

Automated Market Making and Loss-Versus-Rebalancing*

Jason Milionis
Department of Computer Science
Columbia University
jm@cs.columbia.edu

Ciamac C. Moallemi
Graduate School of Business
Columbia University
ciamac@gsb.columbia.edu

Tim Roughgarden
Department of Computer Science
Columbia University
a16z Crypto
tim.roughgarden@gmail.com

Anthony Lee Zhang
Booth School of Business
University of Chicago
anthony.zhang@chicagobooth.edu

Initial version: July 31, 2022
Current version: August 13, 2022

Abstract

We consider the market microstructure of automated market making and, specifically, constant function market makers (CFMMs), from the economic perspective of passive liquidity providers (LPs). In a frictionless, continuous-time Black-Scholes setting and in the absence of trading fees, we decompose the return of an LP into a instantaneous market risk component and a non-negative, non-decreasing, and predictable component which we call “loss-versus-rebalancing” (LVR, pronounced “lever”). Market risk can be fully hedged, but once eliminated, LVR remains as a running cost that must be offset by trading fee income in order for liquidity provision to be profitable. We show how LVR can be interpreted in many ways: as the cost of commitment, as the time value for giving up future optionality, as the compensator in a Doob-Meyer decomposition, as an adverse selection cost in the form the profit of arbitrageurs trading against the pool, and as an information cost because the pool does not have access to accurate market prices. LVR is distinct from the more commonly known metric of “impermanent loss” or “divergence loss”; this latter metric is more fundamentally described as “loss-versus-holding” and is not a true running cost. We express LVR simply and in closed-form: instantaneously, it is the scaled product of the variance of prices and the marginal liquidity available in the pool, i.e., LVR is the floating leg of a generalized variance swap. As such, LVR is easily calibrated to market data and specific CFMM structure. LVR provides tradeable insight in both the *ex ante* and *ex post* assessment of CFMM LP investment decisions, and can also inform the design of CFMM protocols.

*The second author thanks Richard Dewey, Craig Newbold, Guillermo Angeris, Tarun Chitra, and Alex Evans for helpful conversations on automated market making. The second author is an advisor to fintech companies. The third author is Head of Research at a16z Crypto, a venture capital firm with investments in automated market making protocols. The first author was supported in part by an unrestricted gift from Gnosis, Ltd. The third author was supported in part by NSF awards CCF-2006737 and CNS-2212745.

1. Introduction

In recent years, automated market makers (AMMs) and, more specifically, constant function market makers (CFMMs) such as Uniswap [Adams et al., 2020, 2021], have emerged as the dominant mechanism for decentralized exchange on blockchains. Compared to electronic limit order books (LOBs), which are the dominant market structure for traditional, centralized exchange-based electronic markets, CFMMs offer some advantages. First of all, they are efficient computationally. They have minimal storage needs, and matching computations can be done quickly, typically via constant-time closed-form algebraic computations. In an LOB, on the other hand, matching engine calculations may involve complex data structures and computations that scale with the number of orders. Thus, CFMMs are uniquely suited to the severely computation- and storage-constrained environment of the blockchain. Second, LOBs are not well-suited to a “long-tail” of illiquid assets. This is because they require the participation of active market makers. In contrast, CFMMs mainly rely on passive liquidity providers (LPs).

In this paper, we consider the market microstructure of CFMMs from the perspective of passive LPs. They contribute assets to CFMM reserves that are subsequently available for trade with liquidity takers, at quoted prices that are algorithmically determined. Our goal is to answer three related questions:

1. CFMMs hold reserves in risky assets. Therefore, their performance is impacted by *market risk*. If this market exposure is hedged, what is the residual value for the LP that remains?
2. In a CFMM, the LPs *commit* to a particular payoff or risky asset demand curve. What is the cost to LPs of giving up this optionality?
3. LPs are compensated with *trading fees*. What is the appropriate rate of fee generation for a CFMM to be a fair investment for LPs?

Our central contribution is the identification and analysis of a running cost component which we call *loss-versus-rebalancing* (LVR, pronounced “lever”), and that simultaneously addresses all these questions.

Specifically, we develop a framework for reasoning about liquidity provision on CFMMs, in a frictionless and continuous-time Black-Scholes setting. Informally, in our framework, the profit-and-loss (P&L) of a liquidity provider can be decomposed according to

$$\text{LP P\&L} = (\text{Rebalancing P\&L}) - \text{LVR} + (\text{Trading Fee Income}). \tag{1}$$

The first term in this decomposition is the P&L of a specific benchmark “rebalancing” strategy. The rebalancing strategy buys and sells the risky asset exactly the same way the CFMM does, but does so at centralized exchange prices, rather than CFMM prices. Thus, an arbitrageur trading against the rebalancing strategy makes zero profits. The rebalancing strategy does not systematically lose money over time: the strategy is exposed to market risk, but this risk can be hedged fully (and costlessly) by dynamically trading the underlying assets.

The second term in the decomposition (1) is a cost term which we call LVR.¹ The identification and analysis of this term is the main contribution of this paper. LVR is defined as the shortfall in the value of the CFMM reserves (exclusive of trading fees, which will be discussed shortly) relative to the value achieved by the dynamic rebalancing strategy. We establish that LVR is a non-negative, non-decreasing, and predictable process. In other words, we quantify how much worse a liquidity provider will do versus the alternative of dynamically trading the underlying assets. We provide a closed-form expression for LVR in terms of model primitives. Instantaneously, LVR is the scaled product of the (instantaneous) variance of asset prices, and the marginal liquidity available in the pool. Thus, in derivatives pricing language, LVR would be described as the floating leg of a generalized variance swap. Mathematically, LVR is the compensator in the Doob-Meyer decomposition of the pool value process under the risk-neutral distribution.

The intuition for LVR is as follows: The rebalancing strategy will sell the risky asset as the price increases, and buy the risky asset as the price drops, both at centralized exchange prices. An LP in the CFMM pool, on the other hand, purchases and sells equal amounts of the risky asset as the rebalancing strategy, but at systematically worse prices than market prices. In a sense, arbitrageurs monetize the fact that the CFMM does not know current asset prices, to trade against the pool in a zero-sum fashion to exploit their superior information, and their arbitrage profits manifest as LVR losses for the CFMM LPs. In this way, LVR can be viewed as an *adverse selection* or information cost.

Another perspective is that passively investing as an LP in a CFMM can be thought of as committing to buying the risky asset in the future if the price decreases, and selling the asset if the price increases. This strategy thus has payoffs that resemble that of a short straddle position, that is, a strategy which sells call options and sells put options. A short straddle position generates a profit if prices end at the same point where they started, due to the premium from selling the options, and loses money if prices increase or decrease substantially, since the strategy loses money on either the call or the put. A passive LP position, in contrast, makes nothing if prices end where they started, but loses money if prices diverge. Holding an LP position is thus analogous to giving away a straddle: losing from the volatility exposure, without collecting the upfront premium. LVR measures the forgone value from failing to collect the premium for selling options.

Of course, CFMMs also have trading fee income, which is the third term in the decomposition (1). These fees are paid by liquidity seeking agents or “noise traders”, that trade against the pool for at least partially idiosyncratic reasons. Since the rebalancing P&L in (1) can be perfectly hedged, our framework suggests that what remains when evaluating an LP investment in a CFMM is the comparison between fee income and LVR. By comparing these two quantities, our framework provides tradeable insight into CFMM LP investment decisions. For example, consider a constant product market maker pool (e.g., Uniswap V2), trading a risky asset with a daily volatility of $\sigma = 5\%$ versus the numéraire (e.g., ETH versus USD). Our analysis suggests that LVR is a daily

¹LVR is distinct from the more commonly known metric of “impermanent loss” or “divergence loss”. In our framework, this latter metric is more accurately described as “loss-versus-holding”, and is not a true running cost. We discuss the distinction in Section 5.

running cost of 3.125 (bp) of the pool value. If the pool has a 30 (bp) fee on traded volume (as in Uniswap V2), then the pool must turn over $3.125/30 \approx 10\%$ of its value on a daily basis, in order to break even. To a first order, assuming that the volatility is fixed and known, investing in the pool is an *ex ante* assessment as to the level of the future realized trading volume relative to this break even quantity.

Similarly, when evaluating LP performance *ex post*, rather than measuring raw LP P&L, one should consider only P&L arising from LVR and fee income, quantities which can be easily computed. This provides a clearer metric for pool performance since hedgeable market risk has been eliminated. LVR can also be used by CFMM protocol designers for guidance to set fees. This is because in a competitive market for liquidity provision, there should be no excess profits for LPs, and hence fees should balance with LVR. For example, since LVR scales with variance, one might imagine fee mechanisms that also scale with variance. Or, alternatively, protocols could be constructed that compare LVR versus fee income in a backward looking window, increasing fees if they are below LVR, and decreasing fees if they are above LVR. More speculatively, our results suggest a potential approach to redesign CFMMs to reduce or eliminate LVR: a CFMM which had access to a reliable and high-frequency price oracle could in principle quote prices arbitrarily close to market prices for the risky asset, thus achieving payoffs arbitrarily close to that of the rebalancing strategy.

To be clear, many of the phenomena discussed above are, to some degree, known formally or informally in the literature or by practitioners (e.g., applying options pricing models to specific CFMMs and observing negative convexity, or analyzing arbitrage profits). We discuss this in detail shortly in Section 1.1. The novelty in the present paper is the careful identification of LVR as a unifying concept and its crisp characterization in closed-form, in a way that rigorously generalizes broadly across CFMM designs and asset pricing models. Beyond this, as described above, our work has simple and direct empirical consequences, to the analysis of CFMM investment decisions, the design of CFMMs, and the quantification of trading fees, for example, beyond what has appeared in the literature.

1.1. Literature Review

Automated market makers have their origin in the classic literature on prediction markets and market scoring rules; see Pennock and Sami [2007] for a survey of this area. Constant function market makers, which are characterized by a invariant or bonding function, build on the utility-based market making framework of Chen and Pennock [2007]. In that framework, utility indifference conditions define a bonding function for binary payoff Arrow-Debreu securities.

More recent interest in AMMs has been prompted by an entirely new application: its functioning as a decentralized exchange mechanism, first proposed by Buterin [2016] and Lu and Köppelmann [2017]. The latter authors first suggested a constant product market maker, this was first analyzed by Angeris et al. [2019]. Angeris and Chitra [2020] and Angeris et al. [2021a,b] apply tools from convex analysis (e.g., the pool reserve value function) to study the more general case of constant function market makers, we employ some of those tools here. Angeris et al. [2021b] also analyze

arbitrage profits, but do not relate them to the rebalancing strategy or express them in closed-form. A separate line of work seeks to design specific CFMMs with good properties by identifying good bonding functions [Port and Tiruvilumala, 2022, Wu and McTighe, 2022, Forgy and Liu, 2021, Krishnamachari et al., 2021].

Black-Scholes-style options pricing models, like the ones developed in this paper, have been applied to weighted geometric mean market makers over a finite time horizon by Evans [2020], who also observes that constant product pool values are a super-martingale because of negative convexity. Clark [2020] replicates the payoff of a constant product market over a finite time horizon in terms of a static portfolio of European put and call options. Tassy and White [2020] compute the growth rate of a constant product market maker with fees.

The above literature typically consider CFMM liquidity providers in isolation. Another direction is the consideration of equilibrium models [Aoyagi, 2020, Aoyagi and Ito, 2021, Capponi and Jia, 2021]. O’Neill [2022] empirically characterizes CFMM performance from the perspective of LPs, see also the work of Lehar and Parlour [2021]. Arbitrage profits are a form of miner extractable value (MEV), Qin et al. [2022] empirically quantifies this and other types of CFMM related MEV including “sandwich” attacks. Sandwich attacks are also considered by Zhou et al. [2021] empirically and by Park [2021] theoretically.

2. Model

In what follows, we describe the frictionless, continuous-time Black-Scholes setting of our model.

Assets. Fix a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{Q})$ where \mathbb{Q} is a risk-neutral or equivalent martingale measure, satisfying the usual assumptions. Suppose there are two assets,² a risky asset x and a numéraire asset y . Without loss of generality, assume that the risk-free rate is zero. Working over continuous times $t \in \mathbb{R}_+$, assume that there is observable external market price P_t at each time t . The price P_t evolves exogenously according to a geometric Brownian motion that is a continuous \mathbb{Q} -martingale, i.e.,

$$\frac{dP_t}{P_t} = \sigma dB_t^{\mathbb{Q}}, \quad \forall t \geq 0,$$

with volatility $\sigma > 0$, and where $B_t^{\mathbb{Q}}$ is a \mathbb{Q} -Brownian motion.³

Trading Strategies. A trading strategy is a process (x_t, y_t) defining holdings in the risky asset and numéraire at each time t . For a trading strategy to be *admissible*, we require that it be adapted, predictable, and satisfy

$$\mathbb{E}^{\mathbb{Q}} \left[\int_0^t x_s^2 P_s^2 ds \right] < \infty, \quad \forall t \geq 0. \quad (2)$$

²This assumption is without loss of generality, we describe the multi-dimensional case where there are $n \geq 2$ assets, none of which need be the numéraire, in Appendix A.

³Subject to appropriate technical conditions, we could allow for considerably more general dynamics, for example allowing the risk-free rate and volatility to be adapted random processes rather than fixed values. We focus on a baseline Black-Scholes model for maximum clarity.

We further restrict admissible trading strategies to be *self-financing*, i.e., to satisfy

$$\underbrace{x_t P_t + y_t - (x_0 P_0 + y_0)}_{\text{P\&L}} = \int_0^t x_s dP_s, \quad \forall t \geq 0. \quad (3)$$

Equation (3) states that the change in the profit of the strategy in a small period of time is equal to the holdings of the risky asset, x_s , times the change in price, dP_s . The total P&L of the strategy is just the integral of these instantaneous changes. Intuitively, a self-financing strategy executes all rebalancing trades at market prices; hence, trades do not affect the profit of the strategy, and no money needs to be injected into the trading strategy. In the special case where P_t is a martingale, so the expected profit from holding the risky asset is zero, any self-financing strategy makes zero profits in expectation, since any strategy which dynamically trades a asset with zero expected returns also has zero expected returns.⁴

The self-financing condition allows us to specify $\{x_t\}_{t \geq 0}$, along with y_0 , and determine $\{y_t\}$ implicitly. The P&L of the resulting self-financing strategy can be directly expressed in terms of $\{x_t\}$ via the right side of (3). Note that, since P_t is a \mathbb{Q} -martingale, the P&L process given by (3) is also \mathbb{Q} -martingale, and by (2) it is square-integrable.

CFMM Pool. The state of a CFMM pool is characterized by the reserves $(x, y) \in \mathbb{R}_+^2$, which describe the current holdings of the pool in terms of the risky asset and the numéraire, respectively. Define the feasible set of reserves \mathcal{C} according to


$$\mathcal{C} \triangleq \{(x, y) \in \mathbb{R}_+^2 : f(x, y) = L\},$$

where $f: \mathbb{R}_+^2 \rightarrow \mathbb{R}$ is referred to as the *bonding function* or *invariant*, and $L \in \mathbb{R}$ is a constant. In other words, the feasible set is a level set of the bonding function. The pool is defined by a smart contract which allows an agent to transition the pool reserves from the current state $(x_0, y_0) \in \mathcal{C}$ to any other point $(x_1, y_1) \in \mathcal{C}$ in the feasible set, so long as the agent contributes the difference $(x_1 - x_0, y_1 - y_0)$ into the pool, see Figure 1a.

Note that we are ignoring any trading fees collected by the pool; these will be discussed later in Section 6. To simplify our analysis, we will also assume that, aside from trading with arriving liquidity demanding agents, the pool is static otherwise. In particular, we assume that the liquidity providers do not add (mint) or remove (burn) reserves over the time scale of our analysis. In other words, LPs are *passive*. Further, we ignore the details of the underlying blockchain on which the pool operates, e.g., any blockchain transaction fees such as “gas” fees, the discrete-time nature of block updates, etc.

We make the following assumption:

⁴In the general case, the risky asset may have positive expected returns due to risk premia; self-financing strategies may thus have positive expected profits, proportional to how much portfolio weight they put on the risky asset. However, the positive expected returns of the strategy derive only from the risk premia on the underlying asset: if the strategy is delta-hedged, it makes zero expected profits.



(a) Transitions between any two points on the bonding curve $f(x, y) = L$ are permitted, if an agent contributes the difference into the pool.

(b) The pool value optimization problem relates points on the bonding curve to supporting hyperplanes defined by prices.

Figure 1: Illustration of a CFMM.

Assumption 1. *There is a population of arbitrageurs, able to frictionlessly trade at the external market price, continuously monitoring the CFMM pool.*

When an arbitrageur interacts with the pool, we assume they maximize their immediate profit by exploiting any deviation from the external market price. In other words, they transfer the pool to a point in the feasible set \mathcal{C} that allows them to extract maximum value assuming that they unwind their trade at the external market price P . Equivalently, because trading in the pool is zero-sum between the arbitrageur and the liquidity provider, arbitrageurs can be viewed as minimizing the value of the assets in the pool reserves. Motivated by this, define the *pool value function* $V: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by the optimization problem [see, e.g., Angeris and Chitra, 2020, Angeris et al., 2021b]

$$V(P) \triangleq \underset{(x,y) \in \mathbb{R}_+^2}{\text{minimize}} \quad Px + y \quad (4)$$

subject to $f(x, y) = L.$

Geometrically, the pool value function implicitly defines a reparameterization of the pool state from primal coordinates (reserves) to dual coordinates (prices); this is illustrated in Figure 1b.

If we denote by V_t the value of the pool reserves at time t , Assumption 1 translates mathematically to the fact that $V_t = V(P_t)$. We assume that the pool value function satisfies:

Assumption 2. (i) *An optimal solution $(x^*(P), y^*(P))$ to the pool value optimization (4) exists for every $P \geq 0$.*

(ii) *The pool value function $V(\cdot)$ is everywhere twice continuously differentiable.*

(iii) *For all $t \geq 0$,*

$$\mathbb{E}^{\mathbb{Q}} \left[\int_0^t x^*(P_s)^2 P_s^2 ds \right] < \infty.$$

Parts (i)–(ii) are easily verified for many CFMMs, see Section 4 for examples. Part (iii) is a square-integrability condition that will be used in Section 3. Assumption 2(i)–(ii) is a sufficient condition for the following:

Lemma 1. For all prices $P \geq 0$, the pool value function satisfies:

(i) $V(P) \geq 0$.

(ii) $V'(P) = x^*(P) \geq 0$.

(iii) $V''(P) = x'^*(P) \leq 0$.

Proof. The first part follows from the fact that $\mathcal{C} \subset \mathbb{R}_+^2$ and $P \geq 0$. The second part is the envelope theorem or Danskin's theorem [Bertsekas, 1971]. The third part follows from the concavity of $V(\cdot)$, as a pointwise minimum of a collection of affine functions. ■

Lemma 1(ii) establishes that the slope of the pool value function is equal to the reserves in the risky asset. Lemma 1(iii) establishes that the pool value function is concave. Note that this concavity does not depend on the nature of the feasible set \mathcal{C} or the bonding function $f(\cdot)$. This part also establishes that the second derivative of the pool value function is the marginal liquidity available at the price level.

3. Loss-Versus-Rebalancing

To understand the economics of liquidity provision, we'd like to understand the evolution of the CFMM pool value process $V_t \triangleq V(P_t)$. The pool value is clearly subject to market risk, since the pool intrinsically holds the risky asset. In order to understand and disentangle the impact of market risk, consider a *rebalancing strategy* which seeks to replicate the risky holdings of the pool in order to mirror the market risk. Intuitively, the rebalancing strategy buys exactly the same quantity of the risky asset as the CFMM does, but does so at the external market price, rather than the CFMM price. Formally, we define the rebalancing strategy to be the self-financing trading that starts initially holding $(x^*(P_0), y^*(P_0))$ (the same position as the CFMM), and continuously and frictionlessly rebalances to maintain a position in the risky asset given by $x_t \triangleq x^*(P_t)$. Then, applying the self-financing condition (3) the rebalancing portfolio has value

$$R_t = V_0 + \int_0^t x^*(P_s) dP_s, \quad \forall t \geq 0. \tag{5}$$

Because of Assumption 2(iii), the rebalancing strategy is admissible and R_t is a square-integrable \mathbb{Q} -martingale.

Define the *loss-versus-rebalancing* (LVR) to be the difference in value between the rebalancing portfolio and the CFMM pool, i.e.,

$$\text{LVR}_t \triangleq R_t - V_t.$$

The following theorem, which is our main result, characterizes the loss-versus-rebalancing:

Theorem 1. *Loss-versus-rebalancing takes the form*

$$\text{LVR}_t = \int_0^t \ell(P_s) ds, \quad \forall t \geq 0, \tag{6}$$

where we define, for $P \geq 0$, the instantaneous LVR by⁵

$$\ell(P) \triangleq -\frac{\sigma^2 P^2}{2} V''(P) \geq 0. \quad (7)$$

In particular, LVR is a non-negative, non-decreasing, and predictable process.

Proof. The smoothness condition of Assumption 2(ii) allows us to apply Itô's lemma to $V(\cdot)$ to obtain

$$\begin{aligned} dV_t &= V'(P_t) dP_t + \frac{1}{2} V''(P_t) (dP_t)^2 \\ &= V'(P_t) dP_t + \frac{1}{2} V''(P_t) \sigma^2 P_t^2 dt \\ &= x^*(P_t) dP_t + \frac{1}{2} V''(P_t) \sigma^2 P_t^2 dt, \end{aligned} \quad (8)$$

where the last step follows from Lemma 1(ii). Comparing with (5), we obtain (6). Finally, the fact that $\ell(P) \geq 0$ follows from Lemma 1(iii). ■

In what follows, we consider Theorem 1 from several perspectives:

Hedging Interpretation. Combining Theorem 1 with (5), the pool value evolves according to

$$V(P_t) = V_t = R_t - \text{LVR}_t = V_0 + \int_0^t x^*(P_s) dP_s - \int_0^t \ell(P_s) ds, \quad \forall t \geq 0. \quad (9)$$

Writing this as a stochastic differential equation,

$$dV_t = dR_t - d\text{LVR}_t = x^*(P_t) dP_t - \ell(P_t) dt, \quad \forall t \geq 0. \quad (10)$$

Here, it is clear that the evolution of the CFMM pool value is driven by two components:

- *Market risk component:* $dR_t = x^*(P_t) dP_t$.
Observe that, since P_t is a \mathbb{Q} -martingale, so is R_t . In this way, this term captures unpredictable market risk. Moreover, the “delta” or market exposure is exactly the same as that of the rebalancing strategy, i.e., $x^*(P_t)$.
- *Running cost component:* $d\text{LVR}_t = \ell(P_t) dt$.
Since $\ell(P_t) \geq 0$, this component will be non-negative and non-decreasing. In this way, it can be interpreted as a running cost. Further, this component is differentiable and hence locally predictable.

Observe that since $\text{LVR}_t \geq 0$, we have that $R_t \geq V_t$. In other words, the rebalancing strategy has dominating payoffs and is thus a “super-replicating” hedging strategy. It can also be viewed as the instantaneous delta hedging strategy for the CFMM: by “shorting” the rebalancing strategy

⁵Note that LVR can also be expressed directly in terms of the bonding function $f(\cdot)$ rather than in terms of the pool value function $V(\cdot)$, see Theorem 4 in Appendix D.

(taking the opposite positions) in concert with their CFMM investment, the LPs can eliminate all local or instantaneous market risk, and what remains is a predictable increasing loss process.

Doob-Meyer Decomposition. Observe that, since $V(\cdot)$ is concave, and P_t is a \mathbb{Q} -martingale, applying Jensen’s inequality,

$$\mathbb{E}^{\mathbb{Q}} [V_t | \mathcal{F}_s] = \mathbb{E}^{\mathbb{Q}} [V(P_t) | \mathcal{F}_s] \leq V \left(\mathbb{E}^{\mathbb{Q}} [P_t | \mathcal{F}_s] \right) = V(P_s) = V_s, \quad \forall 0 \leq s \leq t.$$

Hence, V_t is a non-negative \mathbb{Q} -supermartingale. In this case, the classic Doob-Meyer decomposition theorem states that V_t can be *uniquely* decomposed according to $V_t = M_t - A_t$, where M_t is a \mathbb{Q} -martingale, and A_t (the “compensator”) is a predictable non-decreasing process starting from zero. Here, we have $M_t = R_t$ and $A_t = \text{LVR}_t$.

Relationship to Option Strategies. In order to give intuition for the cause of the LVR, first consider the rebalancing strategy. Observe that, from Lemma 1(iii), $x^*(\cdot)$ is non-increasing. If prices decrease slightly from P_0 to $P_t < P_0$, the rebalancing strategy responds by buying the risky asset. The rebalancing strategy thus makes a profit, relative to simply holding the initial position $x^*(P_0)$, if prices increase back to P_0 , and makes a loss if prices decrease further from P_t . This argument holds symmetrically for price decreases, implying that the rebalancing strategy makes losses if prices diverge from P_0 , and profits when prices make small movements away from P_0 and back. In the special case where the risky asset’s price is a random walk, the rebalancing strategy thus breaks even on average. In contrast, when prices move away from P_0 and back, the CFMM reverts to the initial value $V(P_0)$, exactly breaking even: there is no profit from price convergence, to offset the losses the CFMM makes when prices diverge from P_0 .

For another intuition behind LVR, note that the CFMM’s asset position and value are both path-independent: if the price is P_T at time T , the CFMM holds $x^*(P_T)$ of the risky asset and has pool value $V_T = V(P_T)$, regardless of the path that prices took to reach P_T . We can thus think of replicating the payoffs of the CFMM using European options, which also have path-independent payoffs. Relative to the initial holding $x^*(P_0)$, the CFMM commits to selling the risky asset if prices increase, and buys the asset if prices decrease. Thus, the CFMM behaves like a short straddle position: a strategy which sells European call options and put options expiring at time T . A short straddle has positive payoffs if P_T is close to P_0 , since the seller collects the premium from selling options, but does not lose much on the options. The straddle loses money if P_T is much larger or smaller than P_0 , since either the call or the put will expire deep in-the-money. A CFMM LP position similarly loses money if P_T is very different from P_0 , but makes nothing if $P_T = P_0$. The LP position can thus be thought of as giving away a short straddle position, without collecting the upfront option premium.

To make this analogy precise, consider an LP that commits to a locked investment in a pool over the interval $[0, T]$. Define $W_{0:T}$ to be the value of this investment at time zero. Applying the

risk-neutral pricing equation,

$$W_{0:T} = \mathbb{E}^{\mathbb{Q}}[V_T] = \mathbb{E}^{\mathbb{Q}}[R_T - \text{LVR}_T] = R_0 - \mathbb{E}^{\mathbb{Q}}[\text{LVR}_T] = V(P_0) - \mathbb{E}^{\mathbb{Q}}[\text{LVR}_T],$$

where we use the fact that R_T is a \mathbb{Q} -martingale. In the final equation, the first term $V(P_0)$ corresponds to the “intrinsic value” of the pool and is its initial reserve value. The second term, $\mathbb{E}^{\mathbb{Q}}[\text{LVR}_T]$ is the “time value”, i.e., the discount due at time zero in exchange for commitment. In other words, expected LVR is analogous to the premium due to a seller of covered calls. This can typically be computed in closed-form via

$$\mathbb{E}^{\mathbb{Q}}[\text{LVR}_T] = \int_0^T \mathbb{E}^{\mathbb{Q}}[\ell(P_t)] dt. \quad (11)$$

Arbitrage Profits and Adverse Selection. A natural question to consider in our setting is the level of arbitrage profits. This is also considered by Angeris et al. [2021b], who informally argue that:⁶

Proposition 1. *The profit earned by arbitrageurs over the time interval $[0, T]$ is given by*

$$\text{ARB}_T \triangleq V(P_0) + \int_0^T x^*(P_t) dP_t - V(P_T).$$

Comparing with (5), observe that $\text{ARB}_T = R_T - V_T$ in our notation. Then, Proposition 1 combined with Theorem 1 implies that:

Corollary 1. *LVR measures arbitrage profits, i.e., $\text{ARB}_T = \text{LVR}_T$.*

In this way, LVR is a form of “adverse selection”: it is an information cost paid by the LPs to agents with superior information (in this case, arbitrageurs with knowledge of external market prices).

Parametric Dependence and Variance Swap Interpretation. Applying Lemma 1(iii), the instantaneous LVR of (7) can be re-written as

$$\ell(P) = \frac{1}{2} \times (\sigma P)^2 \times |x^{*'}(P)|.$$

Here, the first component, $(\sigma P)^2$, is the instantaneous variance of the price, i.e., for small Δt , $\text{Var}[P_{t+\Delta t}|P_t = P] \approx (\sigma P)^2 \Delta t$. Recalling that $x^*(P)$ is the total quantity of risky asset held by the pool if the price is P , the second component, $|x^{*'}(P)|$ corresponds to the *marginal* liquidity available from the pool at price level P .

Now, integrating over time, we have that

$$\text{LVR}_t = \frac{1}{2} \int_0^t (\sigma P_s)^2 \times |x^{*'}(P_s)| ds, \quad \forall t \geq 0.$$

This expression is the payoff of the floating leg of a continuously sampled generalized variance

⁶For completeness, we provide a rigorous proof in the present setting in Appendix C.

swap [see, e.g., Carr and Lee, 2009], specifically a price variance swap that is weighted by marginal liquidity.

4. Examples

In this section, we consider a number of specific CFMM examples.

Example 1 (Weighted Geometric Mean Market Maker / Balancer). *Consider the bonding function $f(x, y) \triangleq x^\theta y^{1-\theta}$, for $\theta \in (0, 1)$. Solving the pool value optimization (4) allows us to obtain the closed-form optimal solutions*

$$x^*(P) = L \left(\frac{\theta}{1-\theta} \frac{1}{P} \right)^{1-\theta}, \quad y^*(P) = L \left(\frac{1-\theta}{\theta} P \right)^\theta.$$

Then,

$$V(P) = \frac{L}{\theta^\theta (1-\theta)^{1-\theta}} P^\theta, \quad V''(P) = -L \theta^{1-\theta} (1-\theta)^\theta \frac{1}{P^{2-\theta}},$$

and

$$\ell(P) = \frac{\sigma^2}{2} \theta (1-\theta) V(P).$$

In the weighted geometric mean case,⁷ the instantaneous LVR normalized per dollar of pool reserves is constant, i.e.,

$$\frac{\ell(P)}{V(P)} = \frac{\sigma^2}{2} \theta (1-\theta). \quad (12)$$

In fact, with a minor caveat, weighted geometric market makers are the *only* CFMMs for which this is true. We discuss this in Appendix B. Finally, observe that LVR is maximized when $\theta = 1/2$, and goes to zero as $\theta \rightarrow \{0, 1\}$.

Example 2 (Constant Product Market Maker / Uniswap V2). *Taking $\theta = 1/2$ in Example 1, we have that*

$$V(P) = 2L\sqrt{P}, \quad \ell(P) = \frac{L\sigma^2}{4}\sqrt{P}, \quad \frac{\ell(P)}{V(P)} = \frac{\sigma^2}{8}.$$

Assume, for example, that $\sigma = 5\%$ (daily) — this is a reasonable value, for example, for trading ETH versus USD. Then,

$$\frac{\ell(P)}{V(P)} = 3.125 \text{ (bp, daily)}.$$

Example 3 (Uniswap V3 Range Order). *For prices in the liquidity range $[P_a, P_b]$, consider the bonding function of Adams et al. [2021],*

$$f(x, y) \triangleq \left(x + L/\sqrt{P_b} \right)^{1/2} \left(y + L\sqrt{P_a} \right)^{1/2}.$$

⁷See also Proposition 1 of Evans [2020], evaluating a weighted geometric mean market maker over a finite horizon using risk-neutral pricing.

Solving the pool value optimization (4),

$$x^*(P) = L \left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b}} \right), \quad y^*(P) = L \left(\sqrt{P} - \sqrt{P_a} \right).$$

Then, for $P \in (P_a, P_b)$,

$$V(P) = L \left(2\sqrt{P} - P/\sqrt{P_b} - \sqrt{P_a} \right), \quad V''(P) = -\frac{L}{2P^{3/2}},$$

so that

$$\ell(P) = \frac{L\sigma^2}{4} \sqrt{P}.$$

Observe that the instantaneous LVR is the same in Example 2. However, the pool value $V(P)$ is lower. Indeed $V(P) \rightarrow 0$ if $P_a \uparrow P$ and $P_b \downarrow P$, so

$$\lim_{\substack{P_a \rightarrow P \\ P_b \rightarrow P}} \frac{\ell(P)}{V(P)} = +\infty,$$

i.e., the instantaneous LVR per dollar of pool reserves can be arbitrarily high in this case, if the liquidity range is sufficiently narrow. This is consistent with the idea that range orders “concentrate” liquidity.

Example 4 (Linear Market Maker / Limit Order). For $K > 0$, consider the linear bonding function $f(x, y) \triangleq Kx + y$. Solving the pool value optimization (4),

$$x^*(P) = \begin{cases} L/K & \text{if } P < K, \\ 0 & \text{if } P \geq K, \end{cases} \quad y^*(P) = \begin{cases} 0 & \text{if } P < K, \\ L & \text{if } P \geq K. \end{cases}$$

Hence, this pool can be viewed as similar to a resting limit order⁸ that is, depending on the relative value of the price P_t versus limit price K , either an order to buy (if $P_t \geq K$) or an order to sell (if $P_t < K$) up to L/K units of the risky asset at price K . In this case,

$$V(P) = L \min \{ P/K, 1 \}.$$

Observe that $V(\cdot)$ does not satisfy the smoothness requirement of Assumption 2(ii): the first derivative is discontinuous at the limit price $P = K$. Thus, the characterization of Theorem 1 does not apply.⁹

⁸While the linear market maker is *statically* identical to a resting limit order, observe that they are *dynamically* different. In particular, once the price level K is crossed, in a traditional LOB, the limit order is filled and removed from the order book. With a linear market maker, the order remains in the pool at the same price and quantity, but with opposite direction.

⁹Note that the pool value function remains concave and the pool value process is a super-martingale. Hence, from the Doob-Meyer decomposition (see Section 3), a non-negative monotonic running cost process exists. However, this process is not described by (6)–(7). Instead, it can be constructed using the concept of “local time” and the Itô-Tanaka-Meyer formula, but we will not pursue such a generalization here [see, e.g., Carr and Jarrow, 1990].

5. Other Benchmarks and Loss-Versus-Holding

In this section, we consider the possibility of alternative benchmarks aside from the rebalancing strategy. Specifically, given an adapted, predictable process \bar{x}_t , we can define a self-financing trading strategy according to which: (i) initial holdings match the pool, i.e., $(\bar{x}_0, \bar{y}_0) \triangleq (x^*(P_0), y^*(P_0))$; and (ii) risky holdings are held according to \bar{x}_t . We assume that \bar{x}_t satisfies the square-integrability condition (2), so that the resulting trading strategy is admissible. Denote the value of that strategy by \bar{R}_t , so that

$$\bar{R}_t = V_0 + \int_0^t \bar{x}_s dP_s, \quad \forall t \geq 0.$$

Then, we can define the *loss-versus-benchmark* according to $\text{LVB}_t \triangleq \bar{R}_t - V_t$.

One benchmark of particular interest is a strategy that simply holds the initial position, i.e., $x_t^{\text{HODL}} \triangleq x^*(P_0)$, with value

$$R_t^{\text{HODL}} = V_0 + \int_0^t x^*(P_0) dP_s = V_0 + x^*(P_0)(P_t - P_0), \quad \forall t \geq 0.$$

Loss versus the HODL benchmark is often discussed among practioners as “**impermanent loss**” or “**divergence loss**” [e.g., Engel and Herlihy, 2021]. Motivated by the aforementioned analysis, in our view this is more accurately described as “**loss-versus-holding**”: $\text{LVH}_T \triangleq R_T^{\text{HODL}} - V_T$.

The following result characterizes the loss process LVB_t as a function of the underlying benchmark strategy \bar{x}_t :

Corollary 2. *For all $t \geq 0$,*

$$\text{LVB}_t = \text{LVR}_t + \underbrace{\int_0^t [\bar{x}_s - x^*(P_s)] dP_s}_{\triangleq \Delta(\bar{x})_t}.$$

The loss process has quadratic variation

$$[\text{LVB}]_t = [\Delta(\bar{x})]_t = \int_0^t [\bar{x}_s - x^*(P_s)]^2 \sigma^2 P_s^2 ds \geq [\text{LVR}]_t = 0.$$

Therefore, among all benchmark strategies, the rebalancing strategy uniquely defines a loss process with minimal (zero) quadratic variation.

Proof. The first part is an immediate corollary of Theorem 1 and (5). The second part follows from the Itô isometry. ■

Here, we see that, for any choice of benchmark, LVB contains the LVR cost as a component. However, there is a second component, $\Delta(\bar{x})$, with exposure to market risk whenever $\bar{x}_s \neq x^*(P_s)$. This market risk component represents an exposure to the risky asset, but is not a loss. It is a zero-mean \mathbb{Q} -martingale, meaning that it has zero expected return if the underlying risky asset has no risk premium. Even if the risky asset has positive expected returns, the market risk component has

zero payoff when perfectly delta-hedged, meaning that all expected returns derive from exposure to the risk premium of the underlying risky asset. Because of the market risk component, for any non-rebalancing benchmark, LVB is not a true “running cost” — it can be positive or negative; it can revert and is indeed “impermanent”.

The choice of the rebalancing strategy as a benchmark is the *unique* choice which removes all instantaneous market risk. The residual that remains is a true running cost: non-negative and monotonic. Mathematically, the rebalancing benchmark is the *extremal choice* of benchmark that minimizes the quadratic variation of the corresponding loss process to zero.¹⁰

6. Trading Fees

CFMM pools also earn trading fees, which we have ignored up to now. In this section, we will see how fees can balance LVR. Trading fees are typically charged to traders as a fixed fraction of the traded volume. We will consider a simplified and stylized model of trading fees, under the following assumption:

Assumption 3. (i) Fees are charged as a fixed percentage of the traded volume in the numéraire, and are realized as a separate cashflow (in the numéraire) to the LPs (as opposed to, for example, being added to the pool reserves).

(ii) Assumption 1 still holds, i.e., there are arbitrageurs continuously monitoring the pool and they **do not pay fees**. As before, the arbitrageurs enforce the relationship that the value of the pool reserves is given by the pool value function, i.e., $V_t = V(P_t)$.

(iii) There exists a separate population of noise traders. They trade only in the CFMM pool, and trade at least partially for idiosyncratic reasons. The noise traders **do pay fees**. However, they do not impact the pool reserves, since the reserves are immediately rebalanced by arbitrageurs, cf. Assumption 3(ii). Denote by F_T the cumulative fees paid by noise traders up to time T . We assume that F_T is adapted, $\mathbb{E}^{\mathbb{Q}}[F_T^2] < \infty$, and that F_T can be priced with the risk-neutral distribution \mathbb{Q} .

Consider a locked investment in a CFMM pool over a finite time horizon $[0, T]$. The terminal payoff of this investment is the terminal pool reserves’ value plus the cumulative generated fees. Under risk-neutral pricing, the value $W_{0:T}$ at time 0 of this locked investment is given by

$$W_{0:T} = \mathbb{E}^{\mathbb{Q}} [V_T + F_T] = \mathbb{E}^{\mathbb{Q}} [R_T - \text{LVR}_T + F_T] = V(P_0) - \mathbb{E}^{\mathbb{Q}} [\text{LVR}_T] + \mathbb{E}^{\mathbb{Q}} [F_T],$$

where we have used the fact that R_T is a \mathbb{Q} -martingale. This implies that:

¹⁰Note that since it is a continuous, monotonic process, it immediately follows that LVR has zero quadratic variation. See also the discussion of the Doob-Meyer decomposition in Section 3.

Corollary 3. *If the expected fees equal the expected LVR for a time horizon T , i.e.,*

$$\mathbb{E}^{\mathbb{Q}} [F_T] = \mathbb{E}^{\mathbb{Q}} [\text{LVR}_T] = \int_0^T \mathbb{E}^{\mathbb{Q}} [\ell(P_t)] dt,$$

*then the pool is **fairly priced**, i.e., the initial cost of establishing the pool (the value $V(P_0)$ of the reserves) is equal to the time 0 value of the locked, finite-horizon investment, $W_{0:T}$.*

In other words, cumulative fees must be commensurate with LVR for the pool to be fairly priced. Although $\mathbb{E}^{\mathbb{Q}} [\text{LVR}_T]$ can be computed in closed-form for many CFMMs (cf. Section 4), over short time horizons T , we can make the approximation that

$$\mathbb{E}^{\mathbb{Q}} [\text{LVR}_T] = \int_0^T \mathbb{E}^{\mathbb{Q}} [\ell(P_t)] dt \approx \ell(P_0) \times T.$$

Applying this result, we can do a back-of-the-envelope calculation to determine the required trading volume to pay off the LVR cost as in the following example:

Example 2 (Constant Product Market Maker / Uniswap V2, continued). *Assuming $\sigma = 5\%$ (daily) volatility, we have the instantaneous LVR per unit of pool value given by*

$$\frac{\ell(P)}{V(P)} = \frac{\sigma^2}{8} = 3.125 \text{ (bp, daily)}.$$

Assuming trading fees of 30 (bp) of trading volume, as in Uniswap V2, the fair pricing condition implies that the break even daily expected volume must be approximately $3.125/30 \approx 10\%$ of the pool value.

Note that the required volume varies with the instantaneous variance of returns (σ^2). This can vary dramatically depending on market conditions. For example, ETH/USD realized volatility varies from 1.5%–15% (daily) over the time period 1/2020–5/2022 — this is a factor of 10. Over the same interval, LVR costs would vary by a factor of $10^2 = 100$. This implies that trading fees *must* similarly vary over two orders of magnitude for *fair pricing* to hold.

Since LVR changes based on market conditions (and can change dramatically), this suggests the utilization of **dynamic trading fee rules**, e.g., adjusting trading fees based on volatility or variance, or adjusting trading fees with LVR. These adjustments can be based on the recent past (e.g., historical realized volatility or realized LVR), or on future predictions (e.g., options-implied volatility). More speculatively, it may be possible to redesign CFMMs to behave more like the rebalancing benchmark, thus reducing or eliminating LVR. If an AMM had access to a high-frequency oracle for P_t , the AMM could in principle quote prices for trades arbitrarily close to P_t , up to the desired asset position $x^*(P_t)$. Quoting prices this way would reduce arbitrageurs' profits, allowing the AMM to achieve a payoff arbitrarily close to that of the rebalancing strategy. We believe that these are interesting directions for future AMM design research.

References

- Hayden Adams, Noah Zinsmeister, and Dan Robinson. Uniswap v2 core, 2020.
- Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core, 2021.
- Guillermo Angeris and Tarun Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, 2020.
- Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of uniswap markets. *arXiv preprint arXiv:1911.03380*, 2019.
- Guillermo Angeris, Alex Evans, and Tarun Chitra. Replicating market makers. *arXiv preprint arXiv:2103.14769*, 2021a.
- Guillermo Angeris, Alex Evans, and Tarun Chitra. Replicating monotonic payoffs without oracles. *arXiv preprint arXiv:2111.13740*, 2021b.
- Jun Aoyagi. Liquidity provision by automated market makers. *SSRN 3674178*, 2020.
- Jun Aoyagi and Yuki Ito. Coexisting exchange platforms: Limit order books and automated market makers. *SSRN 3808755*, 2021.
- Dimitri P Bertsekas. *Control of uncertain systems with a set-membership description of the uncertainty*. PhD thesis, Massachusetts Institute of Technology, 1971.
- Vitalik Buterin. Let’s run on-chain decentralized exchanges the way we run prediction markets, October 2016. URL www.reddit.com/r/ethereum/comments/55m04x/lets_run_onchain_decentralized_exchanges_the_way/.
- Agostino Capponi and Ruizhe Jia. The adoption of blockchain-based decentralized exchanges. *arXiv preprint arXiv:2103.08842*, 2021.
- Peter Carr and Roger Lee. Volatility derivatives. *Annu. Rev. Financ. Econ.*, 1(1):319–339, 2009.
- Peter P Carr and Robert A Jarrow. The stop-loss start-gain paradox and option valuation: A new decomposition into intrinsic and time value. *The review of financial studies*, 3(3):469–492, 1990.
- Y. Chen and D.M Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, pages 49–56, Vancouver, BC, Canada, 2007.
- Joseph Clark. The replicating portfolio of a constant product market. *Available at SSRN 3550601*, 2020.
- Daniel Engel and Maurice Herlihy. Composing networks of automated market makers. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 15–28, 2021.
- Alex Evans. Liquidity provider returns in geometric mean markets. *arXiv preprint arXiv:2006.08806*, 2020.
- Eric Forgy and Leo Liu. A family of multi-asset automated market makers. *arXiv preprint arXiv:2111.08115*, 2021.
- Bhaskar Krishnamachari, Qi Feng, and Eugenio Grippo. Dynamic automated market makers for decentralized cryptocurrency exchange. In *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–2, 2021. doi: 10.1109/ICBC51069.2021.9461100.
- Alfred Lehar and Christine A Parlour. Decentralized exchanges. Technical report, Working paper, 2021.
- Alan Lu and Martin Köppelmann. Building a Decentralized Exchange in Ethereum, March 2017. URL <https://blog.gnosis.pm/building-a-decentralized-exchange-in-ethereum-eea4e7452d6e>.

Peter O’Neill. Can Markets be Fully Automated? Evidence From an “Automated Market Maker”, January 2022. URL https://raw.githubusercontent.com/petero1111/website/gh-pages/ONeill_JMP_2022.pdf.

Andreas Park. The conceptual flaws of constant product automated market making. *Available at SSRN 3805750*, 2021.

David M Pennock and Rahul Sami. Computational aspects of prediction markets. *Algorithmic game theory*, pages 651–674, 2007.

Alexander Port and Neelesh Tiruvilumala. Mixing constant sum and constant product market makers. *arXiv preprint arXiv:2203.12123*, 2022.

Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214. IEEE, 2022.

Martin Tassy and David White. Growth rate of a liquidity provider’s wealth in $xy = c$ automated market makers, 2020.

Mike Wu and Will McTighe. Constant power root market makers. *arXiv preprint arXiv:2205.07452*, 2022.

Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 428–445. IEEE, 2021.

A. Multi-Dimensional Generalization

In this section, we describe the multi-dimensional generalization of our results. Specifically, denote by vectors $x \in \mathbb{R}_+^n$ the reserves in $n \geq 2$ assets (none of which need be the numéraire), and $P_t \in \mathbb{R}_+^n$ a vector of prices (in terms of the numéraire). We assume that the price vector evolves according to geometric Brownian motion, i.e.,

$$dP_t = \text{diag}(P_t)\Sigma^{1/2} dB_t^{\mathbb{Q}}, \quad \forall t \geq 0,$$

with covariance matrix of returns $\Sigma \in \mathbb{R}^{n \times n}$, $\Sigma \succeq 0$, and where $B_t^{\mathbb{Q}}$ is a standard \mathbb{Q} -Brownian motion on \mathbb{R}^n .

Given a bonding function $f: \mathbb{R}_+^n \rightarrow \mathbb{R}$, define the pool value function $V: \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ according to

$$\begin{aligned} V(P) \triangleq & \underset{x \in \mathbb{R}_+^n}{\text{minimize}} && P^\top x \\ & \text{subject to} && f(x) = L. \end{aligned}$$

Analogous to Assumption 2, we will assume that an optimal solution $x^*(P)$ exists for all $P \in \mathbb{R}_+^n$, that $V(\cdot)$ is twice continuously differentiable, and a suitable square-integrability condition on $x^*(\cdot)$.

Analogous to Lemma 1, we have

Lemma 2. *For all prices $P \in \mathbb{R}_+^n$, the pool value function satisfies:*

- (i) $V(P) \geq 0$.

(ii) $\nabla V(P) = x^*(P) \geq 0$.

(iii) $\nabla^2 V(P) = \nabla x^*(P) \preceq 0$.

Define the rebalancing strategy by $x_t = x^*(P_t)$, with value

$$R_t = V_0 + \int_0^t x^*(P_s)^\top dP_s, \quad \forall t \geq 0.$$

Then, we have the following multi-dimensional analog of Theorem 1:

Theorem 2. *Loss-versus-rebalancing takes the form*

$$\text{LVR}_t = \int_0^t \ell(P_s) ds, \quad \forall t \geq 0,$$

where we define, for $P \geq 0$, the instantaneous LVR

$$\ell(P) \triangleq -\frac{1}{2} \text{tr} \left[\text{diag}(P) \Sigma \text{diag}(P) \nabla^2 V(P) \right] \geq 0.$$

In the case where $\Sigma = \sigma^2 I$, i.e., i.i.d. assets, we have that

$$\ell(P) = -\frac{\sigma^2}{2} \text{tr} \left[\text{diag}(P)^2 \nabla^2 V(P) \right] = -\frac{\sigma^2}{2} \sum_{i=1}^n P_i^2 \frac{\partial^2}{\partial P_i^2} V(P) \geq 0.$$

In particular, LVR is a non-negative, non-decreasing, and predictable process.

Proof. Applying Itô's lemma to $V_t = V(P)$,

$$\begin{aligned} dV_t &= \nabla V(P_t)^\top dP_t + \frac{1}{2} (dP_t)^\top \nabla^2 V(P_t) dP_t \\ &= x^*(P_t)^\top dP_t + \frac{1}{2} \text{tr} \left[\Sigma^{1/2} \text{diag}(P) \nabla^2 V(P_t) \text{diag}(P) \Sigma^{1/2} \right] dt \\ &= dR_t - \ell(P_t) dt. \end{aligned}$$

The rest of the result follows as in the proof of Theorem 1. ■

B. Weighted Geometric Mean Market Makers

Weighted geometric mean market makers have the special property that the instantaneous LVR per dollar of pool value, i.e., $\ell(P)/V(P)$, is a constant. The following theorem establishes that these are essentially the only CFMMs for which this is true:

Theorem 3. *Suppose a CFMM satisfies*

$$\frac{\ell(P)}{V(P)} = c, \quad \forall P \geq 0. \tag{13}$$

Then, we have

$$V(P) = L_1 P^\theta + L_2 P^{1-\theta}, \tag{14}$$

for free constants $L_1, L_2 \geq 0$, where

$$\theta \triangleq \frac{1 - \sqrt{1 - 8c/\sigma^2}}{2} \leq \frac{1}{2}.$$

Comparing with Example 1, observe that (14) states that is the pool can only be the “sum” of θ and $1 - \theta$ weighted geometric mean market makers. The two degrees of freedom are intuitive, since θ and $1 - \theta$ are exchangeable in (12).

Proof of Theorem 3. We construct the following ODE from the (13) along with (7),

$$P^2 V''(P) + \bar{c}V(P) = 0,$$

with constant $\bar{c} \triangleq 2c/\sigma^2$. Make the substitution $P = e^z$, to arrive at the equivalent ODE,

$$V''(z) - V'(z) + \bar{c}V(z) = 0,$$

which when solved, along with the known limit condition from (4) that $V(z) \rightarrow 0$ as $z \rightarrow -\infty$, by the usual method of linear ODEs results in the generic solution,

$$V(P) = L_1 P^{\frac{1 - \sqrt{1 - 4\bar{c}}}{2}} + L_2 P^{\frac{1 + \sqrt{1 - 4\bar{c}}}{2}} = L_1 P^\theta + L_2 P^{1 - \theta}.$$

Note that the above calculation is allowed because the quantity under the root is necessarily non-negative, as if it were not, then $V(P)$ would not be everywhere concave, which must be the case by Lemma 1. ■

C. Miscellaneous Proofs

Proof of Proposition 1. We start with a discrete approximation to the arbitrage profit, indexed by $N \geq 1$. Suppose arbitrageurs arrive sequentially, so that the i th arbitrageur arrives at time τ_i , for $1 \leq i \leq N$. For convenience, set $\tau_0 \triangleq 0$ and $\tau_{N+1} \triangleq T$. For each $1 \leq i \leq N$, at time τ_i , the i th arbitrageur observes the price P_{τ_i} , rebalances the pool from $(x^*(P_{\tau_{i-1}}), y^*(P_{\tau_{i-1}}))$ to $(x^*(P_{\tau_i}), y^*(P_{\tau_i}))$. In other words, the arbitrageur purchases $x^*(P_{\tau_{i-1}}) - x^*(P_{\tau_i})$ units of the risky asset from the CFMM at average price

$$P_i^{\text{CFMM}} \triangleq -\frac{y^*(P_{\tau_i}) - y^*(P_{\tau_{i-1}})}{x^*(P_{\tau_i}) - x^*(P_{\tau_{i-1}})}.$$

The arbitrageur can then sell these units on the external market at price P_{τ_i} and earn profits (in the numéraire) from the difference in price according to

$$(P_{\tau_i} - P_i^{\text{CFMM}}) [x^*(P_{\tau_{i-1}}) - x^*(P_{\tau_i})] = P_{\tau_i} [x^*(P_{\tau_{i-1}}) - x^*(P_{\tau_i})] + [y^*(P_{\tau_{i-1}}) - y^*(P_{\tau_i})].$$

Denote by $\text{ARB}_T^{(N)}$ the aggregate arbitrage profits. Summing over $1 \leq i \leq N$, telescoping the sum, and applying summation-by-parts yields

$$\begin{aligned} \text{ARB}_T^{(N)} &\triangleq \sum_{i=1}^N \left\{ P_{\tau_i} \left[x^*(P_{\tau_{i-1}}) - x^*(P_{\tau_i}) \right] + \left[y^*(P_{\tau_{i-1}}) - y^*(P_{\tau_i}) \right] \right\} \\ &= \sum_{i=1}^N P_{\tau_i} \left[x^*(P_{\tau_{i-1}}) - x^*(P_{\tau_i}) \right] + y^*(P_0) - y^*(P_{\tau_N}) \\ &= P_0 x^*(P_0) + y^*(P_0) + \sum_{i=0}^N x^*(P_{\tau_i}) \left[P_{\tau_{i+1}} - P_{\tau_i} \right] - P_T x^*(P_{\tau_N}) - y^*(P_{\tau_N}). \end{aligned}$$

Observe that the sum in the final expression is the discrete approximation of an Itô integral. Assume that the time partition mesh over $[0, T]$ shrinks to zero as $N \rightarrow \infty$. Taking the limit as $N \rightarrow \infty$ and passing to continuous time, the sum converges to an Itô integral, which is well-defined under Assumption 2(iii). Further, $\tau_N \rightarrow T$, so that $P_{\tau_N} \rightarrow P_T$, and $x^*(P_{\tau_N}) \rightarrow x^*(P_T)$, $y^*(P_{\tau_N}) \rightarrow y^*(P_T)$. Thus, it holds that

$$\text{ARB}_T \triangleq \lim_{N \rightarrow \infty} \text{ARB}_T^{(N)} = V(P_0) + \int_0^T x^*(P_t) dP_t - V(P_T). \quad \blacksquare$$

D. Miscellaneous Results

If the bonding function f is sufficiently smooth, we can express LVR in terms of the gradient and Hessian of f :

Theorem 4. *Under suitable technical conditions,*

$$V''(P) = \frac{\partial_y f}{\partial_{xx} f - 2P \partial_{xy} f + P^2 \partial_{yy} f},$$

where all partial derivatives are evaluated at the point $(x^*(P), y^*(P))$. Then,

$$\ell(P) = -\frac{\sigma^2 P^2}{2} \frac{\partial_y f}{\partial_{xx} f - 2P \partial_{xy} f + P^2 \partial_{yy} f}.$$