

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: May 3, 2026

Abstract

Automated Market Makers (AMMs) are both liquidity sources and investment vehicles for market participants. This paper analyzes the risks and returns of liquidity provision (LP) investments in AMMs. In a continuous-time model, we show that LP returns decompose into a beta-like component reflecting market risk exposure, and an alpha-like component reflecting microstructural forces: accrued fees minus losses to arbitrageurs. Applying our decomposition to the Uniswap v2 ETH-USDC pool, we find that over 99.991% of LP return variance is driven by beta exposure to market risk.

1. Introduction

In traditional financial markets, retail traders are generally liquidity takers, transacting against quotes set by large, specialized market-making firms. These firms earn profits through bid-ask spreads, which represent transaction costs for retail investors. In decentralized financial markets, *automated market makers (AMMs)* take on many of the functions traditionally performed

*The second author thanks Richard Dewey, Craig Newbold, Guillermo Angeris, Tarun Chitra, and Alex Evans for helpful conversations on automated market making. We are also grateful to Austin Adams, Jun Aoyagi, Eric Budish, Larry Glosten, Gur Huberman, Mingxuan He, Thomas Rivera, Xin Wan, and Tianyi Zhang for helpful comments. The second author was supported by the Briger Family Digital Finance Lab at Columbia Business School and is an advisor to fintech companies including Paradigm, a venture capital firm with investments in automated market making protocols. 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 NSF award CNS-2212745, and in part by an unrestricted gift from Gnosis, Ltd. The third author was supported in part by NSF awards CCF-2006737 and CNS-2212745.

by market-making firms. Traders can interact with AMMs as liquidity takers, executing trades at posted prices. But unlike in traditional markets, users can also become liquidity providers (LPs) by contributing capital to the AMM’s pools. In return, they earn a share of trading fees and bear the associated risks. AMMs thus play a dual role in DeFi: they are both trading venues and investment vehicles.

A growing literature studies AMMs from a microstructural perspective, analyzing how fee-paying noise traders and informed arbitrageurs impact AMM profitability. This paper asks a complementary question: how do automated market makers behave as *investments* for liquidity providers? What are the risks and returns of LP investments in AMMs, and how do they compare to investing directly in the assets held and traded by AMMs?

A core principle in asset pricing is that the returns of complex risky assets can be decomposed into simpler components. A stock’s return reflects a combination of factor risk premia and a stock-specific alpha; a corporate bond splits into the risk-free rate, term premia, and a credit risk premium; an option embeds directional and volatility exposures to the underlying stock. These decompositions allow us to think of asset returns as a combination of terms common to many assets, and terms unique to the asset at hand: a stock’s alpha, a corporate bond’s credit risk, or an option’s implied volatility.

In this spirit, this paper proposes a model-based method to decompose the returns of AMM LP positions into two components. The first is a “beta-like” component, reflecting the *market risk exposures* of the assets held by the AMM. We call this component of payoffs the *rebalancing strategy*. The second, equal to the difference between the AMM’s returns and the rebalancing strategy’s returns, is an “alpha-like” component, which reflects *microstructural* forces: accrued fees from noise traders minus adverse selection losses to arbitrageurs.

We develop our model in a simple continuous-time model. There is a risky asset and a numéraire, which can be traded on an AMM and an infinitely deep centralized exchange (CEX). As in the classic Black and Scholes [1973] setting, we assume the risky asset’s price follows a geometric Brownian motion, with possibly stochastic volatility. Noise traders trade with the AMM, contributing fees to LPs. Arbitrageurs trade with the AMM and the CEX to maximize profits. We assume arbitrageurs pay no fees, implying that arbitrageurs ensure that the AMM’s price is always equal to the CEX price. In the case that the AMM is a CFMM, the CFMM is described by an invariant curve

$f(x, y) = L$; the CFMM is willing to make any trade such that it stays on this level curve.

Within our model, the profit and loss of the CFMM decomposes into two distinct components. The first reflects the AMM’s role as a *market maker*: LPs profit from any fees that accrue to the pool, and make adverse selection losses from arbitrageurs “sniping” the AMM’s stale quotes when CEX prices move. Adverse selection losses admit a clean expression in our model, depending only on the instantaneous variance of asset prices, and the marginal liquidity of the AMM’s current level set. We call this component “alpha-like”, since it is unique to the microstructural tradeoffs facing the AMM.

The second component reflects the AMM’s dual role as a *mutual fund*. An AMM providing liquidity in the ETH-USDC holds a long position in ETH; it thus profits when ETH prices increase, and makes losses when ETH prices decrease, just as an actively managed mutual fund would. In our model, this component can be *replicated* through an active trading strategy, which always holds the same amount of the risky asset as the AMM does, but makes all trades at CEX prices; we call this replication strategy the *rebalancing strategy*. We call this component “beta-like”, since it reflects simply directional exposure to underlying risky asset prices, rather than forces unique to the AMM as an investment.

In equity pricing models, stock betas are estimated by *projecting* stock returns onto risk factors; the resultant residual defines the stock’s alpha. Analogously, we show that the rebalancing strategy corresponds to the quadratic-variation-minimizing projection of AMM LP returns onto the risky asset’s price movements, and loss-versus-rebalancing corresponds to the projection residual.

This perspective sheds light on the economic differences between the two return components. The “beta-like” rebalancing strategy is a risk-neutral martingale: since it simply holds a position in the risky asset, it only makes profits and losses to the extent that the risky asset has positive risk premia. In other words, the rebalancing strategy can be perfectly delta-hedged, leading its return to be identically zero. The “alpha-like” component is the difference between two strictly increasing processes: accrued fees from noise trade, and accrued losses to arbitrage traders. Just like stock alphas can be positive or negative, our alpha-like component can be greater or smaller than raw AMM LP returns. Either way, the alpha-like component captures the component of AMM returns that is not mechanically driven by beta exposure to the risky asset.

We use our model to empirically analyze the Uniswap v2 ETH-USDC trading pair. Implement-

ing our procedure empirically is straightforward: the analyst only needs a time series of AMM holdings, mints, and burns, and a time series for the price of the risky asset. We construct discrete approximations to the continuous-time rebalancing strategy, which trade to match the AMM’s risky asset holdings at frequencies between 1 minute and multiple months. We then consider how subtracting these strategies’ profits from pool profits changes the volatility of the P&L series.

The value of raw v2 ETH-USDC pool P&L is fairly volatile: daily gains or losses of over \$10 million USD are common. However, the vast majority of this risk is simply market risk: subtracting the profits of the rebalancing strategy reduces the variance of LP returns by four orders of magnitude. Rebalancing frequencies between 1 minute and 1 hour produce similar results; volatility is higher for longer periods between rebalancing, consistent with the idea that more market risk accrues within-periods. Our model also implies that better approximations to the rebalancing strategy should reduce P&L volatility, not to zero, but to the limit of trading fees minus LVR; this prediction holds in our data.

Our results suggest that, in empirical studies of AMM LP returns, it is quantitatively important to hedge out market risk. Suppose an econometrician wishes to study the influence of some x -variable on the microstructural component of AMM LP returns. In our analysis of the v2 ETH-USDC pair, 99.991% of pool P&L variance is driven by “beta” exposure to ETH prices; when this component is subtracted, the remaining “alpha” component of fees minus adverse selection costs, which is presumably the target of the econometrician’s analyses, represents only 0.009% of variance. In best case for the econometrician, where market risk is simply orthogonal noise to the x -variables, using raw P&L as the y -variable would lead to around 108 times larger coefficient standard errors, so around $108^2 = 11,664$ times more data are needed to achieve similar coefficient precision. Moreover, even small correlations between x -variables and market risk can introduce omitted-variable bias which overwhelms any real effect of x -variables on the AMM “alphas”.

A common approach in industry practice is “impermanent loss”, often defined as the difference between the current value of an AMM’s reserves and the value of a static portfolio holding the same assets as the AMM at some fixed point in the past. Impermanent loss has a number of undesirable features as a loss metric. It does not aggregate cleanly over time: it can be positive on two short time intervals, but zero on the two intervals combined. Its value, and even its sign, can be manipulated through the analyst’s arbitrary choice of start point.

Our approach fixes these issues simply by frequently updating the holdings benchmark: our “alpha-like” component is equivalent to dividing a long time interval into many small intervals, calculating impermanent loss on each subinterval, and adding up the results. In mutual funds terms, our approach can be thought of as calculating the fund’s “alpha” on many small time intervals, and adding up these alphas. Like standard return metrics in asset pricing, our loss measure is additive and independent across time: its value in any period does not depend outcomes in the distant past.

Finally, we discuss connections between AMM LP positions and the three classical ways that volatility can be traded: static (European) options, dynamic trading strategies, and variance swaps [Carr and Madan, 2001]. We model the AMM reserves at the static payoff of a pool value function. This relates to Clark [2020], Fukasawa et al. [2022], and Deng et al. [2023], who show that AMM LP payoffs, over any finite time horizon, can be replicated by shorting a bundle of European options. These option positions, in turn, are equivalent to dynamic trading strategies which sells (buys) the asset when prices increase (decrease). In our setting, the rebalancing strategy plays the role of delta-hedging. Finally, a delta-hedged LP position can be thought of as a generalized variance swap, whose payoff over any given time period is equal to realized variance weighted by the marginal liquidity of the AMMs.

1.1. Related Literature

Empirical AMM Literature. Our main contribution to the academic literature is to develop a model-based empirical approach for measuring AMM LP returns. We view this as filling a gap between the existing theoretical and empirical literatures on AMMs. There is a large and developed theoretical literature, mainly analyzing *microstructural* models of AMMs [Aoyagi, 2020, Aoyagi and Ito, 2021, Capponi and Jia, 2021, Lehar and Parlour, 2021, Park, 2021, Lehar et al., 2024, Hasbrouck et al., 2022, Brolley and Zoican, 2023]. Qualitatively, the microstructural forces in our paper are very similar to this literature, as well as classic microstructure models [Glosten and Milgrom, 1985]: AMM LPs benefit from noise trade, and suffer losses to informed arbitrageurs.

The majority of the aforementioned theoretical AMM papers are *one-period* models; without dynamics, we will show that “impermanent loss” is unambiguously the right measure of AMM LP losses. But time is continuous in reality. In mapping the one-period AMM models to data, the

empirical literature has faced a choice of whether to think of time as one long interval, or many shorter intervals. Augustin et al. [2024] takes the former approach, measuring losses relative to holdings at a fixed start date. Aquilina et al. [2024] and Fang [2024] take the latter approach, measuring IL on a number of shorter intervals, and adding the results up. Heimbach et al. [2022] do not specify whether the holdings benchmark is updated. All the above works claim to measure “impermanent loss”, implying that the two approaches measure the same underlying economic object. There is currently no theoretical guidance on which measurement choice is correct, and how these decisions affect the measured economic outcomes.

In a simple dynamic model, we show that AMM LP returns contain both an “alpha-like” microstructural component, and a “beta-like” market risk component. Dynamics are necessary for this decomposition: we show in Section 5 that market risk is not meaningfully present in one-period models. We believe applied researchers should generally try to measure the alpha-like component: the beta-like component reflects the returns to a mechanical trading strategy, and is not specific to the AMM as an investment.

Our decomposition gives clear guidance on how the empirical literature should proceed. The “one-interval” and “many-interval” LP loss metrics in fact measure economically very different things. The “many-interval” metric, used in Aquilina et al. [2024] and Fang [2024], is exactly hedged LP returns, and thus converges to our “alpha-like” component of returns — fees minus arbitrage costs. We think this is generally the right loss measure to use, though we recommend using a higher rebalancing frequency than that used in Aquilina et al. [2024] and Fang [2024]. To avoid confusion, we suggest that researchers using this metric call it *loss versus rebalancing* rather than impermanent loss.

In contrast, “one-interval” impermanent loss, as used in Aquilina et al. [2024], does not fully remove market risk from returns. The resultant metric combines microstructural forces and predictable exposure to risky asset prices, and our empirical analysis suggests that the latter can be much larger than the former. We recommend against using this metric, and we recommend that papers that choose to use this metric argue explicitly why it is meaningful to include market risk exposures in the definition of AMM LP returns, for the purposes of the question being studied.

Theoretical AMM Literature. While the microstructural economics of our model are very similar

to the existing literature, we make a number of modelling simplifications in order to work cleanly in continuous time, which we discuss in detail in Section 3.1. Some of these limitations have been addressed in follow-up work [Milionis et al., 2023a]. Importantly, since our main modeling goal is to provide a framework for measuring LP returns, we purposefully do not impose any equilibrium model of the level of AMM liquidity in this paper: it is not possible to derive a measurement approach which allows detecting abnormal LP returns in a model which strongly assumes that LPs make normal returns. Ultimately, the goal of our model is to explain and motivate our empirical decomposition, which we view as our main contribution to the literature. The empirical approach can be used even if the reader does not agree with all the assumptions of our stylized model.

A number of papers outside the financial economics literature apply tools from convex analysis to analyze CFMMs. Angeris et al. [2019] show analytically and through agent-based simulations that CPMMs are able to closely track the reference prices of underlying assets. In a manner related to this paper, 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. Angeris et al. [2021b] also analyze LP profits, but do not relate them to the rebalancing strategy or perform measurements. Relative to these earlier papers, our contribution is to explicitly decompose LP losses into the returns on the rebalancing strategy, and obtained fees minus LVR for general AMMs, to show how the LVR component corresponds to arbitrageur profits, and to show how these ideas can be used to measure LP losses empirically. Specifically, Angeris et al. [2019] only compute the *expected* value under the risk neutral distribution, only for *constant product market makers*, deriving a dependence to $\sigma^2/8$. We show, for *all constant function market makers*, a *path-by-path* decomposition of the LP returns (across *all paths*) to the rebalancing strategy and characterize the leftover profit-and-loss which is noise trader fees minus LVR.¹

Subsequent to the initial draft of our paper, Hasbrouck et al. [2025] build a model of equilibrium liquidity provision for concentrated-liquidity AMMs, using a continuous-time model with many of the same simplifying assumptions we use. The single-period and continuous-time modelling approaches are thus both used in the finance literature; the approaches each have their strengths

¹In particular, a number of papers have characterized *expected* losses of CFMM LPs—sometimes referred to as expected “impermanent loss”—in more or less general cases.[Angeris et al., 2019, Evans, 2020] As we discuss in Section 8, the (risk-neutral) *expectations* of “impermanent loss” and LVR are equal, but LVR does not contain market risk; we view this as an important conceptual distinction, and it also leads to cleaner empirical measures of LP losses.

and weaknesses, and we expect research to continue in both directions.

Options. Another group of papers draws analogies between CFMMs and options; we discuss the relationship between our results and these papers in detail in Section 9. 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. Lambert [2022] considers a number of related issues. Boueri [2021] considers the profitability of geometric mean market makers under geometric Brownian motion dynamics; his (re-)definition of “impermanent loss” in that setting is equivalent to LVR. Boueri [2021] also does not show the relationship between LVR and arbitrageur profits. Subsequent to the first publicly available version of our paper, Cartea et al. [2022] and Cartea et al. [2023] discussed the concept of “predictable loss”; the “convexity cost” component of predictable loss in their setting is equivalent to LVR.

Industry. Since we introduced the concept, LVR has evolved into a centerpiece in practical applications and analyses, having had wide impact within decentralized finance. LVR and markouts have been used in practice to measure hedged LP returns of AMM pools, including those of concentrated liquidity market makers [Elsts, 2023, CrocSwap, 2022], and a number of industry groups are also working to build AMMs which utilize our insights to attempt to reduce or eliminate LVR [Canidio and Fritsch, 2023, Cata Labs, 2023, Fenbushi Capital, 2023, Algebra Protocol, 2023, Klages-Mundt and Schuldenzucker, 2022, Aori Protocol, 2023, Adams et al., 2024].

2. Institutional Background

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 an 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 CFMMs 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].

An automated market maker (AMM) is a smart contract which allows market participants

to trade one cryptoasset with another directly on the blockchain, rather than using a centralized exchange. Suppose, for example, that a market participant wanted to trade an asset such as ETH held in her blockchain wallet for another asset, such as USDC. She could do so on a custodial centralized exchange (CEX), such as Binance, by “depositing” her ETH by sending it to Binance on the blockchain; trading Binance-custodied ETH for custodied USDC; and then “withdrawing” the USDC, directing Binance to send USDC back into her personal blockchain wallet. This process has a number of costs. The market participant must give the CEX custody over her cryptoassets, exposing her to exchange credit risk. Market participants must have a CEX trading account; many CEXs impose jurisdictional and identification requirements, often imposed by national regulators to satisfy know-your-customer or anti-money-laundering laws, which market participants may be unwilling or unable to satisfy. The CEX must enable trading of the asset pair in question; many tokens with small market caps are simply not listed on CEXs for trade. Moreover, the rules for determining trade priority and prices on CEXs are not always transparent, and CEXs may impose deposit, withdrawal, and settlement fees or delays on their users.

AMMs present an alternative to CEX trading which circumvents many of these costs. Technically, AMMs are “smart contract” wallets, meaning they are blockchain wallets which can hold cryptoassets, but whose behavior is determined fully by blockchain code rather than human discretion. In this example, the AMM wallet holds some inventory of ETH and USDC. When a trader submits a blockchain transaction to trade ETH for USDC, ETH is sent from the user to the AMM’s inventory, and USDC is sent from the AMM’s inventory to a user, in a single atomic transaction. We will discuss how the trade price is determined in Section 3 below. There are no barriers to access to AMMs: any individual can initiate a transaction with any AMM. There is no credit risk, since trades are completed in a single step, and it is not possible to lose the sold assets without gaining the purchased assets. The mechanisms through which trade prices are calculated is comparatively transparent: for example, the Uniswap v2 smart contract consists of under a thousand lines of publicly available code.² Once “deployed” to the blockchain, no one, not even the creator, can modify the code. AMMs thus present an attractive trading option for market participants who do not want exposure to exchange credit risk, who face barriers to accessing CEXs, or who value the transparency of AMM behavior.

²See the Uniswap v2 repository.

The asset inventory AMMs use to trade is provided in a decentralized manner: any market participant can become a “liquidity provider” (LP) in a ETH-USDC AMM, or any other pair of tokens,³ by contributing ETH and USDC to the AMM’s inventory. The AMM pool then trades using these contributed assets as a part of its inventory, and any trading fees accrue proportionately to the liquidity inventory providers depending on her share. At any point, the liquidity providers can withdraw ETH and USDC corresponding to her share of the pool; however, the amounts of ETH and USDC withdrawn will generally differ from what the liquidity provider first contributed, as inventory adjusts with market participants’ trade requests and fees are collected. The economic problem facing LPs, which is the core focus of our paper, is how to evaluate the costs and benefits of providing liquidity to AMMs.

3. 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,⁴ a risky asset x and a numéraire asset y . We assume that the risk-free rate is zero.⁵ There is an infinitely deep centralized exchange, where the risky asset can be traded with zero fees. The price on the centralized exchange is observable, and evolves exogenously according to a geometric Brownian motion that is a continuous \mathbb{Q} -martingale, i.e.,

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

with a stochastic volatility process⁶ given by $\sigma_t > 0$, and where $B_t^{\mathbb{Q}}$ is a \mathbb{Q} -Brownian motion.

This is a standard set of assumptions in the mathematical finance literature, pioneered by Black and Scholes [1973], and accommodates assets with arbitrary risk premia: by Girsanov’s theorem, the

³In particular, a market participant can begin LPing a pair of assets even if no one else is providing liquidity for it, thus effectively listing the pair of assets for trade.

⁴This 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 B.3.

⁵This assumption is without loss of generality. In order to incorporate a non-zero risk-free rate, one need only add a cost of capital for the value of LP assets deposited as pool reserves.

⁶Volatility will be an important input in the analysis that follows. A natural question is how to calibrate volatility as a model parameter. As in the general application of Black-Scholes style models, for *ex ante* analysis of a possible future LP position, an implied volatility is the appropriate input. On the other hand, for *ex post* LP return performance analysis as in Section 7, a historical or realized volatility is appropriate.

risk-neutral measure simply removes the drift term from these assets’ price movements. Throughout the paper, we will mostly state our results using the risk-neutral measure for expositional simplicity.⁷

CFMM pool. The state of a constant function market maker (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.


Example 1. *The constant product market maker is defined by the invariant $\sqrt{xy} = L$.*

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*. We also ignore the details of the underlying blockchain on which the pool operates. In particular, we assume away any blockchain transaction fees such as “gas” fees, and also ignore the discrete-time nature of block updates.

Besides liquidity providers, there are two kinds of agents in the model: arbitrageurs and noise traders.

Arbitrageurs. 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

⁷All of the key probabilistic results of the paper (e.g., Lemma 2, Theorem 1) hold both in the physical and risk-neutral probability measures. The only exception is part of Corollary 1, the implications of this are discussed in Footnote 20.



(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) Arbitrageurs ensures that, when the price is P , pool reserves shift to the point on the bonding curve where the slope is equal to $-P$.

Figure 1: Illustration of a CFMM.

price P . Geometrically, the presence of arbitrageurs implies that, if the price of the risky asset is P , pool reserves will move to the point on the curve $f(x, y) = L$ where the slope of the bonding curve is equal to $-P$, as indicated in Figure 1b.

Noise traders. There is also a population of DEX-specific noise traders. Noise traders trade only in the CFMM pool, and trade for totally idiosyncratic reasons. In most models of equilibrium liquidity provision on decentralized exchanges, it is important that there are at least some noise traders who strictly prefer trading on CFMMs to CEXs, since these traders are the source, since the equilibrium amount of CFMM liquidity is zero in the absence of noise traders [Lehar and Parlour, 2021]. There are many reasons in practice why certain market participants might prefer trading on CFMMs to CEXs. Market participants may not be able or willing to satisfy the know-your-customer requirements imposed by CEXs, or may not trust that centralized exchanges will safely custody their assets. Market participants may, due to jurisdictional restrictions, not have access to a CEX on which a particular asset trades, whereas DEXes have no jurisdictional restrictions; market participants may also value the ability to atomically combine DEX trades with other smart contract operations on blockchains.⁸

Noise traders' trades have an initial impact on CFMM pool prices, but these effects are immediately offset by arbitrageurs, who immediately move the CFMM back to the CEX price P .

⁸There could also be noise trading on the CEX; however, since we have assumed the CEX is infinitely deep, these trades do not affect prices and can be ignored in our model. If there are noise traders who can trade on either the CEX or the DEX, in our model they will choose to trade on the infinitely deep CEX. For the purposes of our model, we only need that there are some noise traders who exogenously prefer to trade on the DEX, with strong enough preferences that they are willing to use the DEX despite higher fees and price impact.

Thus, from the LP’s perspective, noise traders simply contribute a flow of fees. Denote by FEE_T the cumulative fees paid by noise traders up to time T . For simplicity, we assume fees are paid in units of the numéraire; this simplifies the analysis, since fees do not affect the level curve that the CFMM trades on.⁹ In practice, fees are sometimes (e.g., in Uniswap v2 but not in Uniswap v3) reinvested into the pool reserves; another way to think about this assumption is to assume LPs immediately withdraw any accrued fees from the CFMM.

When discussing arbitrageurs, we will distinguish between two conceptually different forms of arbitrage activity. The first, which we call *rebalancing arbitrage*, is arbitrage of a pool when mispricing arises due to movements of the CEX price. The second, which we call *reversion arbitrage*, is arbitrage following the arrival of noise traders who move DEX prices away from P — this type of arbitrage is sometimes called “back-running”. Our model allows us to quantify the magnitude of profits of rebalancing arbitrageurs, but not reversion arbitrageurs.

3.1. Discussion of Assumptions

The basic microstructural components in our model are qualitatively no different from the recent microstructure literature on AMMs [Lehar and Parlour, 2021, Capponi and Jia, 2021, Barbon and Ranaldo, 2021, Hasbrouck et al., 2022, Foley et al., 2023] and classic microstructure models [Glosten and Milgrom, 1985]. As in the literature, AMM LPs earn fees from noise traders — captured in reduced form by the cumulative fee process FEE_t — and incur losses from trades with arbitrageurs.

The key difference between our model and the literature is that we work in continuous-time, whereas most of the microstructure literature uses single-period models. We do this because *market risk* only arises meaningfully in dynamic models. In other words, one-period models suffice to analyze a CFMM’s function as a *market maker*, but its role as a *mutual fund* only becomes evident in multi-period models. We work in continuous-time because it leads to simpler expressions than multi-period discrete-time models. However, in Subsection 5.1, we show that our model’s basic intuitions also hold in a two-step binomial tree.

Our goal is not to construct an *equilibrium model of AMM liquidity provision*, as in much of the recent microstructure literature. Instead, our goal is to build a simple theoretical framework in which LP P&L decomposes into a “beta-like” market risk component, and an “alpha-like” mi-

⁹The same assumption is made by Lehar and Parlour [2021].

crostructural component. Two of our inputs — the liquidity level L and the fee process FEE_t — are typically endogenous in equilibrium models; we instead treat them as primitives. This is because, together with the bonding curve $f(x, y)$ and the volatility process σ_t , these inputs are sufficient statistics for determining the “alpha-like” component of returns. Any two equilibrium models that generate the same L and FEE_t have the same implications for LP P&L.

As an analogy, alphas in classical asset pricing settings capture deviations from frictionless equilibrium models. These alphas may have a variety of drivers, such as new risk factors, bounded rationality, or luck; the measurement approach for alphas does not take a direct stance on the drivers of alphas. Similarly, we provide a framework for measuring LP “alphas” which does not take a stance on the equilibrium models that drive these alphas. Our empirical methodology is thus valid across a wide range of equilibrium models.¹⁰

We assume the CEX is infinitely deep, so trades have no price impact. It is not entirely clear how true this is in practice, and for which assets this is true. On the one hand, Barbon and Ranaldo [2021] and Liao and Robinson [2022] argue that transaction fees are in fact lower on DEXs than CEXs at certain transaction sizes. Moreover, since the barriers to listing tokens on DEXs are so low, many less liquid tokens begin trading on DEXs before they are traded on large CEXs. On the other hand, a large majority of trade volume in the USD-ETH pair, around 90%, occurs on centralized exchanges relative to decentralized exchanges. Moreover, according to our conversations with industry participants, large orders are in practice often executed through OTC desks, so liquidity may be deeper than what is implied by CEX limit order books; the CEX in our model can be thought of as capturing limit order books, OTC desks, and any other sources of liquidity outside the DEX, including in fact other DEXes. Our model is a reasonable approximation so long as any individual DEX does not constitute more than a small fraction of total liquidity in the market available for an asset. Even in settings where this assumption does not hold, our model may still be a useful conceptual benchmark that is cross-sectionally consistent and arbitrage-free,

¹⁰While our primary goal is measurement, our model elements can be used to construct an equilibrium model of AMM liquidity. For example, Hasbrouck et al. [2025] derives a stationary equilibrium liquidity distribution in concentrated-liquidity AMMs by solving for L values that yield zero risk-neutral expected profits for LPs, net of fees. It is also in principle possible to design P&L metrics which compare fees and arbitrage losses in a way incorporate optimizing LP behavior: for example, if the econometrician is willing to make the strong assumption that FEE_t is independent of L , the econometrician could calculate some liquidity level L^* such that fees equal arbitrage losses, and use $L - L^*$ or $\frac{L}{L^*}$ as a loss metric. Most such metrics should reduce to some way to compare accrued fees and arbitrage losses. Different metrics may be more natural for different settings, and we do not pursue such approaches further in this paper.

analogous to the use of option pricing models to value options with illiquid underlying assets, such as employee stock options in privately held companies, or as in real options valuation [Dixit and Pindyck, 1994].

We assume away many frictions to trading which are present in practice: we assume arbitrageurs pay no trading fees on CEXs or DEXs, we ignore gas fees, and we ignore the discrete, block-based nature of trading on blockchain CFMMs. Accounting for fees will imply that the profits arbitrageurs make will tend to be lower than our expressions. In particular, in practice, arbitrageurs tend to engage in “gas races”, bidding high gas fees so that block miners have an incentive to include their arbitrage trades in the blockchain first. These gas races will tend to redirect some of the profits from CEX-DEX arbitrage towards block miners.¹¹ The analysis of arbitrage profits in the presence of fees is the subject of follow-on work [Milionis et al., 2023a], and we defer a more careful discussion of the impact of fees to that work. We also assume noise trader fees are paid in the numéraire, and we assume away minting and burning of LP shares for simplicity. However, we relax both these assumptions in the empirical application in Section 7.

Finally, our stylized model is intended to motivate the decomposition and accompanying empirical methodology, not to serve as a fully realistic representation of AMM trading. This approach follows the tradition of classical asset pricing models like the CAPM [Sharpe, 1964, Lintner, 1965] and APT [Ross, 1976], whose simplifying assumptions motivate tractable empirical tools, despite their lack of literal realism. Our model yields a clean analytical decomposition that clarifies why LP returns can be viewed as the sum of a market-risk component and a microstructural component—two forces that should be separated empirically. Researchers need not accept all of the model’s assumptions to apply the measurement framework it motivates.

4. Market Risk “Beta” and Microstructural “Alpha”

The pool value function $V(P)$. Figure 1b shows that the composition of the CFMM’s reserve pool depends only on the risky asset’s price: at any time t , if the risky asset’s price is P_t , arbitrageurs will move the pool’s reserves to the point on the $f(x, y)$ curve where the slope is $-P_t$. The mark-to-

¹¹CEX-DEX arbitrage is one form of “miner extractable value”, or MEV, a set of circumstances in which miners’ ability to determine the ordering of transactions allows them to extract monetary value; Daian et al. [2020] discusses MEV in detail.

market value of the pool's reserves at any point in time, $P_t x_t + y_t$, is thus also fully determined by the current price P_t . A convenient way to analyze the monetary value of pool reserves at any given point in time is to define the *pool value function* $V: \mathbb{R}_+ \rightarrow \mathbb{R}_+$, as the solution to the optimization problem:

$$\begin{aligned} V(P) \triangleq & \underset{(x,y) \in \mathbb{R}_+^2}{\text{minimize}} && Px + y \\ & \text{subject to} && f(x,y) = L. \end{aligned} \tag{1}$$

Intuitively, at any price P , arbitrageurs maximize profits by choosing the point on the curve $f(x,y) = L$ which minimizes the value of the pool's reserves. This minimizing choice is the tangency point illustrated in Figure 1b; $V(P)$ measures the monetary value of reserves, $Px + y$, at this point. Hence, $V_t = V(P_t)$ at all times.¹²

We assume that the pool value function satisfies:

Assumption 1. 1. *An optimal solution $(x^*(P), y^*(P))$ to the pool value optimization (1) exists for every $P \geq 0$.*

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

3. *For all $t \geq 0$,*

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

Parts 1–2 are easily verified for many CFMMs, see Section 6 for examples. Part 3 is a square-integrability condition that will be used in Section 4. Parts 1–2 are a sufficient condition for the following:

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

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

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

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

¹²Note also that the optimization problem in (1) is isomorphic to the *expenditure minimization problem* from classical demand theory: under price P , a consumer minimizes total expenditures $Px + y$, subject to staying on the indifference curve $f(x,y) = L$. The solution to this problem is to choose the point where the level curve of $f(x,y)$ is tangent to the budget set. The envelope theorem thus gives that the first derivative of the expenditure function is the Hicksian demand function, which is isomorphic to $x^*(P)$, and the second derivative of the expenditure function is the slope of Hicksian demand.

Part 2 of Lemma 1 establishes that the slope of the pool value function is equal to the reserves in the risky asset. Part 3 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.

We call $x^*(P)$ the demand curve of the AMM. We remark that the concavity of the pool value function along with Part 3 of Lemma 1 establish that the demand curves of all CFMMs must be non-increasing functions of the implied pool price. This fact has deep economic roots, in that it is fundamentally implied by and equivalent to Myerson’s lemma when LPs exchanging a risky asset with traders is viewed as an auction for which incentive compatibility of traders reporting their true valuations for the asset is required; this has been the subject of later investigations by Milionis et al. [2024, 2023b], and the monotonicity of demand curves holds for and forms the basis of a broader class of AMMs that are more general than CFMMs, and also include limit order books.

Remark 1. *The setting we accommodate and rely on is that of a locally-smooth demand curve $x^*(P)$. It is neither important nor required that this format comes from a CFMM, i.e., it is not necessary for our results and LVR in particular that there is an invariant curve. So long as for some (instantaneous) interval of time, the LP is committing to a fixed demand curve according to Assumption 1, the definition and characterization of LVR in Lemma 2 carries through precisely as-is with no modification. For one example, this remark means that the LVR calculation of Lemma 2 can be applied with absolutely no adjustments in general concentrated liquidity AMMs, such as Uniswap v3 (in fact, see Example 4 and Example 5, where this is analyzed). This means that the concept of loss-versus-rebalancing, as we present it hereby, applies for any AMM with a locally-smooth demand curve $x^*(P)$. Therefore, LVR does not only refer to and is not a quantity specific to CFMMs; LVR characterizes the behavior of general AMMs; also refer to Footnote 26 for how to apply this to even more general, non-locally-smooth settings.*

The pool value function allows us to write the profit and loss of an CFMM, from time 0 to time t , as:

$$\text{LP P\&L}_t = V_t - V_0 + \text{FEE}_t. \quad (2)$$

In words, LP P\&L_t is the monetary value of the pool reserves at time t , minus the value at time 0,

plus the cumulative fees collected until time t .

Rebalancing strategy R_t . A key component of our decomposition is an object we call the *rebalancing strategy*. Informally, the rebalancing strategy aims to always match the risky asset holdings of the CFMM. Whenever prices move and the CFMM changes its holdings of the risky asset, the rebalancing strategy makes exactly the same net trades; however, it executes these trades at CEX prices, rather than CFMM prices. The rebalancing strategy thus always has the same exposures to risky asset prices as the CFMM, and thus its P&L will correspond to the “market risk” component of our decomposition.

Formally, we first define general 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

$$\mathbb{E}^{\mathbb{Q}} \left[\int_0^t x_s^2 \sigma_s^2 P_s^2 ds \right] < \infty, \quad \forall t \geq 0. \quad (3)$$

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

$$\underbrace{x_t P_t + y_t - (x_0 P_0 + y_0)}_{\text{P\&L}_t} = \int_0^t x_s dP_s, \quad \forall t \geq 0. \quad (4)$$

Self-financing trading strategies can be thought of as strategies which always adjust risky asset positions at market prices. Thus, the profits of any self-financing strategy are unaffected by its *trades* — arbitrary adjustments in x_t do not directly impact profits. Instead, the profits of these strategies are entirely driven by their *market risk exposures*: Equation (4) states that the instantaneous change in a self-financing strategy’s profits is equal to the risky asset position, x_s , times the change in price, dP_s . This also implies that, if we specify y_0 , the initial amount of the numéraire, and $\{x_t\}_{t \geq 0}$, risky asset holdings in all future histories, the path of the numéraire $\{y_t\}$ is implicitly determined through (4). Note that, since P_t is a \mathbb{Q} -martingale, the P&L process given by (4) is also \mathbb{Q} -martingale, and by (3) it is square-integrable.¹³

We then define the rebalancing strategy to be the self-financing trading that starts initially

¹³When the risky asset has no risk premia, P_t is also a physical-probability martingale, so self-financing trading strategies make zero profits in expectation, and conditional on any possible history. In the general case where $\mathbb{Q} \neq \mathbb{P}$ and the asset may have risk premia, the instantaneous expected return on any self-financing strategy is fully driven by the asset’s risk premium, and any self-financing trading strategy makes zero profits when delta-hedged.

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 $x_t \triangleq x^*(P_t)$. Let $R_t = P_t x_t + y_t$ denote the monetary value of the rebalancing strategy at time t ; (4) implies that:

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

Because of Assumption 1 Part 3, the rebalancing strategy is admissible and R_t is a square-integrable \mathbb{Q} -martingale. In particular, being a self-financing strategy, the rebalancing strategy breaks even in expectation under the risk-neutral measure \mathbb{Q} , and only makes expected returns to the extent that the underlying risky asset has nonzero risk premia.

We then express the change in pool value from time 0 to time t as the sum of the rebalancing strategy's profits, and a residual term which we will define as *loss-versus-rebalancing*:

$$V_t - V_0 = (R_t - V_0) - \text{LVR}_t,$$

$$\text{LVR}_t \triangleq R_t - V_t. \quad (6)$$

LVR_t turns out to admit a simple interpretation as the losses of CFMMs to arbitrageurs from price slippage.

Lemma 2. *Loss-versus-rebalancing takes the form:*

$$\text{LVR}_t = \int_0^t \ell(\sigma_s, P_s) ds, \quad \forall t \geq 0, \quad (7)$$

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

$$\ell(\sigma, P) \triangleq \frac{\sigma^2 P^2}{2} |x^{*'}(P)| \geq 0. \quad (8)$$

$\ell(\sigma, P)$ is always positive, so LVR is a non-negative, non-decreasing, and predictable process. Moreover, the cumulative profits of rebalancing arbitrageurs up to time t are equal to LVR_t .

The core intuition behind Lemma 2 is that the CFMM systematically loses money relative to the rebalancing strategy due to *price slippage*: every trade made by the CFMM is made at slightly




Figure 2: LVR and a stylized depiction of CFMM LP price slippage. Suppose prices begin at P_t , the slope of the red line; the CFMM reserves then begin at point A . If prices increase to $P_t + dP_t$, the slope of the brown line, the CFMM trades to point B . The rebalancing strategy trades instead at the price $P_t + dP_t$, to point B^* . LVR is the vertical gap between B and B^* .

worse prices than the rebalancing strategy. LVR is large when volatility σ_t is high, so prices move more; and when $|x^{*'}(P_t)|$ is large, so the CFMM trades more of the risky asset when prices move a given amount. We will refer to $|x^{*'}(P_t)|$ as the *marginal liquidity* of the CFMM at price P_t .¹⁴ The idea that CFMM LP losses are related to price volatility and bonding function curvature is well-understood in the theoretical literature on AMMs [Aoyagi, 2020, Aoyagi and Ito, 2021, Lehar and Parlour, 2021, Capponi and Jia, 2021, O’Neill, 2022]; the relationship between these quantities is particularly simple in continuous-time models [Evans, 2020, Angeris and Chitra, 2020, Angeris et al., 2021a,b].

We prove Lemma 2 rigorously in Appendix A; here, we present an intuitive derivation based on Figure 2. Suppose the market price changes from P_t to $P_t + dP_t$. Arbitrageurs thus trade with the CFMM, moving from point A to point B on the CFMM invariant curve, and selling dx_t of the risky asset, indicated by the horizontal green line. The rebalancing strategy sells exactly the

¹⁴In Appendix A.3, we derive an expression for $|x^{*'}(P)|$ in terms of the CFMM bonding function $f(\cdot)$, and show that $|x^{*'}(P)|$ is related to the curvature of the level sets of $f(\cdot)$: CFMMs with “flatter” bonding curves have higher marginal liquidity. An interesting implication of these results for the design of CFMM invariants is that, in our model, under our assumptions, the only feature of CFMMs which matters for losses is the curvature of the CFMM bonding function, which determines $x^{*'}(P)$. Any two CFMM invariants for an asset pair which have the same local convexity at any given price have the same LVR.

same amount dx_t , but trades at the CEX price, $P_t + dP_t$; it trades along the brown line of slope $P_t + dP_t$ passing through A , reaching point B^* , which is higher than point B . Thus, after the price change, the CFMM and the rebalancing strategy hold the same amount of the risky asset, but the rebalancing strategy holds $B^* - B$ more cash than the CFMM.

The rebalancing strategy trades at the slope of the brown line, $P_t + dP_t$. The CFMM trades at the slope of the purple line — that is, the secant line connecting points A and B . The tangency lines at points A and B have slopes P_t and $P_t + dP_t$ respectively, so the secant line has slope $P_t + \frac{dP_t}{2}$. Thus, the $B - B^*$ line has height:

$$dx_t \times \left((P_t + dP_t) - \left(P_t + \frac{dP_t}{2} \right) \right) = \frac{dx_t dP_t}{2}. \quad (9)$$

This is thus the CFMM's cash loss relative to the rebalancing strategy. It is also equal to arbitrageur profits: arbitrageurs purchase quantity dx_t from the CFMM at $P_t + \frac{dP_t}{2}$, and sell to the CEX at $P_t + dP_t$, earning $\frac{dx_t dP_t}{2}$.

Next, we write the amount traded dx_t as a function of dP_t and the slope of the CFMM's demand function, $x^*(P_t)$:

$$dx_t = \left| \frac{dx^*(P)}{dP} \right| dP_t = |x^{*'}(P_t)| dP_t, \quad (10)$$

Expression (9) then becomes:

$$\frac{dx_t dP_t}{2} = |x^{*'}(P_t)| \frac{(dP_t)^2}{2}. \quad (11)$$

Now, for a geometric Brownian motion, in a small amount of time dt , the quadratic variation $(dP_t)^2 = d[P_t]$ is equal to $\sigma_t^2 P_t^2 dt$, that is, the instantaneous variance $\sigma_t^2 dt$ multiplied by the square of the price. Plugging in to (11), the instantaneous losses of the CFMM LP position are equal to (8) of Lemma 2:

$$|x^{*'}(P_t)| \frac{\sigma_t^2 P_t^2}{2} dt. \quad (12)$$

The derivation illustrates that CFMMs rebalancing — selling when prices rise and buying when prices fall — is not in itself a source of losses. The rebalancing strategy makes exactly the same trades of the risky asset as the CFMM LP position, but does not lose money because it executes all trades at CEX prices.¹⁵ LVR arises from price slippage: the fact that CFMMs execute all trades

¹⁵For that matter, *any* trading strategy which executes all trades at CEX prices exactly breaks even under the

at worse-than-market prices.

The arbitrage losses of CFMMs are closely related to models of “quote sniping” in high-frequency trading. In Budish et al. [2015], price changes render quotes “stale”, triggering a race between market makers trying to cancel their quotes and arbitrageurs seeking to snipe them. CFMMs in our model can be thought of as market makers that never proactively update quotes: CFMM prices become stale whenever CEX prices move, and loss-versus-rebalancing captures the cumulative losses LPs suffer from arbitrageur “quote sniping”.¹⁶

The basic microstructural tradeoff facing the CFMM LP in our model is no different from Budish et al. [2015] and the classic microstructure literature: LPs gain from fees, which we model as an exogenous increasing process, and lose from trades with arbitrageurs. A subtle difference between our model and that of Budish et al. [2015] and the classic literature is that losses in our setting are driven by *volatility* — knowledge of *current* market prices, which only transmits into the CFMM’s prices through arbitrage. Lemma 2 shows that these losses can be quantified based on the volatility of CEX prices and the CFMM’s mechanical price-quantity schedule. In contrast, *informed trading* in classic microstructure models — trading driven by knowledge of *future* market prices — is much more difficult to quantify.

We now state the main result of our model, which follows immediately from Equation (6) and Lemma 2.

Theorem 1. *We have:*

$$\text{LP P\&L}_t = \text{FEE}_t + V_t - V_0 = \underbrace{\int_0^t x^*(P_s) dP_s}_{\text{market risk “beta”}} + \underbrace{\text{FEE}_t - \text{LVR}_t}_{\text{microstructural “alpha”}} \quad (13)$$

Theorem 1 states that LP P&L_t can be decomposed into a “beta-like” market risk component, and an “alpha-like” microstructural component.

The “beta-like” term reflects *market risk*: directional exposure to the risky asset’s price changes.

risk-neutral measure.

¹⁶We model LPs as fully passive. In practice, LPs could attempt to avoid arbitrage losses either by withdrawing liquidity upon large price movements, or attempting to “snipe” their own quotes before arbitrageurs can. In a frictionless model where LPs and arbitrageurs competed on equal footing whenever CEX prices moved, any possible arbitrage profits from sniping would be dissipated through gas fees; thus, we conjecture that LVR would accurately capture LPs losses in such an environment, though it would not capture arbitrageur profits. CEX-DEX arbitrage is more complicated in reality; modelling this accurately is potentially helpful for understanding AMM behavior quantitatively, but is outside the scope of the current paper.

Intuitively, the CFMM behaves like a rule-following actively managed mutual fund or ETF, mechanically adjusting its risky asset holdings $x^*(P)$ as prices change. The CFMM thus makes instantaneous profits equal to the product of its risky asset holdings, $x^*(P_s)$, and instantaneous changes in risky asset prices, dP_s . Like market or factor betas in equity pricing settings, this component is not *unique* to the CFMM; in our model, it can be attained simply by replicating the AMM’s trades on the CEX, holding $x^*(P)$ when the price is P , as the rebalancing strategy does.

The “alpha-like” term reflects the *microstructural* tradeoff facing the CFMM: as in classic microstructure models [Glosten and Milgrom, 1985], LPs make money from fee-paying noise trade, and lose money to informed arbitrage. Like alphas in equity-pricing settings, this term is comparatively *unique* to the CFMM, in the sense that it cannot be replicated simply by trading the underlying asset.¹⁷

In equity-pricing settings, a stock’s alpha can be isolated by hedging exposures to market or factor risk, by subtracting the returns of “tracking portfolios” with zero alphas and equal betas to the stock. Analogously, we will show in Section 7 that the microstructural “alpha” component can be isolated simply by subtracting the profits of the rebalancing strategy from LP P&L.

5. The Rebalancing Strategy as a Projection

In this section, we show that the rebalancing strategy corresponds to the quadratic-variation-minimizing *projection* of AMM profits onto the risky asset’s price path, analogous to how a stock’s betas are calculated by projecting its returns onto market or factor returns. This gives a precise technical foundation for interpreting the rebalancing strategy as capturing the component of AMM P&L “explained by” movements in the underlying asset’s price. We first build intuition using a simple two-step binomial tree, then prove the result formally in continuous time.

5.1. Binomial Tree

We first analyze a simple two-step binomial tree, illustrated in Figure 3. We consider a constant-product CFMM with invariant $f(x, y) = xy$, which begins holding $(x_0, y_0) = (1, 1)$. CEX price

¹⁷A caveat is that Lemma 2 implies that LVR_t is mechanically related to realized variance, and thus can be replicated by trading instruments such as options and variance swaps; we discuss this in Section 9. By assumption in our model, FEE_t is not in general mechanically related to price movements in the underlying asset, and thus cannot in general be replicated by trading the underlying asset.

movements are depicted in Panel A. Prices begin at $P_0 = 1$; at $t = 1$, they increase to $P_1^U = 1.4$ or decrease to $P_1^D = 0.6$; at $t = 2$, they either revert to $P_2^{DU} = P_2^{UD} = 1$, or diverge further to $P^{UU} = 1.8$ or $P^{DD} = 0.4$. The tree is constructed so that, under zero interest rates, risk-neutral probabilities of up and down movements are all 0.5; at these probabilities, prices are martingales, and any strategy that trades at market prices has zero expected return. We ignore noise trade and fees in the tree for expositional simplicity.

We can calculate the asset positions and profits of three strategies: *buy-and-hold*, which simply holds the initial endowment $(x_0, y_0) = (1, 1)$ forever, in Panel B; the rebalancing strategy in Panel C; and the CFMM in Panel D. For the CFMM, we compute what (x_t, y_t) is as a function of P_t and the initial state (x_0, y_0) ; for example, at price $P^U = 1.4$, we must have $x_0 y_0 = 1$ and $\frac{y_0}{x_0} = 1.4$, implying $(x_1, y_1) = (0.845, 1.183)$. The rebalancing strategy simply mimics the CFMM's changes in x_t , but trades from x_t to x_{t+1} at price P_{t+1} , so its cash position updates as:

$$y_{t+1} = y_t - P_{t+1}(x_{t+1} - x_t)$$

The figure shows the asset holdings of each strategy at each node, as well as the market value of holdings, $P_t x_t + y_t$.

We define *loss versus holding* (LVH) as the difference between the CFMM in panel D, and the buy-and-hold strategy in panel B. *Loss versus rebalancing* is instead the difference between the CFMM in panel D and the rebalancing strategy in panel C. The difference between LVH and LVR—the B-C payoff difference—represents the payoffs of the rebalancing strategy relative to buy-and-hold.

The B-C difference is a *risk-neutral martingale*: since it makes all trades at market prices, it breaks even in risk-neutral expectation at each node. At $t = 1$, Panel B and C payoffs are identical, since rebalancing occurs after price movements. Payoffs diverge at $t = 2$. On the upper path, the rebalancing strategy sells to $x_1^U = 0.845$, holding less of the risky asset than buy-and-hold. This increases payoffs if prices revert to P_2^{UD} (2.062 vs 2.000), and lowers them if prices diverge to P_2^{UU} (2.738 vs 2.800). These gains and losses offset exactly: risk-neutral expected profits from P_1^U onwards are 2.400 for both B and C. Similarly, on the downward path, the rebalancing strategy buys to $x_1^D = 1.291$, outperforms at P_2^{DU} and underperforms at P_2^{DD} , and earns 1.6 in expectation.

The C-D difference is a *strictly increasing process*. C makes the same trades as D, so C and D have identical risky asset holdings x_t in all tree nodes. However, C makes trades at CEX prices, and D trades at the worse AMM prices. Thus, D always underperforms C, and the C-D gap is strictly increasing in t .

Figure 3 thus shows that LVH can be thought of as the sum of the B-C gap, a risk-neutral martingale, and the C-D gap, a strictly increasing process. Removing the B-C gap from LVH does not affect risk-neutral expected losses. Rather, it is useful for *denoising*: it takes out directional exposure to the risky asset’s price, node by node, leaving only losses driven by price slippage.¹⁸

A seemingly attractive feature of loss-versus-holding is that, if prices revert to 1 — as in states P_2^{UD} and P_2^{DU} — LVH is equal to zero. This is the basis of the colloquially popular idea that CFMM losses are “impermanent”: losses relative to buy-and-hold vanish, as long as prices eventually revert. Figure 3 shows where this intuition is misleading. The trading strategy of the CFMM — looking at the movements of x_t , common to Panels C and D — is a bet on price convergence. Executed at market prices, as in the rebalancing strategy in Panel C, this trading strategy profits upon mean-reversion, and loses money upon divergence. The CFMM executes all trades at worse-than-market prices, so it is a very inefficient bet on convergence: it breaks even instead of profiting if prices mean-revert, and loses even more than the rebalancing strategy when prices diverge.

5.2. Continuous Time

With some more technical machinery, we can derive a continuous-time equivalent of the decomposition suggested by Figure 3. We consider a broad class of *benchmark strategies*, which begin holding the same position in the risky asset as the CFMM, and adjust holdings at CEX prices. Formally, we define a benchmark as a self-financing trading strategy, described by a position \bar{x}_t in the risky asset. Initial holdings match the pool, i.e., $(\bar{x}_0, \bar{y}_0) \triangleq (x^*(P_0), y^*(P_0))$. We assume that \bar{x}_t satisfies the square-integrability condition (3), so that the resulting trading strategy is admissible.

¹⁸Following tradition [Cox et al., 1979], we analyze the binomial tree under the risk-neutral measure for simplicity. An interpretation of our results, when the underlying asset may have risk premia and physical probabilities differ from risk-neutral probabilities, is that the expected returns of the buy-and-hold and rebalancing strategies are fully driven by their *deltas* — their linear exposures to the risky asset’s price, which here is exactly their x_t values. For example, if the underlying asset had positive risk premia, the rebalancing strategy would overperform (underperform) buy-and-hold from P^D (P^U) onwards, purely because it holds more of the high-return risky asset; hedging this directional exposure would also eliminate the expected return difference. The risk-neutral measure thus allows us to distinguish between return differences driven by the underlying asset’s risk premium, and residual return differences.

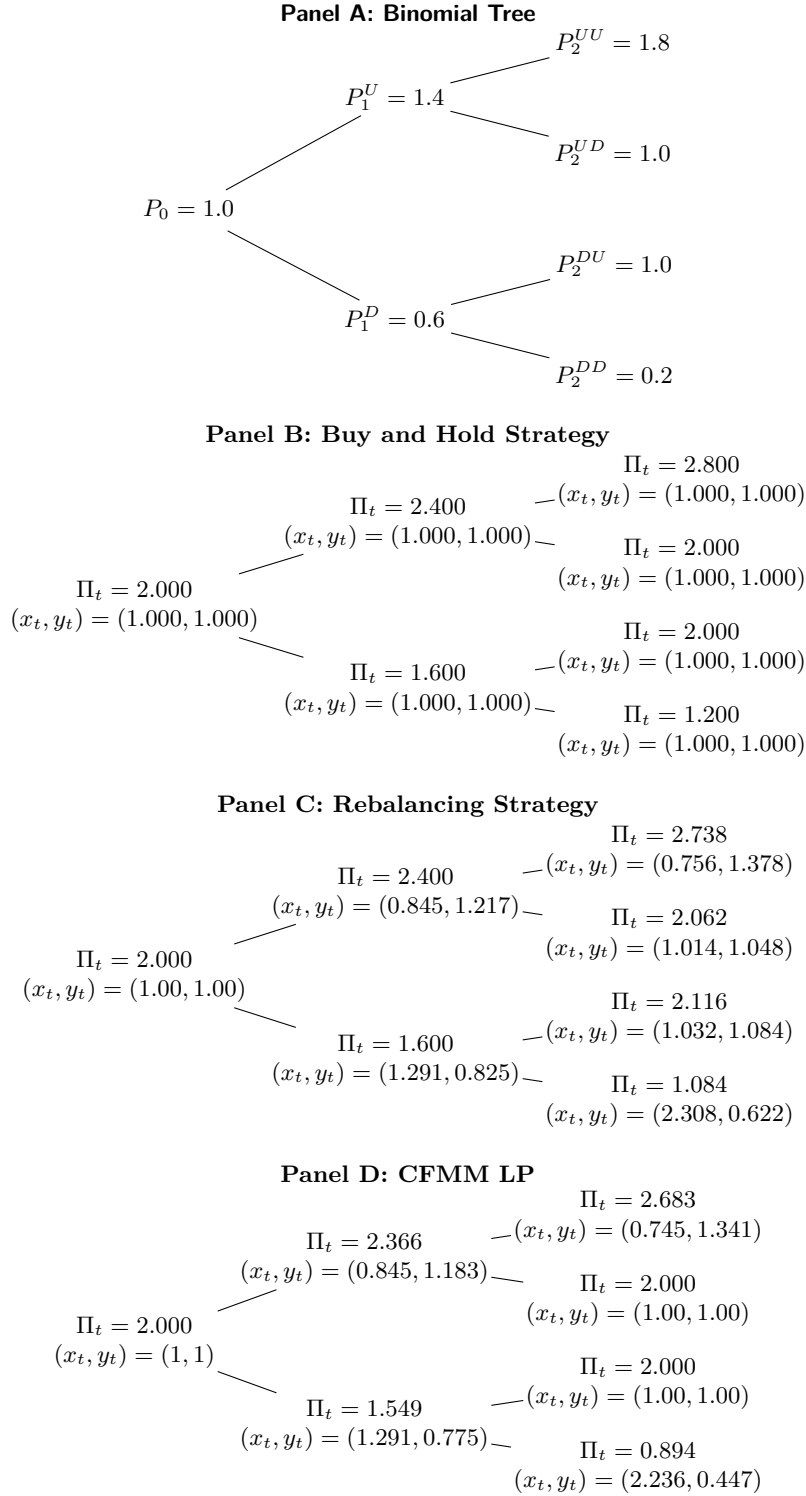


Figure 3: The performance of buy-and-hold, a constant-product CFMM, and the rebalancing strategy, on a two-step binomial tree. Panel A depicts the binomial tree. The performance of the buy-and-hold strategy is shown in Panel B; the rebalancing strategy is shown in panel C, and the constant product CFMM is shown in panel D. In each panel, x_t, y_t are the holdings of the strategy, and Π_t is the pool value, $y_t + P_t x_t$.

Because it is self-financing, the only way that the benchmark strategy can make or lose money is through its exposure to the risky asset. Hence, the *value* of the benchmark strategy, $\bar{R}_t(\bar{x})$, depends on the holding strategy \bar{x}_t and movements in prices, according to (4):

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

We define the AMM's loss versus this benchmark as $\text{LVB}_t(\bar{x}) \triangleq \bar{R}_t(\bar{x}) - V_t$.

Corollary 1. *For all $t \geq 0$,*

$$\text{LVB}_t(\bar{x}) \triangleq \bar{R}_t(\bar{x}) - V_t = \text{LVR}_t + \underbrace{\int_0^t [\bar{x}_s - x^*(P_s)] dP_s}_{\triangleq \Delta(\bar{x})_t}, \quad (14)$$

and at the same time,

$$\mathbb{E}^{\mathbb{Q}}[\text{LVB}_t(\bar{x})] = \mathbb{E}^{\mathbb{Q}}[\text{LVR}_t]. \quad (15)$$

The loss process has quadratic variation¹⁹

$$[\text{LVB}(\bar{x})]_t = [\Delta(\bar{x})]_t = \int_0^t [\bar{x}_s - x^*(P_s)]^2 \sigma_s^2 P_s^2 ds \geq [\text{LVR}]_t = 0. \quad (16)$$

Therefore, among all benchmark strategies, the rebalancing strategy uniquely (up to a set of times of measure zero) defines a loss process with minimal (zero) quadratic variation.

Proof. The first part is an immediate corollary of Lemma 2 and (5). The second part is immediate by the definition of the risk-neutral measure. The third part follows from the Itô isometry. ■

Corollary 1 characterizes the impact of the choice of benchmark self-financing trading strategy. It establishes that the residual loss $\text{LVB}(\bar{x})$ to an arbitrary benchmark \bar{x} that trades as market prices is given by LVR plus an additional term $\Delta(\bar{x})$. By examining the integrand in (16), it is clear that, on a path-by-path basis, the residual loss process $\text{LVB}(\bar{x})$ must have positive quadratic variation, unless $\bar{x}_s = x^*(P_s)$ for almost all $s \in [0, t]$. This condition characterizes the rebalancing strategy.²⁰ In other words, the loss process will have positive quadratic variation along any given

¹⁹Note that (14) and (16) hold almost surely over all sample paths under the risk-neutral measure. Hence, they also hold almost surely over all sample paths under the physical measure.

²⁰The choice of benchmark does not affect risk-neutral *expected* payoffs. All choices of \bar{x} lead to benchmarks with

sample path, unless it is essentially equal to the rebalancing strategy.

Since the square root of quadratic variation $X \mapsto \sqrt{[X]_t}$ is a seminorm over the space of semi-martingales X , we can view Corollary 1 as a “projection” result: consider the following optimization problem over a single sample path

$$\begin{aligned} & \underset{\bar{x}}{\text{minimize}} && \sqrt{[\bar{R}(\bar{x})_t - V_t]_t} \\ & \text{subject to} && \bar{x} \text{ is an admissible, self-financing trading strategy.} \end{aligned} \tag{17}$$

Then, the rebalancing strategy is essentially the unique (up to a set of times of measure zero) optimal solution to (17). In this way, the rebalancing strategy can be viewed as a projection of the AMM’s pool value process onto the set of self-financing trading strategies that trade at market prices. Moreover, this choice of benchmark results in a loss process (LVR) with zero quadratic variation that is predictable.²¹

This derivation thus gives a complementary intuition to the result in Theorem 1. In equity pricing models, asset alphas are isolated by *projecting* individual asset returns onto market or factor returns, picking *beta* values which minimize the expected squared difference between asset and factor returns, and then subtracting the fitted component from returns. This removes variation due to market or factor risk, leaving a residual that reflects asset-specific performance.

Corollary 1 shows that our approach is analogous: the “beta-like” rebalancing strategy is the quadratic-variation-minimizing projection of the AMM pool value process onto changes in P_s . Subtracting the rebalancing strategy’s returns yields an “alpha-like” residual free of linear price exposure to the underlying asset.²²

The goal of CAPM and FF3 models is not to argue that stocks *should* be held in market-neutral strategies, and the goal of Black-Scholes is not to argue that options *should* be delta-hedged. In

equal risk-neutral *expected* returns, cf. (15). Equivalently, any physical expected return differences are purely driven by asset risk premia, and disappear after delta-hedging. (16) is about risk, rather than risk-neutral expected returns: whenever \bar{x}_s and $x^*(P_s)$ differ, *realized* returns of the two strategies will differ on any sample path, purely due to differential exposure to risky asset price movements.

²¹Examining (14), it is clear that $\text{LVB}(\bar{x})$ is a super-martingale and that, as per the Doob-Meyer decomposition, it must decompose uniquely into an increasing predictable process (the “compensator”) and a martingale. LVR is the compensator for $\text{LVB}(\bar{x})$ independent of the choice of benchmark, while $\Delta(\bar{x})$ defines the martingale component and depends on the choice of benchmark.

²²We can also draw an analogy to options-pricing: the rebalancing strategy corresponds to the *delta* component of LP returns. In standard option pricing settings, delta is derived via backward induction and depends on unobserved volatility. Computing deltas for AMMs is much simpler: when the price is P_s , the AMM’s delta is simply its asset holdings, $x^*(P_s)$.

both cases, the decompositions serve primarily conceptual purposes. Of course, the decompositions naturally suggest trading strategies: these model-based insights motivate the factor risk management strategies of long-short hedge funds, and the delta-hedging strategies of option market makers. Following this tradition, this paper focuses on the conceptual P&L decomposition within a frictionless model, leaving a full accounting of the frictional costs and details of implementation to practitioners.

6. Examples

Example 2 (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 (1) 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,

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

and

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

The weighted geometric mean market maker generalizes the constant product market maker. For these market makers, the instantaneous LVR normalized per dollar of pool reserves is constant, i.e.,

$$\frac{\ell(\sigma, P)}{V(P)} = \frac{\sigma^2}{2} \theta(1-\theta). \tag{18}$$

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.2. Finally, observe that LVR is maximized when $\theta = 1/2$, and goes to zero as $\theta \rightarrow \{0, 1\}$.²³

Example 3 (Constant Product Market Maker / Uniswap v2). *Taking $\theta = 1/2$ in Example 2, we have*

²³See also Proposition 1 of Evans [2020], evaluating a weighted geometric mean market maker over a finite horizon using risk-neutral pricing.

that

$$V(P) = 2L\sqrt{P}, \quad \ell(\sigma, P) = \frac{L\sigma^2}{4}\sqrt{P}, \quad \frac{\ell(\sigma, P)}{V(P)} = \frac{\sigma^2}{8}. \quad (19)$$

This example shows that the constant product market maker admits particularly simple expressions for LVR: $\ell(\sigma, P)/V(P)$, the loss per unit time as a fraction of mark-to-market pool value, is simply 1/8 times the instantaneous variance. This formula is straightforward to apply empirically: for example, if the ETH-USDC volatility is $\sigma = 5\%$ (daily), this formula implies that the ETH-USDC LP pool loses approximately $\sigma^2/8 = 3.125$ (bp) in pool value to LVR daily.

Example 4 (Uniswap v3 Range Order). *Given liquidity parameter $L \geq 0$ and prices in the given range $[P_a, P_b]$, consider the bonding function of the “range order” Adams et al. [2021],*

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

Solving the pool value optimization (1),

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

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

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

so that

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

Observe that the instantaneous LVR is the same in Example 3. However, the pool value $V(P)$ is lower. Indeed $V(P) \rightarrow 0$ if $P_a \uparrow P$ and $P_b \downarrow P$, so

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

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 5 (Uniswap v3 Pool). A “concentrated liquidity” pool such as Uniswap v3 [Adams et al., 2021] aggregates the liquidity across a set of range orders $i = 1, \dots, N$, where order i is a range order with liquidity L_i over range²⁴ $[P_a^{(i)}, P_b^{(i)})$. Following Example 4, for each order i , we have that

$$x_i^*(P) = \begin{cases} L_i \left(\frac{1}{\sqrt{P_a^{(i)}}} - \frac{1}{\sqrt{P_b^{(i)}}} \right), & \text{if } P < P_a^{(i)}, \\ L_i \left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{P_b^{(i)}}} \right), & \text{if } P \in [P_a^{(i)}, P_b^{(i)}), \\ 0, & \text{if } P \geq P_b^{(i)}, \end{cases} \quad (20)$$

$$y_i^*(P) = \begin{cases} 0, & \text{if } P < P_a^{(i)}, \\ L_i \left(\sqrt{P} - \sqrt{P_a^{(i)}} \right), & \text{if } P \in [P_a^{(i)}, P_b^{(i)}), \\ L_i \left(\sqrt{P_b^{(i)}} - \sqrt{P_a^{(i)}} \right), & \text{if } P \geq P_b^{(i)}, \end{cases} \quad (21)$$

$$V_i(P) = \begin{cases} L_i \left(\frac{1}{\sqrt{P_a^{(i)}}} - \frac{1}{\sqrt{P_b^{(i)}}} \right) P, & \text{if } P < P_a^{(i)}, \\ L_i \left(2\sqrt{P} - P/\sqrt{P_b^{(i)}} - \sqrt{P_a^{(i)}} \right), & \text{if } P \in [P_a^{(i)}, P_b^{(i)}), \\ L_i \left(\sqrt{P_b^{(i)}} - \sqrt{P_a^{(i)}} \right), & \text{if } P \geq P_b^{(i)}, \end{cases} \quad (22)$$

$$V_i''(P) = x_i^{*'}(P) = \begin{cases} -\frac{1}{2}L_i P^{-3/2}, & \text{if } P \in [P_a^{(i)}, P_b^{(i)}), \\ 0, & \text{otherwise.} \end{cases} \quad (23)$$

The aggregate reserves and pool value are given by [Milionis et al., 2023b]

$$x^*(P) = \sum_{i=1}^N x_i^*(P), \quad y^*(P) = \sum_{i=1}^N y_i^*(P), \quad V(P) = \sum_{i=1}^N V_i(P).$$

When we compute (instantaneous) LVR, we need marginal liquidity $V''(P) = x^{*'}(P)$. Observe from (23) that $V_i''(P)$ is zero when order i is not in-range, i.e., when $P \notin [P_a^{(i)}, P_b^{(i)})$. Hence, only

²⁴In addition, the endpoints of each price range may be constrained to predefined endpoints or “ticks”.

in-range orders contribute to LVR, i.e.,

$$\ell(\sigma, P) = -\frac{\sigma^2 P^2}{2} V''(P) = -\frac{\sigma^2 P^2}{2} \sum_{i=1}^N V_i''(P) = \frac{\sigma^2}{4} \sqrt{P} \times \underbrace{\sum_{i=1}^N L_i \mathbf{1}_{\{P \in [P_a^{(i)}, P_b^{(i)}]\}}}_{\triangleq \bar{L}(P)}.$$

The quantity $\bar{L}(P)$ is the *aggregate in-range liquidity*, given the current price P . From the perspective of the set of range orders which comprises the pool, this is a sufficient statistic for computing LVR.

Example 6 (Linear Market Maker / Limit Order). *For $K > 0$, consider the linear bonding function $f(x, y) \triangleq Kx + y$. Solving the pool value optimization (1),*

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

Hence, this pool can be viewed as similar to a resting limit order²⁵ 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 1 Part 2: the first derivative is discontinuous at the limit price $P = K$. Thus, the characterization of Lemma 2 does not apply.²⁶

²⁵While 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. This merely refers to a superficial change of the *default strategy* after a trade execution.

²⁶Note that the pool value function remains concave and thus the pool value process is a super-martingale. Hence, from the Doob-Meyer decomposition, a non-negative monotonic running cost process (analogous to LVR) exists. However, this process is not described by (7)–(8). 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].




Figure 4: The daily average pool value of the Uniswap v2 WETH-USDC pool.

7. Empirical Analysis

We now bring our model to data, showing how to separate the “alpha-like” and “beta-like” components of LP P&L in practice. This allows us to answer a number of questions. First, how large are the alpha-like and beta-like components, respectively, in terms of their contribution to the variance of returns? Second, how sensitive are the results to the frequency at which we rebalance the AMM’s holdings? Third, as we rebalance at increasingly high frequencies, does the result converge to the “alpha-like” microstructural component of returns — fees minus adverse selection costs — as our model predicts?

7.1. Methodology

We bring the model to data using the WETH-USDC trading pair²⁷ on Uniswap v2 for the period from August 1, 2021 to July 31, 2022. The value of the reserves of this pool over the period are shown in Figure 4. Details of the data sources we use, and how we measure various quantities, are described in Appendix C.

²⁷“WETH”, or “wrapped ETH” is a variation of ETH that is compliant with the ERC-20 token standard. Conversion between the two is a straightforward operation with a minimal fixed cost. For our purposes, we will view ETH and WETH as equivalent.

Repeating (13), we have

$$\overbrace{\text{LP P\&L}_t - \int_0^t x^*(P_s) dP_s}^{\text{hedged P\&L}} = \underbrace{\text{FEE}_t - \text{LVR}_t}_{\text{fees minus LVR}}. \quad (24)$$

market risk

The left side of (24) can be thought of as the P&L from a delta-hedged LP position: the P&L of the LP position, minus that of the rebalancing strategy. This quantity can be estimated empirically under very weak assumptions:

- We can approximate the integral in (24) using *discrete rebalancing strategies*; for a fixed interval length Δ , we define:

$$\text{RB P\&L}_t \triangleq \sum_{s=0, \Delta, 2\Delta, \dots}^t \Delta \text{RB P\&L}_t, \quad (25)$$


$$\Delta \text{RB P\&L}_t \triangleq x_t^{RB} (P_{t+\Delta t} - P_t). \quad (26)$$

Here, x_t^{RB} is the quantity of the risky asset in the pool reserves at time t . Mathematically, (25) is a Riemann sum approximation for rebalancing strategy profits. Economically, (25) can be thought of as the payoffs of a strategy which delta-hedges at market prices to match the AMM’s holdings at time intervals Δ . This strategy holds x_t^{RB} on the interval $[t, t + \Delta t)$, and thus makes dollar profits equal to the holdings times the change in prices $P_{t+\Delta t} - P_t$ on this interval.

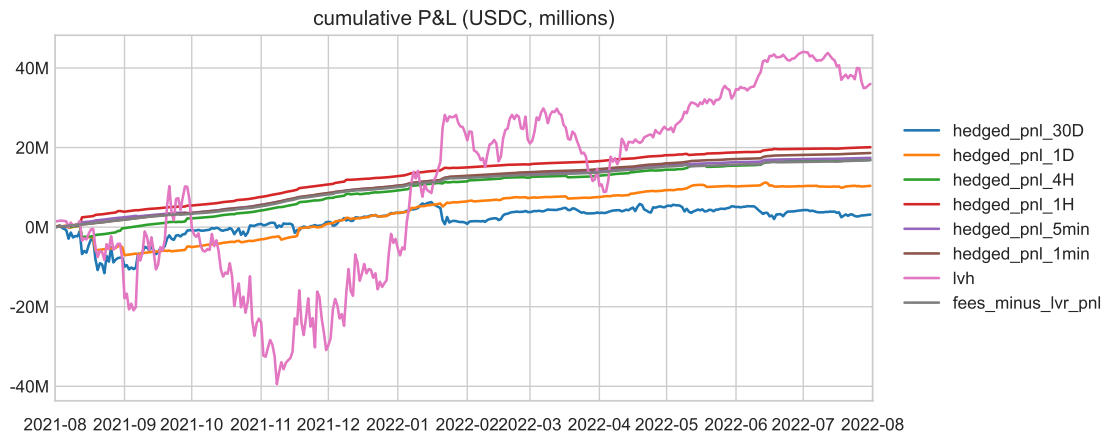
- We measure LP P&L_t, as our model suggests, essentially as $P_t x_t + y_t$, using the observed pool holdings x_t, y_t and the CEX price P_t . In practice, pool reserves adjust over time through “mints” and “burns” of LP positions; we value these mints and burns periodically to CEX prices. Details of these calculations are provided in Appendix C.

7.2. Empirical Results

Figure 5 shows the results of our analysis. In Figure 5a, the `pool_pnl` series shows unhedged LP pool P&L. The pool’s value is fairly volatile: daily gains or losses of over \$10 million USD (on a pool reserves whose value varies between \$100–\$300 million USD, cf. Figure 4) are common.



(a) Including the aggregate pool P&L series.



(b) Excluding the aggregate pool P&L series.

Figure 5: Cumulative pool P&L, delta-hedged P&L, and predicted P&L from our expressions for LVR, for the Uniswap v2 WETH-USDC trading pair. In the first panel, the `pool_pnl` series shows the raw P&L of the LP position. In both panels, the various `hedged_pnl` series show delta-hedged LP P&L, that is, the P&L from a long position in the pool, and a short position in the rebalancing strategy, rebalanced at various frequencies (daily, every four hours, and a short hour, every five minutes). The `fees_minus_lvr_pnl` series shows fees minus LVR, this is the delta-hedged P&L predicted by our theory. The data source for prices is Binance, and LP P&L is calculated using data on Uniswap trades, mints, and burns from the Ethereum blockchain. Details of how we calculate these quantities are in Appendix C.

series	rebalancing freq	avg return (%, annual)	stdev return (%, annual)	Sharpe ratio (annual)
pool_pnl	unhedged	-6.20	42.08	-0.15
lvh	once	17.46	27.08	0.64
hedged_pnl_30D	30 day	1.53	7.27	0.21
hedged_pnl_1D	1 day	5.04	2.87	1.75
hedged_pnl_4H	4 hour	8.21	1.50	5.47
hedged_pnl_1H	1 hour	9.75	0.90	10.81
hedged_pnl_5min	5 min	8.45	0.47	18.16
hedged_pnl_1min	1 min	9.04	0.39	23.33
fees_minus_lvr_pnl	continuous limit	8.16	0.48	17.03

Table 1: Overall return statistics for the cumulative P&L series of Figure 5. This table shows the standard deviation of hedged P&L, where hedging is done at interval lengths ranging from every minute, to every month. The `lvh` series considers hedging a single time, using ETH holdings at the start of our sample period, on August 1 2021, corresponding to “impermanent loss” using a single starting point. The `pool_pnl` series corresponds to raw LP P&L.

However, the vast majority of this risk is simply market risk: the pool holds a long position in ETH, whose price is very volatile over the period we study (cf. Figure 6). The `hedged_pnl` series show the hedged P&L from (24), rebalancing over various frequencies. Visually, these series are so much smaller than the aggregate `pool_pnl` series that their fluctuations are almost invisible in Figure 5a.

Figure 5b zooms in, focusing on the hedged P&L. All of these series are substantially less volatile than the unhedged P&L. This is quantified in the return statistics in Table 1. Quantitatively, the standard deviation of returns of 1-minute hedged P&L is less than 1% of the standard deviation of returns of the unhedged P&L; in variance terms, 0.009% of the variation in returns in unhedged P&L comes from fees and adverse selection, and the remaining 99.991% simply reflects ETH price movements.

As the rebalancing intervals, hedged LP P&L becomes increasingly more volatile, reflecting market risk accrued within intervals. This is consistent with Corollary 1, stating that the rebalancing strategy corresponds to the variance-minimizing projection of LP P&L on market risk: trading strategies which are increasingly coarse approximations to the rebalancing strategy thus leave increasing amounts of variance within the hedged LP time series.

In the `lvh` series of Figure 5 and Table 1, we consider a benchmark which holds ETH holdings at the start of the sample and *never* rebalances, corresponding to “impermanent loss” (cf. Section 8).

The volatility of “impermanent loss” measure is of the same order of magnitude at the unhedged pool. This illustrates why “impermanent loss” fails to account for market risk exposures: it perfectly hedges market risk at the start of the time interval, but accrues market risk as prices and holdings move over time; the result is a series that is almost as polluted by market risk as raw LP P&L.

7.3. Testing the Theory

Our model implies that, as we approximate the profits of the rebalancing strategy increasingly well, the hedged P&L should converge not to zero, but to the “alpha-like”, microstructural component of LP returns on the right side of (24): accrued fees minus accrued adverse-selection costs. We can test this prediction by directly measuring this series, and comparing it to high-frequency hedged LP returns. We directly observe FEE_t , fees paid into the LP pool over any given time period.²⁸ For LVR_t , since Uniswap v2 is a constant-product CFMM, adverse selection losses have the particularly simple expression in (19) of Example 3,

$$LVR_t = \int_0^t \frac{\sigma_s^2}{8} \times V(P_s) ds. \quad (27)$$

We estimate daily realized volatility using Binance prices, see Figure 6. We can then compute LVR simply by plugging in realized daily volatility and pool value to a version of equation (27) that is discretized at a daily frequency.

Note that, empirically, we measure the total fees paid by all kinds of traders. This differs slightly from our model, where we assume arbitrage traders pay no fees. Practically, since fees are simply an increasing process which potentially compensates for LVR, whether fees arise from noise trade or arbitrage trade does not substantially impact the returns on LP positions. If we assumed arbitrage traders paid trading fees in the model, this would decrease the amount of arbitrage: instead of prices on the CFMM moving immediately to match CEX prices at all times, prices would have to move more than fees in order for arbitrage trade to have nonnegative payoffs. The analysis of arbitrage profits in the presence of fees is the subject of follow-on work [Milionis et al., 2023a].

We plot fees minus LVR as the `fees_minus_lvr_pnl` series in Figure 5. Fees minus LVR in our model closely track the pattern of hedged LP P&L. We plot the difference between these series in

²⁸Uniswap v2 also allows for assets to be borrowed for “flash loans” for a fee; however, in v2, these loans are accounted for as swaps, thus FEE_t contains revenue from flash loans. We discuss this briefly in Appendix C.1.




Figure 6: Daily realized volatility for the ETH-USDC pair, computed from Binance minutely closing prices, sampled at 60 minute intervals.

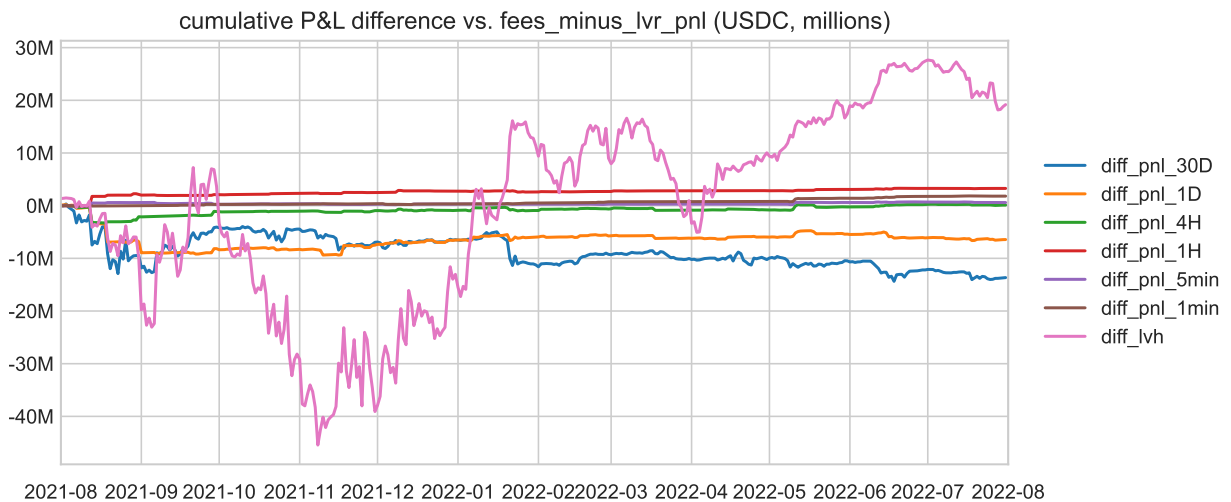


Figure 7: Cumulative differences between the delta-hedged P&L (various `hedged_pnl` series of Figure 5) and predicted P&L from our model (`fees_minus_lvr_pnl` series of Figure 5).

Figure 7. Observe that the gap is smaller for more frequent rebalancing frequencies. In this way, fees minus LVR is an approximate continuous time limit of hedged P&L, just as our model predicts.

This increases our confidence that subtracting the “beta-like” rebalancing strategy from LP P&L indeed leaves an “alpha-like” residual which reflects the contribution of microstructural forces to AMM LP returns. In practice, we recommend that empirical researchers primarily construct hedged LP P&L simply by subtracting the rebalancing strategy’s profits from observed LP P&L: directly subtracting the payoffs of the hedging strategy is more straightforward to interpret, and imposes less theoretical structure on the underlying data.

7.4. Implications for Empirical Analysis of AMM Returns

Figure 5 and Table 1 show that LP P&L consists almost overwhelmingly of the “beta-like” market risk component. This has important implications for empirical analyses of LP returns: if market risk is not explicitly removed from LP returns, any results will predominantly reflect market risk rather than microstructural effects.

Suppose a researcher wishes to study the determinants of LP P&L, and attempts to do so by regressing unhedged LP returns (or fixed-start-point “impermanent loss”, which behaves similarly) on some x -variable. The best-case scenario for the researcher is that market risk is pure noise, uncorrelated with x -variables. OLS coefficient standard errors are proportional to the standard deviation of the regression residual; thus, even in this best case, using unhedged LP returns as the y -variable leads to standard errors roughly 108 times larger than if hedged returns were used. Since standard errors scale with the inverse square root of sample size, achieving equal coefficient standard errors using unhedged LP returns as a y -variable would require around $108^2 = 11,664$ times more data.

In reality, a more severe concern is that market risk may be correlated with any x -variables used. The large relative variance of market risk implies that even a weak correlation between the x -variable and market risk can induce substantial omitted-variable bias, easily swamping any genuine relationship between x and hedged LP returns.

We therefore recommend that empirical researchers analyzing AMM LP returns implement our results by using hedged LP returns as regression inputs. Our procedure has very low data requirements: to construct profits of the rebalancing strategy, (26), the researcher only needs x_t^{RB} ,

the time series of risky asset holdings of the CFMM LP, and P_t , the CEX price of the risky asset. Our results also appear robust to the choice of rebalancing frequency. Failing to hedge can have severe consequences: in empirical designs using unhedged LP returns, any microstructural “signal” — in our model, fees minus adverse selection costs — is likely to be overwhelmed by market noise.

A trader attempting to delta-hedge LP returns in practice would of course face various implementation costs, such as spreads and fees for rebalancing trades, or margin costs for maintaining a short position in the risky asset. Our goal here is not to precisely estimate the *practical* profits from hedged LP-ing; rather, it is mainly to *conceptually* decompose profits into market risk and residual components, analogous to alpha-beta decompositions in asset-pricing factor models. Following practice in this literature, we thus calculate RB P&L_{*t*} assuming frictionless rebalancing.²⁹

7.5. Markouts

Industry participants frequently use “markouts” as an approach to evaluate CFMM LP P&L [see, e.g., CrocSwap, 2022]. Markouts, known as “realized spread” in the microstructure literature [Huang and Stoll, 1996], essentially attribute profits to trades by comparing the price of each trade to a future price, usually at some fixed time offset from the trade (e.g., 10 minutes) either from a CEX, or the CFMM itself. However, delta-hedging and the markout approach are in a sense closely related: the P&L of delta-hedged LPing, when rebalanced at discrete periods, turns out to be exactly equivalent to markout profits, marked to CEX prices at the end of discrete periods of the same frequency. In other words, the main difference between delta-hedging and markout analysis is whether the marking price is obtained in a fixed offset of time in the future after a trade, or based on the ending price of the interval the trade was contained in. These two are likely to be very close, especially when the markout offset or delta-hedging interval are short. Thus, delta-hedged LPing can be thought of as a microfoundation for markout-style analysis.

²⁹While frictions of course play a role in practice, large trading firms have a variety of ways of optimizing these frictional costs; for example, they may rebalance at lower frequencies, optimize rebalancing execution, or “cross-hedge” by netting delta exposures across different parts of the portfolio, to limit trade and shorting costs. These practical factors imply that any attempt we make to quantify the frictional costs of rebalancing is unlikely to be generally accurate across all applied settings.

8. “Impermanent Loss”

It is common in industry practice to evaluate LP P&L against simply holding the LP’s asset mix at some fixed start point; this is the loss metric we call “loss versus holding” (LVH) in earlier sections, and is sometimes called “impermanent loss” in industry practice. In the notation of Section 5.2, defining the benchmark strategy $x_t^{\text{HODL}} \triangleq x^*(P_0)$, then $\text{LVH}_t \triangleq \text{LVB}_t(x_t^{\text{HODL}})$ is defined as the loss process relative to this benchmark. Like LVR, LVH has the property that it is non-negative, and can thus also be viewed as a “cost”:

Proposition 1. *For all $t \geq 0$, $\text{LVH}_t \geq 0$.*

Proof.

$$\text{LVH}_t = \bar{R}_t(x^{\text{HODL}}) - V_t = (P_t x^*(P_0) + y^*(P_0)) - (P_t x^*(P_t) + y^*(P_t)),$$

and, by the optimality of Equation (1) for $P = P_t$, this quantity is always non-negative. ■

LVH seems like a simpler metric than LVR, since it benchmarks profits to a single start point, and does not require calculating the profits of a continuously rebalanced benchmark portfolio. However, this seeming simplicity leads to severe problems with using LVH in practice.

8.1. Problems with Loss-Versus-Holding

First, LVH fails a basic “cumulation” property that is satisfied by most return-like metrics used in finance. Let $\text{LVH}_{t,\tilde{t}}$ denote loss-versus-holding from t to \tilde{t} , evaluated against buying and holding the time- t portfolio. We have, in general:

$$\text{LVH}_{t_1,t_2} + \text{LVH}_{t_2,t_3} \neq \text{LVH}_{t_1,t_3} \tag{28}$$

This non-additivity arises because LVH_{t_2,t_3} is benchmarked against time- t_2 holdings, instead of time- t_1 holdings. A simple example appears in the “reconvergence” paths P^{UD} and P^{DU} on Figure 3. The AMM makes trades and incurs positive LVH on each individual time step: $\text{LVH}_{0,1} > 0$ and $\text{LVH}_{1,2} > 0$. However, the $t = 2$ portfolio is equal to the $t = 0$ portfolio so we have $\text{LVH}_{0,2} = 0$! The econometrician faces a paradox: the AMM appears to suffer positive losses at each individual time step, but no loss over the interval as a whole. Relatedly, LVH is not monotonic: adverse selection

costs may appear to reverse as prices reconverge, even though the AMM has incurred genuine adverse selection costs at every step. By contrast, LVR is monotonically non-decreasing, reflecting the economically natural property that adverse selection costs accumulate irreversibly over time.

Second, the sign of the change in LVH can vary, depending on the start point chosen. Consider an AMM with a fixed level of liquidity; suppose the AMM holds $x = 100$ of the risky asset at t_1 , and $x = 200$ at t_2 . Consider a future time period $[t_3, t_4]$ during which holdings increase from 150 to 160. Then:

$$\text{LVH}_{t_1, t_4} - \text{LVH}_{t_1, t_3} < 0$$

$$\text{LVH}_{t_2, t_4} - \text{LVH}_{t_2, t_3} > 0$$

That is, LVH seems to *increase* relative to the t_1 benchmark, but *decrease* relative to the t_2 benchmark! Both interpretations are internally consistent; however, the econometrician cannot unambiguously conclude whether, in the time period $[t_3, t_4]$, the AMM made or lost money. The magnitude, and even the sign, of the incremental profit depends on an arbitrarily chosen reference point.

Third, LVH is path-independent: for any CFMM, $\text{LVH}_{0, T}$ depends only on the initial and final prices P_0 and P_T , not on the path of prices between them. This is economically problematic, since adverse selection costs should naturally depend on the price path. An LP should lose more to informed trading during a period of high volatility than during a period of low volatility, even if prices start and end at the same levels. By contrast, LVR is explicitly path-dependent: for a constant-product AMM, (27) shows that LVR accumulates faster during high-volatility periods, consistent with the intuition that adverse selection is more costly when prices are more volatile.

Fourth, LVH conflates the “alpha-like” microstructural component of LP returns — adverse selection costs specific to the AMM’s market-making function — with the “beta-like” market risk arising from the AMM’s asset holdings. Because LVH benchmarks against fixed holdings, it accrues market risk as the AMM’s holdings diverge from the benchmark over time (cf. Corollary 1). As our empirical analysis in Section 7 shows, this beta-like component is orders of magnitude larger than the alpha-like component in variance terms, and can easily overwhelm any microstructural signal. LVR, by contrast, cleanly separates these two forces: the rebalancing strategy absorbs all beta-like market risk, isolating the alpha-like microstructural component of LP returns.

A natural way to fix these problems is to re-benchmark performance in each short time interval against holdings at the start of that interval. Fixing some period length Δ , we could define a loss metric L , as:

$$L_{t_1, t_2} \triangleq \text{LVH}_{t_1, t_1 + \Delta} + \text{LVH}_{t_1 + \Delta, t_1 + 2\Delta} + \dots + \text{LVH}_{t_2 - \Delta, t_2} \quad (29)$$

That is, over each length- Δ interval, AMM payoffs are compared to holding the portfolio fixed at the beginning of that interval. This construction guarantees additivity:

$$L_{t_1, t_2} + L_{t_2, t_3} = L_{t_1, t_3} \quad (30)$$

But (29) is exactly the definition of LVR! Thus, LVR can be thought of as simply a *time-discretized, path-consistent* version of LVH: it simply updates the benchmark portfolio at every interval as prices and holdings evolve.

8.2. “Impermanent Loss” in the Empirical AMM Literature

In the one-period models used in the majority of the theoretical literature on AMMs, “impermanent loss” has a simple and unambiguous interpretation. However, empiricists who wish to measure “impermanent loss” in practice must make an important measurement choice: whether to treat time as one long period, or several shorter periods.

In the absence of theoretical guidance, both choices have been made in the empirical literature. Augustin et al. [2024] appears to simply fix a start point and calculate impermanent loss based on holdings at that point in time.³⁰ Other papers, such as Fang [2024] and Aquilina et al. [2024], appear to calculate impermanent loss on a number of shorter time intervals, and add the resulting losses up; specifically, we believe Aquilina et al. [2024] calculates impermanent loss using daily-updated holdings benchmarks,³¹ and that the main specifications in Fang [2024] correspond to weekly-rebalanced impermanent loss.³² Some papers, such as Heimbach et al. [2022], do not specify

³⁰Augustin et al. [2024] write: “The aggregate h-period return to yield farming is thus composed of four components associated with capital gains, impermanent losses, trading fees, and realized farm yields...” The following expression (8) then contains a difference of return metrics $R_{t, t+h}^A$ and $R_{t, t+h}^B$, leading us to believe that the holdings are fixed at a start point t , which appears to be set to September 23, 2020 for all pools.

³¹On page 27, Aquilina et al. write: “We assess Fee Yield and Impermanent Loss on a daily basis, i.e. an LP can decide whether or not to provide liquidity at the end of each day”.

³²Fang [2024] writes: “For the pool-level analyses, I aggregate transactions for each pool at daily, weekly, or monthly frequency. I focus on the weekly frequency because this appears to be the modal horizon for liquidity adjustment.”

whether the holdings benchmark is updated periodically, or what the start point is.

All these papers refer to the object they measure as “impermanent loss”, implicitly assuming that the underlying economic object measured is insensitive to these measurement decisions. This is incorrect. “Impermanent loss” measured using periodically updated holdings, versus fixed holdings, are substantively different metrics. The periodically-updated measure converges to fees minus adverse selection costs, for reasonably high update frequencies. The fixed-start-point measure, in our analysis in Section 7, has over a hundred times greater variance than the periodically-updated measure, and mainly reflects movements in ETH prices rather than any microstructural forces. The periodically-updated measure does not add up to the fixed-start measure, and the latter can easily be zero while the former is positive.

In our view, there is a missing link between the theoretical and empirical literature on AMMs: *market risks* are not salient in the one-period models used in most of the theoretical literature. As we showed in Figure 3, “impermanent loss” (loss versus holdings) and LVR are identical over a single time step; the difference only emerges at the second step, where the AMM’s holdings deviate from the fixed-holding benchmark, and thus LVH begins to accrue market risk. This point has not been explicitly made in the theoretical AMM literature, allowing market risk to creep unnoticed into the metrics used in the empirical AMM literature. Our paper aims to fill this missing link: a simple *dynamic* model of AMMs allows us to clearly separate the distinct contributions of microstructural forces and market risk to AMM P&L.

Our paper gives clear guidance on how the empirical literature should proceed. If the goal of analysis is to measure *microstructural effects* — the performance of the AMM in its market-making capacity — the market-risk component of returns should be treated as a nuisance parameter to be eliminated, akin to how factor models are used to remove known sources of predictability from returns in equity-pricing settings. This can be done simply by subtracting the profits of the rebalancing strategy from pool P&L, as we do in Section 7. This is approximately the approach taken in Fang [2024] and Aquilina et al. [2024], though — aggressively extrapolating from our

We note that (28) implies that there is no guarantee that daily-frequency losses thus calculated add up to weekly or monthly frequency results. Fang also writes: “Instead of loss-vs-holding, Milionis et al. (2022) show that LP returns can be characterized by loss-vs-rebalancing (LVR). My main results are qualitatively the same when using LVR.” However, results using LVR are not reported in the paper. This statement is plausible, since according to our definitions, the weekly-updated-holdings “impermanent loss” measure used in Fang [2024] calls is exactly equivalent to what we call “loss-versus-rebalancing”, rebalanced at a weekly frequency.

empirical analysis — both papers could benefit from slightly more frequent rebalancing, as we find that market risk still dominates microstructural forces at daily and weekly rebalancing frequencies. The loss measure obtained through this approach is mathematically and economically distinct from “impermanent loss” based on a fixed start point. We suggest, of course, calling this updated-holdings measure “loss versus rebalancing” instead of “impermanent loss”.

Of course, different measures serve different purposes, and we cannot say in any formal sense that LVR is the universally “best” metric. However, our model does suggest that fixed-start-point “impermanent loss” is simply equal to LVR, plus the market risk exposures of a well-defined trading strategy, and the market risk component may be much larger than microstructural effects in variance terms. Thus, we believe papers which use fixed-start “impermanent loss” should make an argument for why it is appropriate to include this market risk, for the particular question they consider. As a minimum standard, papers should carefully specify the start-point and updating-frequency decisions they make in their implementation of metrics labelled “impermanent loss”.

9. Option Pricing

We have shown that CFMM LPs behave like a bet on volatility, in the sense that LVR is large when volatility is high. In this section, we briefly discuss the relationship of CFMM LPs to three classical and inter-related ways that volatility can be traded [Carr and Madan, 2001]: static (European) options positions, dynamic trading strategies, and variance swaps. In Appendix B.1, we also demonstrate the equivalences between CFMM LP positions and these concepts in a simple two-step binomial tree model.

9.1. Static Options Positions

Our results are related to Clark [2020], Fukasawa et al. [2022], and Deng et al. [2023], who show that, over any finite time horizon, an AMM LP position’s payoff, ignoring fees, can be replicated by shorting a bundle of European options. Technically, this follows from the facts 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 $V_T = V(P_T)$ corresponding to the final “payoff function” $V(\cdot)$, regardless of the path that prices took to reach P_T . More intuitively, at any time

T , the CFMM simply offers a menu of quantities of the asset to buy or sell at any given price, identically to a portfolio of European options. Ignoring fees, the CFMM exactly breaks even if prices do not move, $P_T = P_0$, and loses money otherwise; hence, the CFMM LP position is essentially equivalent to *giving away* a bundle of European options. This intuition is consistent with the fact that the $V(\cdot)$ is a *concave* function (cf. Lemma 1).

Expected LVR until period T can be thought of as the value of the European options given away. This analogy gives another intuition for the comparative statics of expected LVR. European options are worth more when volatility is higher, so LVR is increasing in the volatility of the underlying asset. When the marginal liquidity of the AMM bonding curve is greater, the replicating portfolio of European options is larger: AMMs that trade more aggressively essentially give away more European options, also increasing LVR.

As previous papers have discussed, the European option replication result also implies that, over any finite time horizon, the exposure of the AMM LP position to underlying prices can be *totally* hedged, by taking a long position in the replicating bundle of European options. This trade — a long position in the AMM LP, plus a short position in the replicating bundle of European options — is essentially a *trading fee swap*, betting on whether accrued trading fees from time 0 to T are greater than European option premia of the replicating portfolio at time 0. The trader enters an LP position, and pays a premium for buying the replicating bundle of European options upfront. The AMM LP position then loses no money from price movements; the total position profits if the accrued trading fees until time T are greater than the European option premia paid upfront, and loses otherwise.

9.2. Dynamic Trading Strategies

Classic options theory implies that static option positions are equivalent to dynamically trading the underlying asset in a certain way. The static option position is a combination of short straddles and strangles, selling out-of-the-money calls and puts. This position is equivalent to a dynamic trading strategy which sells the asset when prices increase, and buys when prices decrease. This is exactly what the AMM LP position does: observe that, from Lemma 1 Part 3, $x^*(\cdot)$ 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 .

9.3. Variance Swaps

Finally, as discussed by Fukasawa et al. [2022], variance can be traded directly by trading swaps on realized variance. The VIX is such a contract, operating on a fixed finite time horizon. Applying Lemma 1 Part 3, the instantaneous LVR of (8) can be re-written as

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

Here, the first component, $(\sigma P)^2$, is the instantaneous variance or quadratic variation 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

$$\text{LVR}_t = \frac{1}{2} \int_0^t (\sigma_s P_s)^2 \times |x^*(P_s)| ds = \frac{1}{2} \int_0^t |x^*(P_s)| d[P]_s, \quad \forall t \geq 0.$$

This expression is the payoff of the floating leg of a continuously sampled generalized variance swap [see, e.g., Carr and Lee, 2009], specifically a price variance swap that is weighted by marginal liquidity.

10. Discussion and Implications

Besides their positive value for understanding and quantifying the losses from AMM LPing, our results also suggest ways that these losses could be reduced or eliminated. We showed that CFMM LPs lose money from price slippage: when CEX prices move, CFMM quotes become “stale” and

are vulnerable to sniping by rebalancing arbitrageurs. This slippage can in principle be eliminated: if an AMM had access to a high-frequency oracle for the CEX price P_t , the AMM could in principle quote prices arbitrarily close to P_t , up to the desired asset position $x^*(P_t)$.³³ Quoting prices this way would reduce arbitrageurs’ profits, allowing the AMM to achieve a payoff arbitrarily close to that of the rebalancing strategy. This design has a number of risks — it relies heavily on the accuracy of the oracle for P_t , and leaves open the potential for oracle manipulation — but in principle an oracle-based AMM could substantially reduce or eliminate LVR.

A related set of AMM designs revolve around selling the right to arbitrage the pool to certain special wallets, and redistributing profits to AMM LPs. This could be done in several ways. Suppose a particular crypto wallet address, which we call the “authorized participant” or AP, had the unique right for the first AMM trade in every block. This first execution right effectively conveys the right to arbitrage the pool [Josojo, 2022]. Alternatively, the AP could be provided the right to trade with the AMM paying zero fees [Jump Crypto, 2022]. The AP would then have a large advantage in arbitrage trading, since the AP could profitably trade against arbitrarily small price movements, whereas prices would have to move at least as much as the AMM’s percentage trading fees for non-AP wallets to profit from arbitrage trade. The AP would thus be able to capture a large fraction of fees from arbitrage trade. In either of these designs, an AMM protocol aiming to reduce LVR could thus run an AP wallet itself, doing CEX-DEX arbitrage, and redistributing arbitrage profits to LPs. Alternatively, AMM protocols could simply run periodic auctions — over longer periods, such as weeks or months — in which wallets can bid for the right to be authorized participants for some period of time. In principle, potential arbitrageurs should bid the *ex ante* expectation of arbitrage profits, which includes LVR; the protocol can then redistribute these profits to LPs. Adams et al. [2024] analyzes this design in detail.³⁴ Both these methods capture either LVR or its expectation, and redistribute the profits to pool LPs. Another strand of work has argued that LVR would be reduced or eliminated if orders were sold in *batch auctions*, rather than placed directly in mempools; this design is discussed in Ramseyer et al. [2022] and Canidio and Fritsch [2023], and related “request-for-quotes” mechanisms are now in use in a number of widely used protocols.³⁵

³³A similar design is proposed in Krishnamachari et al. [2021b].

³⁴Note that one subtlety about the zero fee design is that, in addition to rebalancing arbitrage profits, APs would be able to capture reversion arbitrage profits; hence, LVR provides a lower bound for how much revenue wallets could capture from having preferential access to arbitrage the pool.

³⁵“Batching” is used in CoWswap; a number of AMM protocols have also implemented related RFQ-like systems,

We believe that these are interesting directions for future AMM design research.

An earlier 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 Lau, 2021, Krishnamachari et al., 2021a]. Our work suggests that bonding functions only affect LVR insofar as they change the marginal liquidity of CFMMs. In certain cases, it may be desirable for marginal liquidity to vary with prices in a systematic way — this may be the case, for example, for stablecoin swap pairs. In other cases, different bonding functions, or “universal” AMM designs such as Uniswap v3, have equivalent infinitesimal LVR whenever they have equivalent amounts of marginal liquidity for an asset pair.

11. Conclusion

Automated market makers are *decentralized* market makers. Market participants can both trade with AMMs, and contribute capital to participate in the liquidity provision activities of the AMM. In the second case, the AMM behaves essentially as an investment, similar to mutual funds or ETFs in traditional financial markets. In this paper, we have constructed a framework for understanding the returns to investments in AMM LP positions. We provide a simple model, in which LP P&L cleanly decomposes into a “beta-like” component reflecting market risk, and an “alpha-like” component reflecting the microstructural forces of fee revenue versus adverse selection losses. These components are straightforward to separate empirically: the beta component can be eliminated simply by subtracting the profits of the rebalancing strategy from LP P&L. Our results have implications for how to measure the returns to LP positions as investments.

References

- Austin Adams, Ciamac Moallemi, Sara Reynolds, and Dan Robinson. am-amm: An auction-managed automated market maker. *arXiv preprint arXiv:2403.03367*, 2024.
- Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core, 2021.
- Algebra Protocol. ALGEBRA Ecosystem: Decentralized exchange, 2023. URL <https://algebra.finance/static/Algebra%20Tech%20Paper-15411d15f8653a81d5f7f574bfe655ad.pdf>.
- 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.
-
- such as Uniswap X, linch fusion, and Paraswap AugustusRFQ.

- 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.
- Aori Protocol. Trading on Aori, 2023. URL <https://www.aori.io/docs/trading>.
- 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.
- Matteo Aquilina, Sean Foley, Leonardo Gambacorta, and William Krekel. Decentralised dealers? examining liquidity provision in decentralised exchanges. Technical report, Bank for International Settlements, 2024.
- Patrick Augustin, Roy Chen-Zhang, and Donghwa Shin. Reaching for yield in decentralized financial markets. Working paper, 2024.
- Andrea Barbon and Angelo Ranaldo. On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. *arXiv preprint arXiv:2112.07386*, 2021.
- Dimitri P Bertsekas. *Control of uncertain systems with a set-membership description of the uncertainty*. PhD thesis, Massachusetts Institute of Technology, 1971.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
- Nassib Boueri. G3m impermanent loss dynamics. *arXiv preprint arXiv:2108.06593*, 2021.
- Michael Brolley and Marius Zoican. On-demand fast trading on decentralized exchanges. *Finance Research Letters*, 51:103350, 2023.
- Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621, 2015.
- 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/.
- Andrea Canidio and Robin Fritsch. Arbitrageurs’ profits, lvr, and sandwich attacks: batch trading as an amm design response. *arXiv preprint arXiv:2307.02074*, 2023.
- 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 Carr and Dilip Madan. Towards a theory of volatility trading. *Option pricing, interest rates and risk management, handbooks in mathematical finance*, 22(7):458–476, 2001.
- 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.
- Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *Available at SSRN 4273989*, 2022.

- Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Predictable losses of liquidity provision in constant function markets and concentrated liquidity markets. *Applied Mathematical Finance*, 30(2):69–93, 2023.
- Cata Labs. Optimising lp performance part 2: Dynamic fees, Nov 2023. URL <https://blog.catalyst.exchange/optimising-lp-performance-part-2-dynamic-fees/>.
- 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.
- John C Cox, Stephen A Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979.
- CrocSwap. Usage of markout to calculate lp profitability in uniswap v3, 2022. URL <https://crocswap.medium.com/usage-of-markout-to-calculate-lp-profitability-in-uniswap-v3-e32773b1a88e>.
- Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.
- Jun Deng, Hua Zong, and Yun Wang. Static replication of impermanent loss for concentrated liquidity provision in decentralised markets. *Operations Research Letters*, 51(3):206–211, 2023.
- Avinash K Dixit and Robert S Pindyck. *Investment under uncertainty*. Princeton university press, 1994.
- Atis Elsts. Conceptualizing Uniswap v3 LP profit and loss, January 2023. URL <https://atise.medium.com/conceptualizing-uniswap-v3-lp-profit-and-loss-ecbae6e09644>.
- Alex Evans. Liquidity provider returns in geometric mean markets. *arXiv preprint arXiv:2006.08806*, 2020.
- Chuck Fang. Learning financial innovations: Experience from automated market makers. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*, 2024.
- Fenbushi Capital. Ending LP’s Losing Game: Exploring the Loss-Versus-Rebalancing (LVR) Problem and its Solutions, Oct 2023. URL <https://mirror.xyz/0xbdA5bCe76bF62d97D9C9dF0425CC10079Df1aD75/bW0CccjVC7eoYK0zgmjXFdhWdc8rrUL6Yei-eugF52s>.
- Sean Foley, Peter O’Neill, and Tālis J Putniņš. A better market design? applying ‘automated market makers’ to traditional financial markets. *Available at SSRN 4459924*, 2023.
- Eric Forgy and Leo Lau. A family of multi-asset automated market makers. *arXiv preprint arXiv:2111.08115*, 2021.
- Masaaki Fukasawa, Basile Maire, and Marcus Wunsch. Weighted variance swaps hedge against impermanent loss. *Available at SSRN 4095029*, 2022.
- Lawrence R Glosten and Paul R Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14(1):71–100, 1985.
- Joel Hasbrouck, Thomas J Rivera, and Fahad Saleh. The need for fees at a dex: How increases in fees can increase dex trading volume. *Available at SSRN*, 2022.
- Joel Hasbrouck, Thomas J Rivera, and Fahad Saleh. An economic model of a decentralized exchange with concentrated liquidity. *Management Science*, 2025.
- Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. Risks and returns of uniswap v3 liquidity providers. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, pages 89–101, 2022.

- Roger D. Huang and Hans R. Stoll. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3):313–357, 1996.
- Josojo. MEV capturing AMM (McAMM), 2022. URL <https://ethresear.ch/t/mev-capturing-amm-mcamm/13336>.
- Jump Crypto. Pyth Defender white paper draft. Working paper, September 2022.
- Ariah Klages-Mundt and Steffen Schuldenzucker. Elliptic Concentrated Liquidity Pool (E-CLP): Technical Overview, 2022.
- 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, 2021a. doi: 10.1109/ICBC51069.2021.9461100.
- Bhaskar Krishnamachari, Qi Feng, and Eugenio Grippo. Dynamic curves for decentralized autonomous cryptocurrency exchanges. *arXiv preprint arXiv:2101.02778*, 2021b.
- Guillaume Lambert. Website at <https://lambert-guillaume.medium.com/>, 2022.
- Alfred Lehar and Christine A Parlour. Decentralized exchanges. Technical report, Working paper, 2021.
- Alfred Lehar, Christine A Parlour, and Marius Zoican. Liquidity fragmentation on decentralized exchanges. *Available at SSRN 4267429*, 2024.
- Gordon Liao and Dan Robinson. The dominance of uniswap v3 liquidity. <https://blog.uniswap.org/uniswap-v3-dominance>, 2022. Accessed: 2024-05-01.
- John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, February 1965. doi: 10.2307/1924119.
- 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>.
- Jason Milionis, Ciamac C Moallemi, and Tim Roughgarden. Automated market making and arbitrage profits in the presence of fees. *arXiv preprint arXiv:2305.14604*, 2023a.
- Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. Complexity-Approximation Trade-offs in Exchange Mechanisms: AMMs vs. LOBs. In *Financial Cryptography and Data Security*. Springer International Publishing, 2023b. ISBN 978-3-031-47753-9.
- Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. A Myersonian Framework for Optimal Liquidity Provision in Automated Market Makers. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- 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.

Geoffrey Ramseyer, Mohak Goyal, Ashish Goel, and David Mazières. Augmenting batch exchanges with constant function market makers. *arXiv preprint arXiv:2210.04929*, 2022.

Stephen A Ross. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360, 1976.

William F Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3):425–442, 1964.

Carl P Simon, Lawrence Blume, et al. *Mathematics for economists*, volume 7. Norton New York, 1994.

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.

Appendix

A. Proofs

A.1. Proof of Lemma 1

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.

A.2. Proof of Lemma 2

First, we show that LVR_t is equal to expression (7). The smoothness condition of Assumption 1 Part 2 allows us to apply Itô's lemma to $V(\cdot)$ to obtain

$$\begin{aligned} dV_t &= V'(P_t) dP_t + \frac{1}{2} V''(P_t) (dP_t)^2 \\ &= V'(P_t) dP_t + \frac{1}{2} V''(P_t) \sigma^2 P_t^2 dt \\ &= x^*(P_t) dP_t + \frac{1}{2} V''(P_t) \sigma^2 P_t^2 dt, \end{aligned} \tag{31}$$

where the last step follows from Lemma 1 Part 2. Comparing with (5), we obtain (7). Finally, the fact that $\ell(\sigma, P) \geq 0$ follows from Lemma 1 Part 3.

Next, we show that the cumulative profits of rebalancing arbitrageurs over the time interval $[0, T]$ is equal to LVR_T . 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

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

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 1 Part 3. 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

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

Hence, the cumulative profits of rebalancing arbitrageurs from time 0 to time T are equal to LVR, defined in (6).

A.3. LVR, Marginal Liquidity, and Bonding Function Curvature

For sufficiently smooth CFMM bonding functions, the marginal liquidity can be expressed in terms of derivatives of the CFMM bonding function:

$$\frac{dx}{dP} = \frac{\frac{\partial f}{\partial y}}{\left(\frac{\partial^2 f}{\partial x^2} + P^2 \frac{\partial^2 f}{\partial y^2} - 2P \frac{\partial^2 f}{\partial x \partial y} \right)} \quad (32)$$

Qualitatively, (32) implies that marginal liquidity, $x'(P)$, is related to the curvature of the CFMM invariant curves. The denominator of (32) is equal to P^2 times the negative of the determi-

nant of the bordered Hessian of f . f is strictly quasiconcave — that is, the upper level sets of f are convex — if and only if this determinant is positive; moreover, the magnitude of the determinant is related to the curvature of the level curves of f [p. 542 Simon et al., 1994]. Thus, CFMM invariants with “flatter”, more linear level curves will have greater marginal liquidity $\frac{dx}{dP}$, and also greater LVR.

(32) is also useful because it can be used to calculate $\frac{dx}{dP}$ using analytic expressions for bonding functions, which may be useful for computing LVR in practice.

A.3.1. Derivation of (32)

The Lagrangian of the pool expenditure minimization problem, (1), is:

$$\Lambda = Px + y + \lambda [f(x, y) - L]$$

The optimal solution is characterized by the FOCs:

$$\frac{\partial \Lambda}{\partial x} : P + \lambda \frac{\partial f}{\partial x} = 0 \tag{33}$$

$$\frac{\partial \Lambda}{\partial y} : 1 + \lambda \frac{\partial f}{\partial y} = 0 \tag{34}$$

$$\frac{\partial \Lambda}{\partial \lambda} : f(x, y) - L = 0 \tag{35}$$

Now, we will take $\frac{dx}{dP}$ by applying the implicit function theorem to this system of first-order conditions. The derivatives of the FOCs are:

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial x} : 1$$

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial y} : 0$$

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial \lambda} : 0$$

$$\frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial x} : \lambda \frac{\partial^2 f}{\partial x^2}$$

$$\frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial y} : \lambda \frac{\partial^2 f}{\partial x \partial y}$$

$$\frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial \lambda} : \frac{\partial f}{\partial x}$$

$$\frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial x} : \lambda \frac{\partial^2 f}{\partial x \partial y}$$

$$\frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial y} : \lambda \frac{\partial^2 f}{\partial y^2}$$

$$\frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial \lambda} : \frac{\partial f}{\partial y}$$

$$\frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial x} : -\frac{\partial f}{\partial x}$$

$$\frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial y} : -\frac{\partial f}{\partial y}$$

$$\frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial \lambda} : 0$$

Hence, we wish to solve the following system of equations for $\frac{dx}{dP}$:

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial x} + \frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial x} \frac{dx}{dP} + \frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial x} \frac{dy}{dP} + \frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial x} \frac{d\lambda}{dP} = 0$$

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial y} + \frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial y} \frac{dx}{dP} + \frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial y} \frac{dy}{dP} + \frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial y} \frac{d\lambda}{dP} = 0$$

$$\frac{\partial}{\partial P} \frac{\partial \Lambda}{\partial \lambda} + \frac{\partial}{\partial x} \frac{\partial \Lambda}{\partial \lambda} \frac{dx}{dP} + \frac{\partial}{\partial y} \frac{\partial \Lambda}{\partial \lambda} \frac{dy}{dP} + \frac{\partial}{\partial \lambda} \frac{\partial \Lambda}{\partial \lambda} \frac{d\lambda}{dP} = 0$$

Substituting, we have:

$$1 + \lambda \frac{\partial^2 f}{\partial x^2} \frac{dx}{dP} + \lambda \frac{\partial^2 f}{\partial x \partial y} \frac{dy}{dP} - \frac{\partial f}{\partial x} \frac{d\lambda}{dP} = 0 \quad (36)$$

$$0 + \lambda \frac{\partial^2 f}{\partial x \partial y} \frac{dx}{dP} + \lambda \frac{\partial^2 f}{\partial y^2} \frac{dy}{dP} - \frac{\partial f}{\partial y} \frac{d\lambda}{dP} = 0 \quad (37)$$

$$\frac{\partial f}{\partial x} \frac{dx}{dP} + \frac{\partial f}{\partial y} \frac{dy}{dP} = 0 \quad (38)$$

Now, applying (33) and (34), we have:

$$\frac{\partial f}{\partial x} = P \frac{\partial f}{\partial y} \quad (39)$$

Thus, we can simplify (38) to:

$$\frac{dy}{dP} = -P \frac{dx}{dP}$$

Substituting into (36) and (37), we have:

$$1 + \lambda \frac{\partial^2 f}{\partial x^2} \frac{dx}{dP} + \lambda \frac{\partial^2 f}{\partial x \partial y} \left(-P \frac{dx}{dP} \right) - \frac{\partial f}{\partial x} \frac{d\lambda}{dP} = 0$$

$$0 + \lambda \frac{\partial^2 f}{\partial x \partial y} \frac{dx}{dP} + \lambda \frac{\partial^2 f}{\partial y^2} \left(-P \frac{dx}{dP} \right) - \frac{\partial f}{\partial y} \frac{d\lambda}{dP} = 0$$

Rearranging,

$$\left(\frac{\partial^2 f}{\partial x^2} - P \frac{\partial^2 f}{\partial x \partial y} \right) \lambda \frac{dx}{dP} = \frac{\partial f}{\partial x} \frac{d\lambda}{dP} - 1 \quad (40)$$

$$\left(\frac{\partial^2 f}{\partial x \partial y} - P \frac{\partial^2 f}{\partial y^2} \right) \lambda \frac{dx}{dP} = \frac{\partial f}{\partial y} \frac{d\lambda}{dP} \quad (41)$$

Now, multiply (41) by P , applying (39), and subtract (40), to get:

$$\left(\frac{\partial^2 f}{\partial x^2} + P^2 \frac{\partial^2 f}{\partial y^2} - 2P \frac{\partial^2 f}{\partial x \partial y} \right) \lambda \frac{dx}{dP} = -1$$

Now,

$$\frac{dx}{dP} = \frac{-1}{\lambda \left(\frac{\partial^2 f}{\partial x^2} + P^2 \frac{\partial^2 f}{\partial y^2} - 2P \frac{\partial^2 f}{\partial x \partial y} \right)}$$

Now, to solve for λ , we simply use (34). Hence, we have:

$$\frac{dx}{dP} = \frac{\frac{\partial f}{\partial y}}{\left(\frac{\partial^2 f}{\partial x^2} + P^2 \frac{\partial^2 f}{\partial y^2} - 2P \frac{\partial^2 f}{\partial x \partial y} \right)} \quad (42)$$

This is (32).

B. Other Results

B.1. LVR and Impermanent Loss in a Two-Step Binomial Tree

Relationships to option strategies. The binomial tree example also helps illustrate the analogy between CFMM LP payoffs and European options. The time-2 difference between the payoffs of the rebalancing strategy and the buy-and-hold strategy is positive when prices mean-revert, and negative when prices diverge. Hence, payoffs are similar to those of a short European straddle or strangle position, which involves selling calls and puts which expire after two periods. The positive payoff when prices revert can be thought of as the option premia collected from selling the straddle, and the negative payoffs when prices diverge can be thought of as the payouts to the option buyer, which are made if either the call or the put sold expire in-the-money. The CFMM LP position has a similar pattern of payoffs, but makes 0 profits if prices end where they started. An CFMM LP position, ignoring fees, can thus be thought of like giving away a straddle position, without collecting any upfront option premia. Viewed this way, the equivalence between the rebalancing strategy and the static European strangle on the binomial tree reflects the classic idea that static option positions can be replicated by dynamically trading the underlying asset; in this case, European straddles and strangles are replicated by a strategy which sells the risky asset when prices increase and buys when prices decrease.

B.2. 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., $\ell(\sigma, P)/V(P)$, is a constant. The following theorem establishes that these are essentially the only CFMMs for which this is true:

Theorem 2. *Suppose a CFMM satisfies*

$$\frac{\ell(\sigma, P)}{V(P)} = c(\sigma), \quad \forall P \geq 0. \quad (43)$$

Then, we have

$$V(P) = L_1 P^{\theta(\sigma)} + L_2 P^{1-\theta(\sigma)}, \quad (44)$$

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

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

Comparing with Example 2, observe that (44) states that 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 (18).

Proof of Theorem 2. We construct the following ODE from the (43) along with (8),

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

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 (1) 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. ■

B.3. 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.,

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

with covariance matrix of returns $\Sigma_t \in \mathbb{R}^{n \times n}$, $\Sigma_t \succeq 0$, 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 1, 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 3. *For all prices $P \in \mathbb{R}_+^n$, the pool value function satisfies:*

1. $V(P) \geq 0$.
2. $\nabla V(P) = x^*(P) \geq 0$.
3. $\nabla^2 V(P) = \nabla x^*(P) \preceq 0$.

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 Lemma 2:

Theorem 3. *Loss-versus-rebalancing takes the form*

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

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

$$\ell(\Sigma, P) \triangleq -\frac{1}{2} \text{tr} [\text{diag}(P) \Sigma \text{diag}(P) \nabla x^*(P)] \geq 0,$$

where we have applied Lemma 3. In the case where $\Sigma = \sigma^2 I$, i.e., i.i.d. assets, we have that

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

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

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

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

The rest of the result follows as in the proof of Lemma 2. ■

C. Data and Measurement

C.1. Data

Prices. We download minute-level USDC-ETH prices from the Binance API. We use close prices at the end of each minute for P_t .

Uniswap. We download data on the Uniswap v2 WETH-USDC pool from Dune Analytics, a data provider which aggregates data from the Ethereum blockchain into SQL databases. The queries we use to extract this data are included in Appendix C.2.

Mints and burns. In each minute, we observe the gross amounts of each asset in which are withdrawn through “burns”, and deposited through “mints”. Let $(x_t^{\text{mint}}, y_t^{\text{mint}})$ and $(x_t^{\text{burn}}, y_t^{\text{burn}})$ be the total amounts of each asset x and y which are minted and burned respectively, between time $t - 1$ and time t . We will value minted and burned assets at the time- t closing price P_t . Thus, the monetary value of mints and burns respectively are:

$$\Pi_t^{\text{mint}} \triangleq y_t^{\text{mint}} + P_t x_t^{\text{mint}}, \quad \Pi_t^{\text{burn}} \triangleq y_t^{\text{burn}} + P_t x_t^{\text{burn}}.$$

Let (x_t, y_t) be the total asset holdings of the pool at time t . As in the model, define the pool value at time t by

$$V_t \triangleq y_t + P_t x_t.$$

We can calculate $\Delta LP \text{ P\&L}_t$, the change in P&L of the pool, from period $t - 1$ to t , as

$$\Delta LP \text{ P\&L}_t \triangleq V_t + \Pi_t^{\text{burn}} - \Pi_t^{\text{mint}} - V_{t-1}. \quad (45)$$

In words, this is the value of pool reserves at time t valued at price P_t , plus burned assets and minus minted assets valued at P_t , minus the value of pool reserves at time $t - 1$ valued at P_{t-1} . Note that, in contrast to our simplifying assumption in the model that fees are paid in the numéraire, in practice in Uniswap v2 fees are paid directly into the pool reserves; hence, the pool P&L includes transaction fees paid into the pool.

Rebalancing strategy. We rebalance the pool at different time frequencies. For each rebalancing frequency, we compute the returns of a strategy which at any point in time holds as much ETH as the pool holds at the start of the period. For example, if the rebalancing frequency is daily, we set x_t^{RB} at any minute t equal to the LP pool reserves at the start of the day containing the minute t . We then calculate the returns on the rebalancing strategy using expression (26), that is:

$$\Delta RB \text{ P\&L}_t = x_t^{RB} (P_{t+1} - P_t). \quad (46)$$

Fees. In each minute, we compute the gross amount of each asset in the pair bought and sold. The Uniswap v2 pool has a fixed fee rate of 30bps on the contributed asset; we thus calculate fees in each asset by multiplying the gross amount contributed of each asset by 0.003. Call these fees x_t^{fee} and y_t^{fee} in period t . We value fees at the period t price; thus, the monetary value of fees in period t , which we will call ΔFEE_t , is:

$$\Delta FEE_t \triangleq y_t^{fee} + P_t x_t^{fee}. \quad (47)$$

Note also that Uniswap v2 allows for “flash loans”, in which assets are withdrawn and returned within a single transaction. Flash loans which are accounted for simply as swaps in Uniswap v2; thus, the calculation in (47) also correctly accounts for flash loans revenues.³⁶

LVR. We compute a realized daily volatility using USDC-ETH prices from the Binance API sampled at 60 minute intervals. Let ΔLVR_t be the increment of LVR in period t . As in Example 3, we then

³⁶See the Uniswap v2 documentation of flash swaps.

calculate ΔLVR_t simply as

$$\Delta\text{LVR}_t \triangleq \frac{\hat{\sigma}_t^2}{8} \times V_t \times \Delta t, \quad (48)$$

where $\hat{\sigma}_t$ denotes the realized daily volatility estimate for the day containing period t , and $\Delta t = 1/(24 \times 60)$ corresponds to a one minute period. This is a discrete approximation of (27).

Adding everything up. We then calculate the cumulative returns, the left side of (24), as:

$$\text{LP P\&L}_t - \int_0^t x^*(P_s) dP_s \triangleq \sum_{t=1}^T (\Delta\text{P\&L}_t - \Delta\text{RB P\&L}_t),$$

using the definitions of $\Delta\text{LP P\&L}_t$ and $\Delta\text{RB P\&L}_t$ in (45) and (46) respectively. Note that different rebalancing frequencies give slightly different values of for $\Delta\text{RB P\&L}_t$. We calculate the right side of (24) as:

$$\text{FEE}_t - \text{LVR}_t \triangleq \sum_{t=1}^T (\Delta\text{FEE}_t - \Delta\text{LVR}_t),$$

using the definitions of ΔFEE_t and ΔLVR_t in (47) and (48) respectively.

C.2. Dune SQL Queries

This appendix contains the SQL queries we use on Dune to extract Uniswap v2 WETH-USDC data.

Mints.

```
1 SELECT to_char(evt_block_time, 'YYYY-MM-DD"T"HH24:MI:SSOF') AS ts, *
2 FROM uniswap_v2."Pair_evt_Mint"
3 WHERE contract_address = '\xb4e16d0168e52d35cacd2c6185b44281ec28c9dc'
4 ORDER BY evt_block_number, evt_index ASC
```

Burns.

```
1 SELECT to_char(evt_block_time, 'YYYY-MM-DD"T"HH24:MI:SSOF') AS ts, *
2 FROM uniswap_v2."Pair_evt_Burn"
3 WHERE contract_address = '\xb4e16d0168e52d35cacd2c6185b44281ec28c9dc'
4 ORDER BY evt_block_number, evt_index ASC
```

Trades.

```
1 SELECT
2 date_trunc('minute', evt_block_time) AS minute,
```

```

3 SUM("amount0In") as "amount0In",
4 SUM("amount1In") as "amount1In",
5 SUM("amount00Out") as "amount00Out",
6 SUM("amount10Out") as "amount10Out"
7 FROM uniswap_v2."Pair_evt_Swap"
8 WHERE contract_address = '\xb4e16d0168e52d35cacd2c6185b44281ec28c9dc'
9 GROUP BY 1
10 ORDER BY 1 ASC

```

Pool reserves.

```

1 SELECT
2 minute,
3 latest_reserves[3] AS reserve0,
4 latest_reserves[4] AS reserve1
5 FROM
6 (SELECT date_trunc('minute', evt_block_time) AS minute,
7 (SELECT MAX(ARRAY[evt_block_number, evt_index, reserve0, reserve1]))
8 AS latest_reserves
9 FROM uniswap_v2."Pair_evt_Sync"
10 WHERE contract_address = '\xb4e16d0168e52d35cacd2c6185b44281ec28c9dc'
11 GROUP BY 1) AS day_reserves
12 ORDER BY 1 ASC

```