Quantifying Loss in Automated Market Makers

Jason Milionis
jm@cs.columbia.edu
Columbia University
Computer Science
New York, NY, USA

Ciamac C. Moallemi
ciamac@gsb.columbia.edu
Columbia University
Graduate School of Business
New York, NY, USA

Tim Roughgarden
tim.roughgarden@gmail.com
Columbia University & a16z Crypto
Computer Science
New York, NY, USA

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

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. 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. 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. For a more complete version of this paper, please refer to https://arxiv.org/pdf/2208.06046.pdf.

CCS CONCEPTS

• **Applied computing → Economics**.

KEYWORDS

blockchain; decentralized finance; automated market makers; mathematical finance; trading and market microstructure; portfolio management

DeFi '22, November 11, 2022, Los Angeles, CA, USA

ACM ISBN 978-1-4503-9882-4/22/11…$15.00
https://doi.org/10.1145/3560832.3563441

ACM Reference Format:
Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. Quantifying Loss in Automated Market Makers. In *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security (DeFi '22), November 11, 2022, Los Angeles, CA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3560832.3563441

1 INTRODUCTION

In recent years, automated market makers (AMMs) and, more specifically, constant function market makers (CFMMs) such as Uniswap [1, 2], have emerged as the dominant mechanism for decentralized exchange on blockchains. 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 of a running cost component which we call *loss-versus-rebalancing* (LVR, pronounced "lever") that simultaneously addresses all these questions.

Informally, in our framework for reasoning about liquidity provision on CFMMs, the profit-and-loss (P&L) of a liquidity provider can be decomposed according to

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

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] 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.

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. To a first order, assuming that the volatility is fixed and known, investing in a pool is an *ex ante* assessment as to the level of the future realized trading volume relative to the break even quantity of the volume required to obtain commensurate fee income with LVR.

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 has 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 Section 2. 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.

2 LITERATURE REVIEW

Automated market makers have their origin in the classic literature on prediction markets and market scoring rules; see Pennock and Sami [14] for a survey of this area. Using AMMs as a decentralized exchange mechanism was first proposed by Buterin [6] and Lu and Köppelmann [12]. Angeris and Chitra [3] and Angeris et al. [4, 5] studied the more general case of constant function market makers. Angeris et al. [5] 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 [15, 17, 10, 11].

Black-Scholes-style options pricing models have been applied to weighted geometric mean market makers over a finite time horizon by Evans [9], who also observes that constant product pool values are a super-martingale because of negative convexity. Clark [7] 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 [16] compute the growth rate of a constant product market maker with fees.

3 SUMMARY OF RESULTS

For the full details, please refer to Milionis et al. [13].

3.1 Model

Consider a frictionless, continuous-time Black-Scholes setting, where $\mathbb{Q}$ is a risk-neutral or equivalent martingale measure, we have a risky asset x and a numéraire asset y, and an observable external market price P_t at each time t, evolving exogenously according to a

[1] LVR 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.

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.

A trading strategy is a self-financing process (x_t, y_t) defining holdings in the risky asset and numéraire at each time t. 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 C according to

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

where $f\colon \mathbb{R}_+^2 \to \mathbb{R}$ is referred to as the *bonding function* or *invariant*, and $L \in \mathbb{R}$ is a constant.

Note that we are ignoring any trading fees collected by the pool; these will be discussed later. 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.

If we assume that there is a population of arbitrageurs, able to frictionlessly trade at the external market price, continuously monitoring the CFMM pool, then, to maximize their profits, the arbitrageurs would, at each time t, set the implied price of the pool to be exactly equal to the exogenously-determined price at time t [4]. Based on this, define the value of the reserves of the pool at time t as V_t, and the *pool value function* $V\colon \mathbb{R}_+ \to \mathbb{R}_+$ by the optimization problem [see, e.g., 3, 5]

$$\begin{array}{rll} V(P) \triangleq & \underset{(x,y)\in\mathbb{R}_+^2}{\text{minimize}} & Px + y \\ & \text{subject to} & f(x, y) = L. \end{array} \tag{2}$$

3.2 Loss-Versus-Rebalancing

To understand the economics of liquidity provision, we'd like to understand the evolution of the CFMM pool value process $V(P_t)$. 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 $\big(x^*(P_0), y^*(P_0)\big)$ (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, the rebalancing portfolio has value

$$R_t = V_0 + \int_0^t x^*(P_s)\, dP_s, \quad \forall\, t \geq 0. \tag{3}$$

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

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

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

Theorem 3.1. *Loss-versus-rebalancing takes the form*

$$\mathsf{LVR}_t = \int_0^t \ell(P_s)\, ds, \quad \forall\, t \geq 0, \tag{4}$$

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

$$\ell(P) \triangleq -\frac{\sigma^2 P^2}{2} V''(P) \geq 0. \tag{5}$$

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

3.3 Loss-Versus-Holding

In this section, we show how LVR relates to what is often discussed among practitioners as "**impermanent loss**" or "**divergence loss**" [e.g., 8]. In our view this is more accurately described as "**loss-versus-holding**". More specifically, consider the strategy that simply holds the initial position, i.e., $x_t^{\mathsf{HODL}} \triangleq x^*(P_0)$, with value

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

Then, loss-versus-holding is $\mathsf{LVH}_T \triangleq R_t^{\mathsf{HODL}} - V_t$. By comparing the expressions of (3) with the above, we can see that for all $t \geq 0$,

$$\mathsf{LVH}_t = \mathsf{LVR}_t + \int_0^t \left[x^*(P_0) - x^*(P_s)\right] dP_s.$$

Here, we see that LVH contains the LVR cost as a component. The second component, though, has exposure to market risk whenever $x^*(P_s) \neq x^*(P_0)$. 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. Because of that component, LVH is not a true "running cost" — it can be positive or negative; it can revert and is indeed "impermanent".

ACKNOWLEDGMENTS

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.

REFERENCES

[1] Hayden Adams, Noah Zinsmeister, and Dan Robinson. 2020. Uniswap v2 core. (2020).
[2] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. 2021. Uniswap v3 core. (2021).
[3] Guillermo Angeris and Tarun Chitra. 2020. Improved price oracles: constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 80–91.
[4] Guillermo Angeris, Alex Evans, and Tarun Chitra. 2021. Replicating market makers. *arXiv preprint arXiv:2103.14769*.
[5] Guillermo Angeris, Alex Evans, and Tarun Chitra. 2021. Replicating monotonic payoffs without oracles. *arXiv preprint arXiv:2111.13740*.

[6] Vitalik Buterin. 2016. Let's run on-chain decentralized exchanges the way we run prediction markets. Reddit Post. (Oct. 2016). Retrieved July 31, 2022 from www.reddit.com/r/ethereum/comments/55m04x/lets_run_onchain_decentralized_exchanges_the_way/.

[7] Joseph Clark. 2020. The replicating portfolio of a constant product market. *Available at SSRN 3550601*.

[8] Daniel Engel and Maurice Herlihy. 2021. Composing networks of automated market makers. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 15–28.

[9] Alex Evans. 2020. Liquidity provider returns in geometric mean markets. *arXiv preprint arXiv:2006.08806*.

[10] Eric Forgy and Leo Lau. 2021. A family of multi-asset automated market makers. *arXiv preprint arXiv:2111.08115*.

[11] Bhaskar Krishnamachari, Qi Feng, and Eugenio Grippo. 2021. Dynamic automated market makers for decentralized cryptocurrency exchange. In *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 1–2. DOI: 10.1109/ICBC51069.2021.9461100.

[12] Alan Lu and Martin Köppelmann. 2017. Building a Decentralized Exchange in Ethereum. en. (Mar. 2017). Retrieved July 31, 2022 from https://blog.gnosis.pm/building-a-decentralized-exchange-in-ethereum-eea4e7452d6e.

[13] Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. Automated market making and loss-versus-rebalancing. (2022). DOI: 10.48550/ARXIV.2208.06046.

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

[15] Alexander Port and Neelesh Tiruviluamala. 2022. Mixing constant sum and constant product market makers. *arXiv preprint arXiv:2203.12123*.

[16] Martin Tassy and David White. 2020. Growth rate of a liquidity provider's wealth in $xy = c$ automated market makers. (2020).

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