs' ownership shares. Existing CFMMs (*e.g.*, [1, 14, 7, 10, 19]) provide mechanisms for LPs to add or remove liquidity, but none currently provide the exact functionality needed to support rebalancing.

Rebalancing also raises security concerns. When presented with a rebalancing proposal, it is trivial for a CFMM to check that its own liquidity does not decrease. However, verifying that the proposal maximizes the global objective function is expensive in gas costs. For this reason, the LPs might instead rely on a trusted third party to manage rebalancing, in the same way DEXes currently engage services like Chainlink [8] to provide oracles, or LayerZero [21] or Axelar [18] to provide cross-chain messaging. In the context of the Ethereum Virtual Machine, one could instead install a deterministic solver as a “precompiled contract” [9], thereby providing a low-cost way for a CFMM to verify that a proposed rebalancing is fair and optimal.

In practice, successful defensive rebalancing requires executing a rebalancing transaction before competing arbitrage transactions eliminate the opportunity. Mechanisms for achieving this goal, including transaction sequencing, private order flow, auctions, or proposer-builder separation, are beyond this paper's scope. One direction for future work is the design of decentralized mechanisms that detect arbitrage opportunities, compute Pareto-efficient rebalancings, and obtain sufficient execution priority to realize those rebalancings in adversarial transaction-ordering environments.

References

- 1 Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 Core, March 2021. URL: <https://uniswap.org/whitepaper-v3.pdf>.
- 2 Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. Constant function market makers: Multi-asset trades via convex optimization, 2021. [arXiv:2107.12484](#).
- 3 Guillermo Angeris and Tarun Chitra. Improved Price Oracles: Constant Function Market Makers. *SSRN Electronic Journal*, 2020. doi:10.2139/ssrn.3636514.
- 4 Guillermo Angeris, Tarun Chitra, Alex Evans, and Stephen Boyd. Optimal Routing for Constant Function Market Makers, 2022. Version Number: 1. doi:10.48550/arXiv.2204.05238.
- 5 Guillermo Angeris, Alex Evans, and Tarun Chitra. When does the tail wag the dog? Curvature and market making. *arXiv:2012.08040 [q-fin]*, December 2020. arXiv: 2012.08040. [arXiv:2012.08040](#).

12:16 Defensive Rebalancing in Networks of AMMs

- 6 Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. A theory of Automated Market Makers in DeFi. *arXiv:2102.11350 [cs]*, April 2021. arXiv: 2102.11350. arXiv:2102.11350.
- 7 Bullish. Solana | Digital Assets | Bullish. URL: <https://www.bullish.com/us/digital-assets/solana>.
- 8 {Chainlink, Inc.}. Chainlink 2.0 and the future of Decentralized Oracle Networks | Chainlink, December 2025. URL: <https://chain.link/whitepaper>.
- 9 EVM codes. Precompiled contracts, 2022. As of 25 January 2026. URL: <https://www.evm.codes/precompiled>.
- 10 CoinEx. CoinEx Automated Market Making (AMM) User Guide, September 2025. URL: <https://support.coinex.com/hc/en-us/articles/4411281577369-CoinEx-Automated-Market-Making-AMM-User-Guide>.
- 11 Vincent Danos, Julien Prat, and Hamza El Khalloufi. Global Order Routing on Exchange Networks | Financial Cryptography and Data Security. FC 2021 International Workshops, March 2021. URL: https://dl.acm.org/doi/10.1007/978-3-662-63958-0_19.
- 12 Nicola Dimitri. Arbitrage in automated market makers. *Frontiers in Blockchain*, Volume 7 - 2024, 2024. doi:10.3389/fbloc.2024.1455070.
- 13 Bhaskar Krishnamachari, Qi Feng, and Eugenio Grippo. Dynamic Curves for Decentralized Autonomous Cryptocurrency Exchanges. *arXiv:2101.02778 [q-fin]*, January 2021. arXiv: 2101.02778. arXiv:2101.02778.
- 14 Fernando Martinelli and Nikolai Mushegian. Balancer: A non-custodial portfolio manager, liquidity provider, and price sensor. <https://balancer.finance/whitepaper/>, 2019. As of 2 February 2021.
- 15 Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- 16 Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. Extended abstract: The effect of trading fees on arbitrage profits in automated market makers. In Aleksander Essex, Shin'ichiro Matsuo, Oksana Kulyk, Lewis Gudgeon, Arian Klages-Mundt, Daniel Perez, Sam Werner, Andrea Bracciali, and Geoff Goodell, editors, *Financial Cryptography and Data Security. FC 2023 International Workshops - Voting, CoDecFin, DeFi, WTSC, Bol, Brač, Croatia, May 5, 2023, Revised Selected Papers*, volume 13953 of *Lecture Notes in Computer Science*, pages 262–265. Springer, 2023. doi:10.1007/978-3-031-48806-1_17.
- 17 Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated Market Making and Loss-Versus-Rebalancing, September 2022. arXiv:2208.06046 [math, q-fin]. arXiv:2208.06046.
- 18 Axelar Research. Axelar network: Connecting applications with blockchain ecosystems, 2022. As of 8 May 2023. arXiv:2011.12783.
- 19 Skate. Skate AMM Fundamentals. URL: <https://docs.skatechain.org/skate-amm/introduction>.
- 20 Ye Wang, Yan Chen, Haotian Wu, Liyi Zhou, Shuiguang Deng, and Roger Wattenhofer. Cyclic Arbitrage in Decentralized Exchanges, January 2022. arXiv:2105.02784 [cs, q-fin]. arXiv:2105.02784.
- 21 Ryan Zarick, Bryan Pellegrino, and Caleb Banister. Layerzero: Trustless omnichain interoperability protocol, 2022. As of 8 May 2023. URL: https://layerzero.network/pdf/LayerZero_Whitepaper_Release.pdf.

A

 Appendix

► **Lemma 6.** *Any arbitrage-free configuration has a global valuation.*

Proof. Suppose the CFMM network trades token types $(T_1, \dots, T_m)$, and let $p_{jk} > 0$ denote the spot price of T_j in units of T_k at any CFMM trading the pair (T_j, T_k) . We first show that arbitrage freedom implies that the product of spot prices around every cycle is 1.

Consider a cycle

$$T_{i_1}, T_{i_2}, \dots, T_{i_r}, T_{i_1}.$$

If

$$p_{i_1 i_2} p_{i_2 i_3} \cdots p_{i_r i_1} > 1,$$

then, by continuity of the CFMM trading functions, a sufficiently small quantity of T_{i_1} can be traded around the cycle, returning a strictly larger quantity of T_{i_1} and thereby yielding an arbitrage profit. Hence

$$p_{i_1 i_2} p_{i_2 i_3} \cdots p_{i_r i_1} \leq 1.$$

Applying the same argument to the cycle traversed in the opposite direction gives the reciprocal inequality, so

$$p_{i_1 i_2} p_{i_2 i_3} \cdots p_{i_r i_1} = 1.$$

Assume first that the token graph is connected. Choose T_1 as a numéraire and set $V_1 := 1$. For any other token T_j , choose a path

$$T_1 = T_{i_0}, T_{i_1}, \dots, T_{i_r} = T_j$$

and define

$$V_j = \prod_{\ell=0}^{r-1} p_{i_{\ell+1} i_\ell}.$$

This definition is independent of the chosen path. Indeed, if two paths from T_1 to T_j produced different values, then traversing one path forward and the other backward would form a cycle whose product of spot prices was not 1, a contradiction.

Finally, suppose a CFMM trades T_j, T_k . Take any path from T_1 to T_k and append the edge from T_k to T_j . By the definition of V_j ,

$$V_j = p_{jk} V_k,$$

and therefore

$$p_{jk} = \frac{V_j}{V_k}.$$

Thus $V = (V_1, \dots, V_m)$ is a consistent valuation.

If the token graph has multiple connected components, apply the above construction independently to each component, choosing an arbitrary token in each component as a numéraire and assigning it value 1. ◀