Reserve-Matching Portfolios and Benchmark Losses in AMMs

Gustavo Grivol
Josh Reed

New York University
Stern School of Business

May 22, 2026

1 Introduction

Automated market makers (AMMs) in decentralized finance maintain liquidity pools whose reserves are constrained to satisfy a specified bonding curve. As external market prices evolve, arbitrageurs trade against the AMM and thereby continuously synchronize the reserve state of the pool with prices on centralized exchanges. This synchronization process changes the reserve composition of the liquidity pool and generates benchmark losses relative to frictionless dynamically rebalanced portfolios.

A central benchmark introduced in the recent literature is loss-versus-rebalancing (LVR), which compares the value of a continuously synchronized liquidity pool to the value of a continuously synchronized self-financing portfolio holding the same reserve exposure. The classical LVR framework shows that benchmark underperformance arises from the interaction between price volatility and the curvature of the AMM bonding curve.

Continuous synchronization of the benchmark portfolio, however, is only one possible rebalancing mechanism. Throughout this paper, the benchmark portfolio is permitted to rebalance its reserve exposure intermittently. The timing of these benchmark rebalancing events may itself affect benchmark performance. This raises a natural question: how should liquidity pool performance be evaluated relative to dynamically rebalanced benchmark portfolios when benchmark rebalancing may occur intermittently rather than continuously?

In this paper, we introduce a general reserve-matching framework for studying benchmark performance under arbitrary rebalancing rules. Rather than continuously synchronizing reserve exposure with the AMM, a reserve-matching portfolio updates its risky holdings only at designated synchronization times and maintains those holdings between synchronization events. The resulting class of benchmark portfolios includes continuously synchronized reserve-matching portfolios, discretely synchronized portfolios, absolutely continuous synchronization policies, and singular synchronization policies within a unified framework.

A key structural result of the paper is that admissible reserve-matching portfolios are dominating portfolios: their wealth processes remain pathwise greater than or equal to the marked-to-market value of the liquidity pool. Geometrically, reserve synchronization causes benchmark wealth to evolve along tangent lines to the concave liquidity-pool value function. This tangent-line structure plays a central role throughout the paper and provides a unified geometric interpretation of reserve-matching benchmarks.

The reserve-matching framework naturally leads to a generalized benchmark problem. We define the loss-versus-optimal-rebalancing (LVOR) benchmark by optimizing expected benchmark performance over all admissible reserve-matching synchronization policies. Classical LVR then arises as the special case corresponding to continuous synchronization.

A central contribution of the paper is a decomposition of benchmark underperformance into two distinct components. The first is a volatility component generated by the curvature of the liquidity-pool value function together with the quadratic variation of the reference price process. This corresponds precisely to the classical LVR effect. The second is a timing component generated by synchronization under nonmartingale price dynamics. This decomposition shows that classical LVR captures only part of the benchmark effects associated with liquidity provision when synchronization timing is endogenous.

We then characterize optimal reserve-matching synchronization policies under several classes of price dynamics. When the reference price process is a martingale, synchronization timing has no expected value and every admissible reserve-matching policy achieves the same expected benchmark performance. Under submartingale and supermartingale dynamics, however, synchronization timing becomes economically valuable, and optimal policies are characterized by synchronization at running minima or running maxima of the reference price process.

Our framework also highlights the importance of global properties of price dynamics for AMM benchmark analysis. In particular, we study a class of strict local martingale price processes that exhibit no explicit drift term in their stochastic differential representation, yet nevertheless generate positive synchronization timing value through their global supermartingale structure. This demonstrates that benchmark performance may depend not only on the local stochastic behavior of the price process, but also on whether the process is globally a martingale or merely a strict local martingale.

Finally, we illustrate the economic magnitude of synchronization timing effects using historical ETH-USDC data from Binance and Uniswap V2. The empirical results show that realized benchmark performance may vary substantially across synchronization policies, particularly during periods exhibiting directional price movements.

Our paper contributes to the growing literature on AMM benchmark losses and liquidity provision. Relative to the existing LVR framework, our contribution is to endogenize synchronization timing through a general class of reserve-matching portfolios and to characterize the interaction between volatility-driven benchmark effects and synchronization timing effects under general semimartingale price dynamics.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the liquidity-pool framework and reserve dynamics. Section 4 defines reserve-matching portfolios and establishes their domination properties. Section 5 introduces the LVOR benchmark and develops its decomposition. Section 6 characterizes optimal reserve-matching policies under several classes of price dynamics. Section 7 studies strict local martingale price processes, and Section 8 presents the empirical analysis. Section 9 concludes.

2 Literature Review

The growing literature on liquidity provision in CFMMs has attracted significant attention in recent years [Angeris et al.(2019), Drossos et al.(2025)]. This literature studies the economic trade-offs faced by liquidity providers (LPs), who earn trading fees while simultaneously bearing losses generated by reserve rebalancing as prices move across trading venues. In particular, when the external reference price moves away from its initial level, the reserve composition of the liquidity pool changes in a way that may reduce the pool's value relative to alternative benchmark portfolios. As a result, the interaction between fee income, reserve exposure, and benchmark underperformance gives rise to complex liquidity-provision and equilibrium dynamics [Hasbrouck et al.(2026)].

Although there is rich literature on designing fee mechanisms to incentivize liquidity providers, [Adams et al.(2025)] and [Baggiani et al.(2025)], our focus will be in measuring the opportunity cost of not investing in an alternative similar portfolio. Initial research in AMMs introduced the concept of impermanent loss, defined as the difference between the return of an LP position and that of a static portfolio with the same initial allocation [Pintail(2019), Loesch et al.(2021)]. To further isolate the risks associated with arbitrage and dynamic rebalancing, [Milionis et al.(2022)] proposed the loss-versus-rebalancing (LVR) metric, which compares LP returns without fees to those of a continuously rebalanced portfolio. In settings with a risk-free asset, an additional opportunity cost emerges from the inability to realize and reinvest excess returns into the risk-free asset [Cartea et al.(2023)]. These loss metrics have motivated a broad literature on optimal liquidity provision, including studies on optimal withdrawal timing [Capponi and Zhu(2024)] and the design of CFMM bonding curve curvature [Goyal et al.(2023)].

A more recent development incorporates block-generation dynamics into the valuation of LP holdings and arbitrage losses. This perspective has been studied under Poisson block arrivals in [Milionis et al.(2023)] and under more general block-time processes in [Nezlobin and Tassy(2025)]. The dynamic behavior of LP portfolios under these execution constraints is central to our analysis. In particular, we depart from the standard assumption of comparison with a continuously rebalanced portfolio and instead consider a broader class of benchmark portfolios. The construction of replicating portfolios as hedging instruments for LPs is an active area of research [Deng et al.(2023), Maire and Wunsch(2024)], including approaches based on options and other derivative instruments [Clark(2020), Fukasawa et al.(2023)].

3 Model Setup

3.1 Reference Price Process

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions. We interpret $\mathcal{F}_t$ as the public information available by time t. Let $p = (p_t)_{t\geq 0}$ denote the external reference price of a risky asset, expressed in units of a numeraire asset. Throughout the paper, p may be interpreted as the quoted price of the risky asset on an infinitely liquid centralized exchange.

We assume that p is a nonnegative, continuous semimartingale adapted to $(\mathcal{F}_t)$. Accordingly, p admits the canonical decomposition

$$p_t \quad = \quad p_0 + M_t + A_t, \tag{1}$$

where M is a local martingale and A is an adapted finite-variation process. Define the hitting time

$$\tau_0 = \inf\{t \geq 0 : p_t = 0\}. \tag{2}$$

We allow the possibility of the reference price reaching zero and assume that zero is absorbing, so that $p_t = 0$ for all $t \geq \tau_0$. Stronger assumptions, including martingale, submartingale, and Markovian diffusion specifications, are imposed in later sections when deriving sharper characterization results.

3.2 Trading Venues

We next consider an economy with two trading venues for the risky asset. The first venue is a centralized exchange (CEX) that is assumed to be infinitely liquid and continuously quotes the exogenous reference price p_t. Agents may trade any desired quantity of the risky asset on the centralized exchange at a price p_t, without transaction costs or price impact. The fundamental value of the risky asset is therefore instantaneously incorporated into prices on the centralized exchange.

The second venue is a decentralized exchange (DEX) implemented as an automated market maker. At each time t, the DEX holds reserves (x_t, y_t), where x_t denotes units of the risky asset and y_t denotes units of the numeraire asset. Throughout the paper, we assume that the liquidity pool remains synchronized with the external reference price process p_t. Accordingly, the reserve configuration is determined by the bonding curve together with the prevailing external price, so that the reserve state evolves along the bonding curve as p_t changes over time. The benchmark gains and losses studied in this paper arise from comparing the liquidity pool to dynamically rebalanced benchmark portfolios operating under different synchronization policies.

3.3 Pool Reserves and Pricing Rule

The decentralized exchange is characterized by a reserve state (x, y) of its liquidity pool, where x denotes units of the risky asset and y denotes units of the numeraire asset. Trades against the pool change the reserve states subject to its trading rule. Specifically, the set of feasible reserve states is specified by an invariant function

$$F(x, y) = k,$$

where $k > 0$ represents the liquidity level of the pool. The reserve states are therefore constrained to lie on the level set associated with k. This formulation includes, for example, the constant-product market maker

$$xy = k.$$

Throughout the paper, we fix the liquidity level k and represent the corresponding level set as the graph of a bonding curve. Thus the feasible reserve states are written as

$$\mathcal{R} = \{(x, f(x)) : x > 0\},$$

where $f : (0, \infty) \to (0, \infty)$ gives the required numeraire reserve as a function of the risky-asset reserve. This reduction from the invariant representation $F(x, y) = k$ to the graph representation $y = f(x)$ allows us to study the pool through a one-dimensional reserve schedule.

Definition 3.1. *A function $f : (0, \infty) \to (0, \infty)$ is called a bonding curve if:*

1. *f is strictly decreasing, strictly convex, and twice continuously differentiable;*
2. $\lim_{x \downarrow 0} f(x) = \infty$ *and* $\lim_{x \uparrow \infty} f(x) = 0$.

The marginal price of the risky asset in units of the numeraire is determined by the slope of the bonding curve. If the pool is at reserve state $(x, f(x))$, then an infinitesimal withdrawal of Δx units of the risky asset requires an approximately offsetting deposit

$$\Delta y \approx -f'(x)\Delta x$$

of the numeraire asset. We therefore define the pool price associated with reserve level x by

$$p(x) = -f'(x).$$

Since f is strictly convex, $p(x)$ is decreasing in x: higher risky-asset reserves correspond to lower prices, while lower risky-asset reserves correspond to higher prices. The local price-impact coefficient is

$$p'(x) = -f''(x),$$

so that changes in reserves move prices according to the curvature of the bonding curve.

3.4 Synchronized Reserve State

Given an external reference price $p > 0$, the synchronized reserve state of the liquidity pool is defined as the reserve pair $(x(p), y(p))$ satisfying

$$-f'(x(p)) = p,$$

with

$$y(p) = f(x(p)).$$

The quantity $x(p)$ therefore represents the risky-asset reserve level consistent with the external price p. Since f is strictly convex and decreasing, the map $x(\cdot)$ is decreasing. Higher prices correspond to lower risky-asset reserves, while lower prices correspond to larger risky holdings by the pool.

3.5 Liquidity Pool Value

Given a synchronized reserve state, we now define the corresponding marked-to-market value of the liquidity pool. Since the pool holds both risky and numeraire reserves, its total value depends on the prevailing external reference price together with the synchronized reserve pair.

Suppose the liquidity pool is synchronized with an external reference price $p > 0$. The marked-to-market value of the liquidity pool is then

$$W(p) = f(x(p)) + p\,x(p),$$

where $x(p)$ denotes the risky reserve satisfying

$$-f'(x(p)) = p.$$

The first term represents the numeraire reserves held by the pool, while the second term represents the market value of the risky reserves evaluated at the external reference price p.

Differentiating yields

$$W'(p) = f'(x(p))x'(p) + x(p) + p\,x'(p).$$

Using the identity

$$f'(x(p)) = -p,$$

the terms involving $x'(p)$ cancel, giving

$$W'(p) = x(p).$$

Differentiating once more yields

$$W''(p) = x'(p).$$

Since $x(\cdot)$ is decreasing, the liquidity-pool value function W is thus concave.

The identity

$$W'(p) = x(p)$$

plays a central role throughout the paper. In particular, the risky reserve holdings of the liquidity pool coincide with the derivative of the pool value with respect to the external reference price.

The analysis in later sections requires repeated applications of Itô calculus to functions of the liquidity-pool value function. We therefore impose the following regularity condition.

Assumption 3.2. *The liquidity-pool value function W belongs to $C^3((0,\infty))$.*

This assumption is purely technical and is used to justify the semimartingale calculations appearing later in the paper.

3.6 Boundary Behavior of the Liquidity Pool Value

Because the reference price process may reach the boundary $p = 0$, we impose a regularity condition on the liquidity-pool value function near the boundary.

Assumption 3.3. *The liquidity-pool value function*

$$W(p) = f(x(p)) + px(p)$$

admits a finite limit as $p \downarrow 0$. That is,

$$\lim_{p \downarrow 0} W(p) = W(0) < \infty.$$

Under Definition 3.1, the synchronized risky reserve level satisfies

$$x(p) \uparrow \infty \quad \text{as } p \downarrow 0,$$

while the synchronized numeraire reserve level satisfies

$$f(x(p)) \downarrow 0 \quad \text{as } p \downarrow 0.$$

Thus, near the boundary, the liquidity pool accumulates arbitrarily large quantities of the risky asset while the numeraire reserves vanish. The assumption above requires that the total marked-to-market value of the liquidity pool nevertheless remains finite.

This condition is satisfied by many standard bonding curves, including the constant-product rule. It ensures that the liquidity pool value remains finite as the reference price approaches zero. Additional technical assumptions are imposed later where needed.

4 Dominating and Reserve-Matching Portfolios

We now introduce the class of benchmark portfolios whose values are used to compare against the value of the liquidity pool. Our starting point is the observation that the liquidity pool value process may be dominated by self-financing portfolios that dynamically trade the risky asset outside the AMM.

We first study general dominating portfolios. We then introduce a more specialized class of reserve-matching portfolios whose risky holdings are linked to the reserve holdings of the AMM. These reserve-matching portfolios will form the basis for the benchmark loss measures studied in subsequent sections.

4.1 Self-Financing Portfolios

We begin with a general class of dynamically traded benchmark portfolios. A portfolio is self-financing if changes in its wealth arise solely from gains and losses on its risky-asset position. Any self-financing portfolio may be specified by its trading strategy, which is a predictable process $h = (h_t)_{t \geq 0}$, where h_t denotes its holdings of the risky asset at time t. Let R_t denote the associated wealth process measured in units of the numeraire asset. Accordingly,

$$R_t = R_0 + \int_0^t h_s \, dp_s.$$

Definition 4.1. *A predictable process $h = (h_s)_{s \geq 0}$ is called admissible on a finite horizon $[0, t]$ if:*

1. *h is p-integrable on $[0, t]$, so that*
$$\int_0^u h_s \, dp_s$$
is well defined for every $u \in [0, t]$;

2. *the stochastic integral*
$$\int_0^u h_s \, dM_s$$
is a true martingale on $[0, t]$.

A sufficient condition for condition 2 is

$$\mathbb{E}\left[\int_0^t h_s^2 \, d\langle M\rangle_s\right] < \infty,$$

which ensures that

$$\int_0^u h_s \, dM_s$$

is a square-integrable martingale on $[0, t]$.

Condition 1 above ensures that the self-financing wealth process is well defined, while condition 2 rules out pathological trading strategies that extract nonzero expected gains from the martingale component of the reference price process. In particular, condition 2 implies that

$$E\left[\int_0^u h_s \, dM_s\right] = 0, \qquad 0 \leq u \leq t,$$

so that admissible strategies cannot systematically profit from martingale fluctuations alone.

4.2 Dominating Portfolios

We say that a self-financing portfolio with wealth process R dominates the liquidity pool on $[0, t]$ if

$$R_s \geq W_s, \qquad 0 \leq s \leq t,$$

where $W_s = W(p_s)$ denotes the marked-to-market value of the liquidity pool.

Dominating portfolios provide pathwise benchmark comparisons for liquidity provision. In particular, they represent self-financing trading strategies whose wealth remains greater than or equal to the value of the liquidity pool along every realized price path.

4.3 Examples of Dominating Portfolios

We now present several examples illustrating different mechanisms through which pathwise domination may arise. Some portfolios require continuous reserve adjustment, while others rebalance only intermittently or intervene only when the dominance constraint becomes nearly binding.

Example 4.2 (Continuous reserve-matching portfolio)**.** *Suppose the strategy*

$$h_t = x(p_t), \qquad t \geq 0,$$

is admissible. The associated self-financing wealth process is

$$R_t = R_0 + \int_0^t x(p_s) \, dp_s.$$

Applying Itô's formula to the pool value $W(p_t)$ yields

$$W(p_t) = W(p_0) + \int_0^t x(p_s)\, dp_s + \frac{1}{2}\int_0^t W''(p_s)\, d\langle p\rangle_s.$$

Hence,

$$R_t - W(p_t) = R_0 - W(p_0) - \frac{1}{2}\int_0^t W''(p_s)\, d\langle p\rangle_s.$$

If $R_0 \geq W(p_0)$ and W is concave, then $W'' \leq 0$, so the above difference is nonnegative for all t. Thus the continuously rebalanced reserve-matching portfolio dominates the liquidity pool.

Example 4.3 (Discrete rebalancing dominating portfolio)**.** *Pathwise domination does not require continuous reserve adjustment. Let*

$$0 = \tau_0 \leq \tau_1 \leq \tau_2 \leq \cdots$$

be an increasing sequence of stopping times. Suppose that on each interval $[\tau_n, \tau_{n+1})$, the portfolio holds a constant risky position h_n, so that

$$R_t = R_{\tau_n} + h_n(p_t - p_{\tau_n}), \qquad t \in [\tau_n, \tau_{n+1}).$$

Thus, between rebalancing times, the portfolio wealth evolves linearly as a function of the reference price. Pathwise domination requires this wealth line to remain above the graph of the liquidity-pool value function W. Suppose first that

$$R_{\tau_n} = W(p_{\tau_n}).$$

Since W is concave, the unique supporting tangent line through

$$(p_{\tau_n}, W(p_{\tau_n}))$$

is the tangent line to W at p_{τ_n}, which has slope

$$W'(p_{\tau_n}) = x(p_{\tau_n}).$$

Hence choosing

$$h_n = x(p_{\tau_n})$$

preserves domination on $[\tau_n, \tau_{n+1})$.

Now suppose instead that

$$R_{\tau_n} > W(p_{\tau_n}).$$

Under strict concavity of W, there exist unique prices

$$p_n^- < p_{\tau_n} < p_n^+$$

such that the tangent lines to W at p_n^- and p_n^+ both pass through

$$(p_{\tau_n}, R_{\tau_n}).$$

The existence of these tangent points follows from the boundary behavior

$$x(p) \uparrow \infty \quad \text{as } p \downarrow 0, \qquad x(p) \downarrow 0 \quad \text{as } p \uparrow \infty,$$

which implies that the graph of W becomes arbitrarily steep near zero and asymptotically flat at large prices. The corresponding tangent slopes are

$$x(p_n^-) = W'(p_n^-)$$

and

$$x(p_n^+) = W'(p_n^+).$$

Since $x(\cdot)$ is decreasing,

$$x(p_n^+) \leq x(p_{\tau_n}) \leq x(p_n^-).$$

Therefore, any slope

$$x(p_n^+) \leq h_n \leq x(p_n^-)$$

generates a wealth line lying above the graph of W, and therefore preserves domination between rebalancing times. Thus positive surplus over the liquidity-pool value permits a nontrivial range of admissible risky holdings while maintaining pathwise domination.

Example 4.4 (ε-boundary policy)**.** *Fix $\varepsilon > 0$, and let*

$$S_t := R_t - W(p_t)$$

denote the surplus relative to the liquidity pool. Consider a portfolio that follows an arbitrary self-financing strategy while $S_t \geq \varepsilon/2$.

When the surplus falls to $\varepsilon/2$, the portfolio switches to a supporting-line policy. More precisely, if

$$\sigma = \inf\{t \geq 0 : S_t \leq \varepsilon/2\},$$

then after time σ the portfolio holds the constant risky position

$$h_t = x(p_\sigma),$$

until the surplus again reaches ε. By the supporting-line argument of Example 4.2, the resulting wealth process continues to dominate the liquidity pool value during the supporting-line phase. The portfolio then returns to the arbitrary self-financing strategy until the surplus next falls to $\varepsilon/2$, at which point the procedure repeats. The gap between the intervention threshold $\varepsilon/2$ and the release threshold ε prevents immediate repeated switching.

This example illustrates that domination need not require continuous reserve adjustment. A portfolio may follow a flexible strategy away from the boundary and switch to a supporting-line policy only when its surplus becomes small.

These examples demonstrate that dominating portfolios form a broad class. Some require continuous adjustment, others rebalance only intermittently, and still others need intervene only when the domination constraint becomes nearly binding.

4.4 Reserve-Matching Portfolios

We now introduce a class of self-financing portfolios whose risky holdings track the reserve schedule of the liquidity pool. The central idea is that the portfolio need not continuously match the pool reserves at every instant. Rather, reserve holdings are updated only at selected synchronization times, and between such times the portfolio continues to hold the most recently matched reserve position.

Let $x(p)$ denote the quantity of the risky asset held by the liquidity pool when the external reference price is p. A reserve-matching portfolio takes risky holdings of the form

$$h_t = x(p_{\tau(t)}),$$

where $\tau(t) \leq t$ denotes the most recent synchronization time relevant at time t. Thus the portfolio holds the reserve exposure corresponding to the most recent synchronized price level.

To model synchronization rules in a unified way, let $L = (L_t)_{t \geq 0}$ be a nondecreasing, càdlàg, adapted process satisfying $L_0 = 0$. The points of increase of L represent times at which reserve synchronization occurs. Associated with L is the map

$$\tau(t) := L^{-1}(L(t+)), \qquad t \geq 0,$$

where L^{-1} denotes the left-continuous inverse

$$L^{-1}(a) := \inf\{s \geq 0 : L_s \geq a\}, \qquad a \geq 0.$$

The time $\tau(t)$ identifies the synchronization time whose reserve level determines the portfolio holdings at time t.

Definition 4.5. *A reserve-matching policy is any nondecreasing, càdlàg, adapted process L satisfying $L_0 = 0$. The associated reserve-matching portfolio has the trading strategy*

$$h_t = x(p_{\tau(t)}),$$

where

$$\tau(t) = L^{-1}(L(t+)).$$

A reserve-matching policy L is called admissible on the finite horizon $[0, t]$ if the associated trading strategy

$$h_t = x(p_{\tau(t)})$$

is admissible in the sense of Definition 4.1. We denote by $\mathcal{R}_t$ the class of admissible reserve-matching policies on the horizon $[0, t]$.

This framework permits highly general timing rules. In particular, a policy may generate finitely many, countably many, or even infinitely many synchronization events over a finite interval. While the latter need not be literally implementable, such policies are analytically useful and may often be approximated by discrete strategies.

A useful perspective is provided by the classical decomposition of monotone functions: every nondecreasing càdlàg process may be written as the sum of a jump component, an absolutely continuous component, and a singular continuous component. In the present setting, these correspond respectively to discrete synchronization times, continuous synchronization activity, and more irregular timing rules concentrated on thin sets of times.

We now present three representative examples.

Example 4.6 (Discrete reserve-matching policies)**.** *Let $(\tau_n)_{n\geq 1}$ be an increasing sequence of stopping times representing synchronization times. A simple discrete reserve-matching policy is*

$$L_t = \sum_{n\geq 1} \mathbf{1}_{\{\tau_n \leq t\}}.$$

Under this policy, reserve synchronization occurs only at the stopping times τ_n, and the portfolio holds the reserve exposure corresponding to the most recent synchronization date between updates.

Example 4.7 (Absolutely continuous reserve-matching policies)**.** *Let $g = (g_t)_{t\geq 0}$ be a nonnegative locally integrable adapted process and define*

$$L_t = \int_0^t g_s \, ds.$$

Synchronization then occurs continuously over time whenever $g_t > 0$, while intervals on which $g_t = 0$ correspond to periods with no reserve updating. Absolutely continuous policies provide a continuous-time analogue of gradual reserve adjustment, in contrast to the discrete synchronization rules of Example 4.6.

Example 4.8 (Singular reserve-matching policies)**.** *A reserve-matching policy may also be continuous and increasing while having derivative equal to zero almost everywhere. Such singular policies generate synchronization on sets of Lebesgue measure zero, yet may still produce infinitely many synchronization events over finite horizons.*

Canonical examples arise from local time processes of diffusions. These policies are mathematically more irregular than either the discrete or absolutely continuous cases, and illustrate that reserve synchronization may be concentrated on exceptionally thin sets of times.

The reserve-matching framework therefore accommodates a broad range of synchronization rules, ranging from discrete rebalancing times to continuously accumulating and singular synchronization mechanisms.

4.5 Pathwise Dominance of Reserve-Matching Portfolios

We now establish that reserve-matching portfolios are also dominating portfolios. In particular, when the portfolio updates its risky holdings according to the reserve schedule of the liquidity pool, its wealth process dominates the marked-to-market value of the liquidity pool.

Let $L \in \mathcal{R}_t$ be an admissible reserve-matching policy on the horizon $[0, t]$ with associated synchronization map τ, and let R^L denote the corresponding self-financing wealth process

$$R_t^L = R_0 + \int_0^t x(p_{\tau(s)})\, dp_s.$$

Recall that the liquidity-pool value process is

$$W_t = W(p_t).$$

We compare the reserve-matching portfolio to the liquidity pool through the surplus process

$$S_t^L := R_t^L - W_t.$$

Theorem 4.9. *Let $L \in \mathcal{R}_t$ be an admissible reserve-matching policy on the horizon $[0, t]$. If*

$$R_0 \geq W(p_0),$$

then

$$R_u^L \geq W(p_u), \qquad 0 \leq u \leq t.$$

That is, every admissible reserve-matching portfolio on $[0, t]$ *is also a dominating portfolio on* $[0, t]$.

Proof. Fix $u \in [0, t]$, and define

$$q_s := p_{\tau(s)}, \qquad 0 \le s \le u.$$

Consider the tangent process

$$T_s := W(q_s) + x(q_s)(p_s - q_s).$$

Since W is concave and $x = W'$, the graph of W lies below its tangent lines. Hence

$$T_s \ge W(p_s), \qquad 0 \le s \le u.$$

We now compute the dynamics of T. By Assumption 3.2, $W \in C^3((0, \infty))$. Applying Itô's formula to $W(q_s)$ gives

$$dW(q_s) = x(q_s)\, dq_s + \frac{1}{2} W''(q_s)\, d\langle q\rangle_s.$$

Next, applying the semimartingale product rule,

$$d\big(x(q_s)(p_s - q_s)\big) = x(q_s)\, d(p_s - q_s) + (p_s - q_s)\, dx(q_s) + d\langle x(q), p - q\rangle_s.$$

Applying Itô's formula to $x(q_s) = W'(q_s)$,

$$dx(q_s) = W''(q_s)\, dq_s + \frac{1}{2} W'''(q_s)\, d\langle q\rangle_s.$$

Substituting into the product rule yields

$$\begin{aligned} d\big(x(q_s)(p_s - q_s)\big) = {} & x(q_s)(dp_s - dq_s) \\ & + (p_s - q_s)\left(W''(q_s)\, dq_s + \frac{1}{2} W'''(q_s)\, d\langle q\rangle_s\right) \\ & + W''(q_s)\, d\langle q, p - q\rangle_s. \end{aligned}$$

Combining with the expression for $dW(q_s)$, we obtain

$$\begin{aligned} dT_s = {} & x(q_s)\, dp_s + \frac{1}{2} W''(q_s)\, d\langle q\rangle_s \\ & + (p_s - q_s)\left(W''(q_s)\, dq_s + \frac{1}{2} W'''(q_s)\, d\langle q\rangle_s\right) \\ & + W''(q_s)\, d\langle q, p - q\rangle_s. \end{aligned}$$

Now, whenever q_s varies, the synchronization rule implies

$$q_s = p_s.$$

Hence the terms multiplied by $(p_s - q_s)$ vanish identically. Moreover,

$$d\langle q, p - q\rangle_s = d\langle q, p\rangle_s - d\langle q\rangle_s.$$

Since q only varies when $q = p$, we have

$$d\langle q, p\rangle_s = d\langle q\rangle_s,$$

and therefore

$$d\langle q, p - q\rangle_s = 0.$$

Consequently,

$$dT_s = x(q_s)\, dp_s.$$

Integrating from 0 to u,

$$T_u = T_0 + \int_0^u x(q_s)\, dp_s.$$

On the other hand, the reserve-matching wealth process satisfies

$$R_u^L = R_0 + \int_0^u x(q_s)\, dp_s.$$

Therefore,

$$R_u^L - T_u = R_0 - T_0.$$

Since

$$T_0 = W(p_0),$$

it follows that

$$R_u^L - W(p_u) = R_0 - W(p_0) + T_u - W(p_u).$$

Finally, recall that

$$T_u \geq W(p_u).$$

Hence

$$R_u^L - W(p_u) \geq R_0 - W(p_0) \geq 0,$$

which implies

$$R_u^L \geq W(p_u).$$

Because $u \in [0, t]$ was arbitrary,

$$R_s^L \geq W(p_s), \qquad 0 \leq s \leq t.$$

Thus every admissible reserve-matching portfolio is a dominating portfolio on $[0, t]$. □

Theorem 4.9 shows that reserve-matching portfolios inherit the geometric advantage generated by the concavity of the liquidity-pool value function. Reserve synchronization causes the portfolio wealth to evolve along tangent lines to W, which remain above the marked-to-market value of the liquidity pool.

5 Benchmark Loss Measures

We now introduce benchmark loss measures for liquidity provision based on the reserve-matching framework developed in Section 4. The classical loss-versus-rebalancing benchmark corresponds to a single synchronization rule—continuous reserve matching. By contrast, Section 4 identified a broader class of admissible reserve-matching portfolios indexed by synchronization policies. This naturally leads to generalized benchmark notions obtained by optimizing over synchronization rules within the reserve-matching class.

Throughout this section, let

$$W_t = W(p_t)$$

denote the marked-to-market value of the liquidity pool, and for each admissible reserve-matching policy $L \in \mathcal{R}_t$, let

$$R_t^L = R_0 + \int_0^t x(p_{\tau(s)})\, dp_s$$

denote the associated reserve-matching wealth process.

5.1 Loss Versus Rebalancing (LVR)

We begin with the classical loss-versus-rebalancing benchmark introduced in [Milionis et al.(2022)]. The LVR benchmark compares the liquidity pool to a continuously synchronized benchmark portfolio that instantaneously updates its reserve exposure to match that of the pool.

Fix a finite horizon $t > 0$ and assume for the remainder of this section that the continuously synchronized reserve-matching policy is admissible on every finite horizon $[0, t]$. Accordingly, consider the continuously synchronized policy

$$\tau(s) = s, \qquad s \geq 0.$$

The associated reserve-matching portfolio holds

$$h_s = x(p_s),$$

and has wealth process

$$R_t^c = R_0 + \int_0^t x(p_s)\, dp_s.$$

Definition 5.1. *The loss-versus-rebalancing benchmark at time t is defined by*

$$LVR_t := R_t^c - W_t.$$

Thus LVR_t measures the benchmark gap of the liquidity pool relative to the continuously synchronized reserve-matching portfolio. Using Itô's formula together with the identity $W'(p) = x(p)$, we obtain

$$LVR_t = R_0 - W(p_0) - \frac{1}{2}\int_0^t W''(p_s)\, d\langle p\rangle_s.$$

In particular, since $W(\cdot)$ is concave, if $R_0 \geq W(p_0)$, then

$$LVR_t \geq 0.$$

5.2 Policy-Specific Benchmark Loss

Continuous synchronization is only one feasible synchronization rule. For any admissible reserve-matching policy $L \in \mathcal{R}_t$, we may instead compare the liquidity pool to the reserve-matching portfolio generated by that policy.

Definition 5.2. *For $L \in \mathcal{R}_t$, the loss versus reserve-matching policy L at horizon t is defined by*

$$LVR_t(L) := R_t^L - W_t.$$

By Theorem 4.9, if $R_0 \geq W(p_0)$, then

$$LVR_t(L) \geq 0$$

for every admissible reserve-matching policy $L \in \mathcal{R}_t$. The quantity $LVR_t(L)$ therefore measures how benchmark underperformance depends on the chosen synchronization rule.

5.3 Loss Versus Optimal Reserve-Matching (LVOR)

We now optimize over the reserve-matching class introduced in Section 4. Since continuous reserve matching is only one member of $\mathcal{R}_t$, this produces a stronger benchmark than the classical LVR measure.

Definition 5.3. *The loss-versus-optimal-reserve-matching benchmark at horizon t is defined by*

$$LVOR_t := \sup_{L \in \mathcal{R}_t} E[LVR_t(L)].$$

When the supremum above is finite, any policy $L^* \in \mathcal{R}_t$ attaining the supremum is called an optimal reserve-matching policy.

The expectations above are permitted to take the value $+\infty$, and therefore $LVOR_t$ may also equal $+\infty$.

The quantity $LVOR_t$ measures the largest expected benchmark gap of the liquidity pool relative to any admissible reserve-matching strategy evaluated at horizon t. Equivalently, $LVOR_t$ asks how poorly the pool performs, in expectation, when compared not merely to continuous synchronization, but to the best synchronization rule available within the reserve-matching class.

Remark 5.4. *(Dominating benchmarks). One may also consider the broader class $\mathcal{D}_t$ of admissible self-financing portfolios satisfying*

$$R_t \geq W_t$$

pathwise. Optimizing over $\mathcal{D}_t$ yields a more general domination benchmark. We focus on reserve-matching portfolios because they are economically tied to the AMM reserve rule and lead to tractable timing problems.

5.4 Relation Between LVR and LVOR

Because continuous reserve matching is admissible by assumption, the classical benchmark is nested within the generalized benchmark.

Proposition 5.5. *For every horizon t,*

$$LVOR_t \geq E[LVR_t].$$

Proof. The continuously synchronized reserve-matching policy

$$\tau(s) = s, \qquad 0 \leq s \leq t,$$

belongs to $\mathcal{R}_t$. Hence $E[LVR_t]$ is one feasible value in the supremum defining $LVOR_t$. The result follows immediately. □

5.5 Economic Interpretation

The distinction between LVR_t and $LVOR_t$ is economically meaningful. The classical benchmark LVR_t evaluates liquidity provision relative to an investor who continuously synchronizes reserve holdings with those of the liquidity pool. By contrast, $LVOR_t$ compares the liquidity pool to the best admissible reserve-matching strategy, allowing the benchmark investor to choose synchronization times strategically.

Accordingly, LVR_t captures losses generated purely by continuous reserve synchronization and the curvature of the liquidity-pool value function, while $LVOR_t$ additionally incorporates the value of synchronization timing. In particular, $LVOR_t$ measures how poorly the liquidity pool performs relative not merely to instantaneous reserve tracking, but to the most favorable synchronization rule available within the admissible reserve-matching class.

When the reference price is a martingale, synchronization timing often has no expected directional value, and under suitable integrability conditions the two notions coincide:

$$LVOR_t = E[LVR_t].$$

In this case, the benchmark losses arise entirely from volatility and convexity effects.

Under nonmartingale price dynamics, synchronization timing may strictly improve expected benchmark performance relative to continuous reserve matching. In such environments,

$$LVOR_t > E[LVR_t]$$

may hold, revealing benchmark effects that are invisible under the classical continuous-rebalancing metric.

This additional value may be interpreted as a form of timing optionality. A benchmark investor who is free to choose when reserves are synchronized can benefit from predictable price movements in ways unavailable to the continuously synchronized benchmark.

From the perspective of liquidity providers, the distinction is operationally important. If benchmark portfolios with superior timing performance exist, then the classical LVR benchmark may understate the magnitude of benchmark underperformance under nonmartingale price dynamics.

Section 5.6 formalizes these ideas by decomposing $LVOR_t$ into a volatility component inherited from classical LVR and a directional component generated by optimal synchronization timing.

5.6 Drift and Volatility Decomposition

The benchmark $LVOR_t$ contains two economically distinct sources of liquidity-provider underperformance: a volatility component generated by quadratic variation together with the curvature of the liquidity-pool value function, and a directional component arising from predictable movements in the reference price.

To make this distinction precise, suppose the reference price process admits the continuous semimartingale decomposition

$$dp_t = dM_t + dA_t,$$

where M is a local martingale and A is a continuous finite-variation process. Let $L \in \mathcal{R}_t$ be an admissible reserve-matching policy on $[0, t]$. Since

$$R_t^L = R_0 + \int_0^t x(p_{\tau(s)})\, dp_s,$$

we may write

$$R_t^L = R_0 + \int_0^t x(p_{\tau(s)})\, dM_s + \int_0^t x(p_{\tau(s)})\, dA_s.$$

The martingale term has zero expectation, and therefore

$$E[R_t^L] = R_0 + E\left[\int_0^t x(p_{\tau(s)})\, dA_s\right].$$

Subtracting the liquidity-pool value $W_t = W(p_t)$, and applying Itô's formula to $W(p_t)$, yields

$$E[LVR_t(L)] = R_0 - W(p_0) - \frac{1}{2}E\left[\int_0^t W''(p_s)\, d\langle M\rangle_s\right] + E\left[\int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dA_s\right].$$

This decomposition separates benchmark loss into two components.

(i) Volatility component:

$$-\frac{1}{2}E\left[\int_0^t W''(p_s)\, d\langle M\rangle_s\right].$$

Since W is concave, this term is nonnegative and corresponds to the classical loss-versus-rebalancing effect generated by volatility together with the curvature of the liquidity-pool value function.

(ii) Directional timing component:

$$E\left[\int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dA_s\right].$$

This term captures the benchmark effect generated by synchronization timing under predictable price movements.

For the continuously synchronized policy $\tau(s) = s$, the directional term vanishes, and therefore

$$E[LVR_t] = R_0 - W(p_0) - \frac{1}{2}E\left[\int_0^t W''(p_s)\, d\langle M\rangle_s\right].$$

Hence, for any admissible reserve-matching policy $L \in \mathcal{R}_t$,

$$E[LVR_t(L)] = E[LVR_t] + E\left[\int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dA_s\right].$$

Optimizing over $L \in \mathcal{R}_t$ yields

$$LVOR_t = E[LVR_t] + \sup_{L \in \mathcal{R}_t} E\left[\int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dA_s\right].$$

Thus $LVOR_t$ equals the classical expected LVR benchmark plus the maximal expected value generated by synchronization timing under the finite-variation component of the price process.

When the reference price is a martingale, one has $A \equiv 0$. Thus, since for any admissible policy the stochastic integral term is a true martingale, it follows that

$$LVOR_t = E[LVR_t].$$

In this case, all benchmark effects arise purely from volatility and convexity effects.

By contrast, when the price process contains a nontrivial finite-variation component, the directional timing component may be nonzero. In such environments, synchronization timing may affect expected benchmark performance relative to continuous reserve matching.

The decomposition therefore clarifies the economic content of $LVOR_t$: it extends classical LVR by incorporating an endogenous timing component generated by predictable price dynamics. Section 6 characterizes optimal synchronization rules under several classes of price processes.

6 Finite-Horizon Optimal Reserve-Matching Policies

We now study the finite-horizon LVOR optimization problem. Recall that for a finite horizon $t > 0$, the loss-versus-optimal-reserve-matching benchmark is defined by

$$LVOR_t = \sup_{L \in \mathcal{R}_t} E[LVR_t(L)].$$

Section 5 derived a decomposition of $LVOR_t$ separating volatility effects from synchronization timing effects. Specifically, the expected benchmark loss decomposes into a volatility component driven by the quadratic variation of the reference price process and the curvature of the liquidity-pool value function, together with a directional timing component generated by the finite-variation part of the reference price process.

This suggests viewing the finite-horizon LVOR optimization problem as a synchronization timing problem. Different reserve-matching policies correspond to different synchronization rules, and these rules may produce different expected benchmark performance depending on the dynamics of the reference price process.

When the reference price is a martingale, synchronization timing has no effect on expected benchmark performance. However, whenever the price process contains a nontrivial finite-variation component, synchronization timing may affect expected benchmark performance relative to continuous reserve matching.

The goal of this section is to characterize optimal reserve-matching policies over finite horizons under several important classes of price dynamics.

6.1 Martingale Prices

We first consider the case in which the reference price process p is a continuous martingale on $[0, t]$. That is,

$$E[p_s \mid \mathcal{F}_r] = p_r, \qquad 0 \le r \le s \le t.$$

In this environment, synchronization timing carries no expected directional value. Our first result shows that every admissible reserve-matching policy generates the same expected benchmark performance.

Proposition 6.1. *Suppose that p is a continuous martingale on $[0, t]$ and let $L \in \mathcal{R}_t$ be an admissible reserve-matching policy. Then*

$$E[\mathrm{LVR}_t(L)] = -\frac{1}{2} E\left[\int_0^t W''(p_s)\, d\langle p\rangle_s\right].$$

In particular, the expected benchmark loss is independent of the synchronization policy L, and therefore

$$\mathrm{LVOR}_t = E[\mathrm{LVR}_t].$$

Hence every admissible reserve-matching policy is optimal in expectation.

Proof. By the decomposition derived in Section 5.6,

$$\mathrm{LVR}_t(L) = \int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dp_s - \frac{1}{2}\int_0^t W''(p_s)\, d\langle p\rangle_s.$$

Since p is a continuous martingale and L and the continuously synchronized policy are both admissible, it follows that the stochastic integral is a true martingale. Consequently,

$$E\left[\int_0^t \big(x(p_{\tau(s)}) - x(p_s)\big)\, dp_s\right] = 0.$$

Taking expectations yields

$$E[\mathrm{LVR}_t(L)] = -\frac{1}{2} E\left[\int_0^t W''(p_s)\, d\langle p\rangle_s\right],$$

which is independent of L. Taking the supremum over admissible reserve-matching policies gives

$$\mathrm{LVOR}_t = E[\mathrm{LVR}_t].$$

□

Proposition 6.1 shows that under martingale dynamics, synchronization timing affects only the distribution of benchmark losses and not their expectation. All expected benchmark underperformance arises entirely from volatility and curvature effects rather than from timing opportunities.

6.2 Submartingale Prices

We now consider the case in which the reference price process p is a continuous submartingale on $[0, t]$. That is,

$$\mathbb{E}[p_s \mid \mathcal{F}_r] \geq p_r, \qquad 0 \leq r \leq s \leq t.$$

Equivalently, by the Doob–Meyer decomposition, we may write

$$p = M + A,$$

where M is a continuous martingale and A is a continuous nondecreasing adapted process. This corresponds to an environment in which the reference price exhibits nonnegative directional drift on average.

Recall that under a reserve-matching policy L, the risky holdings process is given by

$$h_s = x(p_{\tau(s)}),$$

where $\tau(s)$ is the most recent synchronization time relevant at date s. Since $x(\cdot)$ is decreasing, synchronizing at lower prices generates larger risky holdings, whereas synchronizing at higher prices generates smaller holdings. Under submartingale dynamics, future price increments have nonnegative conditional expectation, so higher expected benchmark wealth is achieved by carrying as large a risky position as possible. This suggests synchronizing at the lowest price observed so far.

Our main result confirms that intuition.

Proposition 6.2. *Suppose that p is a continuous submartingale on $[0, t]$. Assume that the running-minimum policy*

$$L^*(s) = p_0 - \inf_{0 \leq u \leq s} p_u$$

is admissible on $[0, t]$. Then L^* is optimal for the finite-horizon benchmark problem

$$LVOR_t = \sup_{L \in \mathcal{R}_t} E[LVR_t(L)].$$

Proof. By the decomposition in Section 5.6, for any admissible reserve-matching policy $L \in R_t$,

$$\mathbb{E}[LVR_t(L)] = \mathbb{E}[LVR_t] + \mathbb{E}\left[\int_0^t \left(x(p_{\tau(s)}) - x(p_s)\right) dA_s\right].$$

The first term on the right-hand side is independent of L. It therefore suffices to maximize

$$\mathbb{E}\left[\int_0^t \left(x(p_{\tau(s)}) - x(p_s)\right) dA_s\right].$$

Since A is nondecreasing and $x(\cdot)$ is decreasing, the integrand is maximized pointwise by minimizing $p_{\tau(s)}$ for each $s \in [0, t]$. This is achieved by synchronizing whenever the reference price reaches a new running minimum.

Equivalently, the optimal reserve-matching policy increases precisely when the reference price attains a new running minimum:

$$L^*(s) = p_0 - \min_{0 \le u \le s} p_u.$$

This policy maximizes the directional timing component and therefore attains the supremum defining $LVOR_t$. □

Remark 6.3. *The admissibility of the running-minimum policy is not automatic. It depends on the behavior of the reserve schedule $x(\cdot)$ near the boundary $p = 0$, together with the integrability properties of the reference price process. In particular, if $x(p)$ grows rapidly as $p \downarrow 0$, then the associated reserve-matching strategy*

$$h_s = x\left(\inf_{0 \le u \le s} p_u\right)$$

may fail to be admissible. A sufficient condition for admissibility is

$$E\left[\int_0^t x^2\left(\inf_{0 \le r \le s} p_r\right) d\langle M \rangle_s\right] < \infty.$$

Proposition 6.2 shows that under upward directional dynamics, delayed synchronization becomes economically valuable. Rather than continuously updating reserves, the benchmark investor benefits from synchronizing at historically low prices, thereby locking in larger risky holdings prior to subsequent upward movements.

This sharply contrasts with the martingale case, where synchronization timing has no effect on expected benchmark performance.

6.3 Supermartingale Prices

We finally consider the case in which the reference price process p is a continuous supermartingale on $[0, t]$. That is,

$$\mathbb{E}[p_s \mid \mathcal{F}_r] \le p_r, \qquad 0 \le r \le s \le t.$$

Equivalently, by the Doob–Meyer decomposition, we may write

$$p = M + A,$$

where M is a continuous martingale and A is a continuous nonincreasing adapted process. This corresponds to an environment in which the reference price exhibits nonpositive directional drift on average.

Recall that under a reserve-matching policy L, the risky holdings process is given by

$$h_s = x(p_{\tau(s)}),$$

where $\tau(s)$ is the most recent synchronization time relevant at date s. Since $x(\cdot)$ is decreasing, synchronizing at higher prices generates smaller risky holdings, whereas synchronizing at lower prices generates larger holdings. Under supermartingale dynamics, future price increments are nonpositive on average, so expected benchmark wealth is improved by carrying as small a risky position as possible. This suggests synchronizing at the highest price observed so far.

Our main result confirms that intuition.

Proposition 6.4. *Suppose that p is a continuous supermartingale on $[0,t]$. Assume that the running-maximum policy*

$$L^*(s) = \max_{0 \le u \le s} p_u - p_0, \qquad 0 \le s \le t,$$

is admissible on $[0,t]$. Then L^* is optimal for the finite-horizon benchmark problem

$$LVOR_t = \sup_{L \in \mathcal{R}_t} E[LVR_t(L)].$$

Proof. By the decomposition in Section 5.6, the optimization problem again reduces to maximizing the directional timing component. Since A is nonincreasing and $x(\cdot)$ is decreasing, the integrand is maximized pointwise by maximizing $p_{\tau(s)}$ for each $s \in [0,t]$. This is achieved by synchronizing whenever the reference price reaches a new running maximum, namely

$$L^*(s) = \max_{0 \le u \le s} p_u - p_0.$$

Therefore the running-maximum policy attains the supremum defining $LVOR_t$. □

Remark 6.5. *In contrast to the running-minimum policy of Proposition 6.2, the running-maximum* *policy naturally generates bounded risky exposure relative to its initial reserve level. Since*

$$\sup_{0 \le r \le s} p_r \ge p_0$$

and $x(\cdot)$ is decreasing,

$$x\left(\sup_{0 \le r \le s} p_r\right) \le x(p_0).$$

Consequently,

$$\int_0^t x^2\left(\sup_{0 \le r \le s} p_r\right) d\langle M \rangle_s \le x^2(p_0)\langle M \rangle_t.$$

Thus, if

$$E\left[x^2(p_0)\langle M \rangle_t\right] < \infty,$$

then the running-maximum policy satisfies the square-integrability condition for admissibility.

Proposition 6.4 is the mirror image of Proposition 6.2. Under upward drift, synchronization at running minima is optimal because it maximizes future exposure. Under downward drift, synchronization at running maxima is optimal because it minimizes future exposure.

Thus, when prices decline on average, patience again has positive expected value, but now in the opposite direction: the benchmark investor waits for high-price states, locks in smaller risky holdings, and thereby outperforms continuous reserve matching during subsequent downward movements.

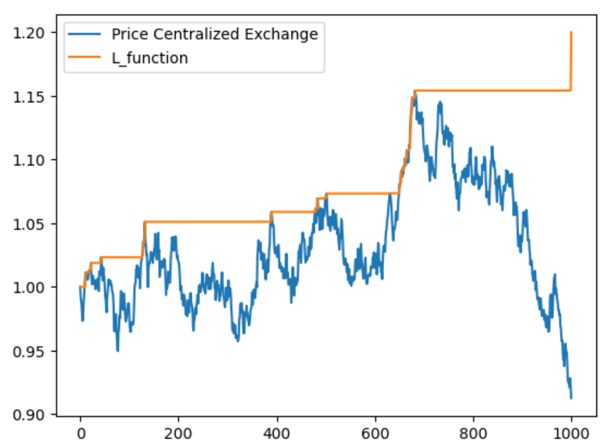

Figure 1: The reference price on the centralized exchange and its optimal L function in the supermartingale case.

Figure 1 provides a realization of a supermartingale reference price process along with its optimal reserve-matching policy L^*. Recall that the points of increase of L^* are the times when synchronization occurs.

7 Bubble Example

We now consider a class of strict local martingale price processes. These processes provide a canonical example in which synchronization timing may acquire positive expected value despite the absence of directional drift in the stochastic differential representation of prices. The example therefore illustrates that the distinction between local martingale and true martingale behavior is economically relevant for reserve-matching benchmarks.

7.1 Model

Suppose the reference price satisfies

$$dp_t = p_t^{\alpha} dB_t, \qquad p_0 > 0,$$

where $\alpha > 1$. The process p is a positive local martingale and therefore a supermartingale. For $\alpha > 1$, however, it fails to be a true martingale. Although the stochastic differential contains no explicit drift term, the process admits the Doob–Meyer decomposition

$$p_t = p_0 + N_t - A_t,$$

where N_t is a martingale and A_t is a nontrivial increasing predictable process. This class of diffusions is closely related to reciprocal Bessel processes and provides a canonical example of bubble-like price dynamics.

By contrast, when $\alpha \leq 1$, the diffusion remains a true martingale. In particular, $\alpha = 1$ corresponds to geometric Brownian motion, while for $\alpha < 1$ the origin becomes accessible and the process may hit zero in finite time.

7.2 Optimal Reserve Matching

Since p is a strict local martingale and therefore a supermartingale when $\alpha > 1$, Proposition 6.4 implies that the optimal finite-horizon reserve-matching policy synchronizes at running maxima:

$$L^*(s) = \max_{0 \leq u \leq s} p_u - p_0.$$

Equivalently, the benchmark portfolio updates its reserve exposure only when the reference price attains a new running maximum.

The notable feature of this example is that the stochastic differential representation of the price process contains no explicit drift term. Nevertheless, the failure of the local martingale to be a true martingale induces global supermartingale behavior through the compensator in the Doob–Meyer decomposition. Consequently, synchronization timing acquires positive expected value despite the absence of drift in the stochastic differential equation itself.

By delaying synchronization until new price highs are reached, the benchmark portfolio locks in smaller risky-asset exposure on average, thereby improving expected benchmark performance relative to continuous reserve matching.

7.3 Benchmark Implications

This example illustrates that liquidity pool benchmark performance depends not only on local price dynamics, but also on whether the reference price process is a true martingale. Because the stochastic differential representation

$$dp_t = p_t^\alpha dB_t$$

contains no explicit drift term, the process locally resembles a martingale diffusion. One might therefore expect the martingale benchmark results of Section 6.1 to apply.

Globally, however, the process is a positive strict local martingale and therefore a supermartingale when $\alpha > 1$. Consequently, synchronization timing acquires positive expected value even though the stochastic differential equation itself contains no drift term.

In particular, strict local martingale dynamics may generate strictly positive synchronization timing value:

$$LVOR_t > E[LVR_t].$$

Thus the generalized benchmark detects economically meaningful timing effects that are invisible under the classical continuously synchronized LVR benchmark.

The bubble example therefore demonstrates that benchmark performance in AMMs may depend on subtle global properties of price dynamics that are not visible from the local stochastic differential representation alone. More broadly, the example highlights the importance of distinguishing between local martingale and true martingale behavior when evaluating benchmark losses under intermittent reserve synchronization.

8 Empirics

To illustrate the economic magnitude of synchronization timing effects, we compare the realized performance of several reserve-matching benchmark policies using historical market data. The objective of this section is not to identify implementable trading strategies or statistically estimate structural price dynamics, but rather to illustrate empirically how synchronization timing may affect benchmark performance under different realized price paths.

We collected second-by-second price data for the ETH-USDC pair from the Binance centralized exchange over the period June 2024 through October 2024. We additionally collected reserve data for the Uniswap V2 ETH-USDC liquidity pool at each block.

To study different synchronization environments, we partitioned the data into nonoverlapping one-hour intervals and classified intervals according to their realized directional behavior. Intervals with positive cumulative returns were classified as upward-trending periods, while intervals with negative cumulative returns were classified as downward-trending periods. Intervals with comparatively small net price movements were treated as approximately martingale-like environments.

The purpose of this classification is not to statistically identify true submartingale or supermartingale dynamics, but rather to compare synchronization policies across intervals exhibiting different realized directional movements.

Throughout each interval, reserve levels were normalized using the initial liquidity level in order to isolate synchronization timing effects from changes in liquidity conditions over time. Throughout each interval, the reserve schedule $x(p)$ was held fixed at its initial level, while synchronized reserve states were evaluated along the realized price path. Using the observed price paths, we then evaluated the realized benchmark performance of several reserve-matching synchronization policies corresponding to the theoretical results of Section 6, together with a passive HODL benchmark:

- continuous synchronization,
- synchronization at running minima,
- synchronization at running maxima,
- a HODL portfolio holding the initial reserve allocation.

For each synchronization rule, we computed the realized benchmark gain relative to the liquidity pool over the interval. The resulting sample means, measured in USDC, are reported in Table 1.

	Cont.	Running Max.	Running Min.	HODL
supermartingale-like	1084.12	3379.26	889.17	3158.97
martingale-like	479.80	361.67	407.15	300.87
submartingale-like	1016.65	841.86	3196.01	3005.25

Table 1: Average benchmark gains relative to the liquidity pool over one-hour intervals.

The empirical results are broadly consistent with the theoretical predictions developed in Section 6. During upward-trending intervals, synchronization at running minima generates the largest benchmark gains, while during downward-trending intervals synchronization at running maxima performs best. By contrast, approximately martingale-like intervals exhibit substantially smaller dispersion across synchronization policies.

Figure 2 illustrates these effects for a representative interval classified as downward-trending. The upper panel plots cumulative benchmark gains relative to the liquidity pool under each synchronization rule. All synchronization policies generate positive benchmark gains relative to the liquidity pool, consistent with the pathwise domination property of reserve-matching portfolios.

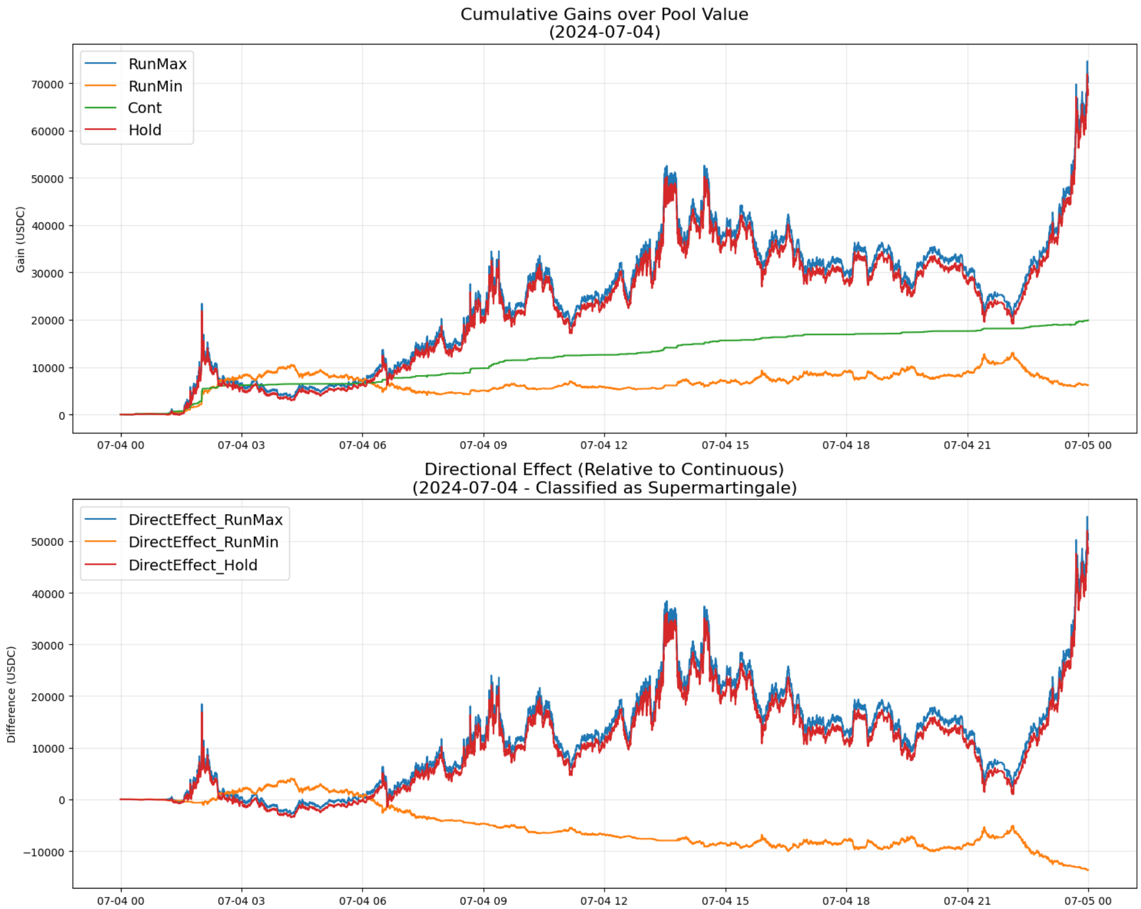

Figure 2: Benchmark gains under alternative reserve-matching policies during a downward-trending interval. The upper panel reports cumulative gains relative to the liquidity pool. The lower panel isolates the directional timing component by measuring gains relative to continuous synchronization.

The lower panel isolates the directional timing component by plotting benchmark gains relative to the continuously synchronized portfolio. Consistent with Proposition 6.4, the running-maximum policy generates substantial additional benchmark gains during the downward-trending period, while the running-minimum policy underperforms continuous synchronization.

Overall, the empirical results support the central prediction of the LVOR framework: synchronization timing becomes economically important when realized price paths exhibit directional behavior. In contrast, periods with limited directional movement display comparatively small differences across synchronization policies, consistent with the martingale benchmark results of Section 6.1.

The empirical exercise is intentionally stylized and abstracts from transaction costs, execution frictions, endogenous arbitrage interactions, and equilibrium effects. Nevertheless, the results suggest that the timing component identified by the LVOR framework may be quantitatively significant in realized market environments.

9 Conclusion

This paper introduces a general reserve-matching framework for studying benchmark performance in automated market makers under intermittent synchronization. Classical loss-versus-rebalancing (LVR) compares a liquidity pool to a continuously synchronized benchmark portfolio. We generalized this benchmark by introducing a broad class of reserve-matching portfolios indexed by synchronization policies and by defining the associated loss-versus-optimal-rebalancing (LVOR) benchmark.

A central insight of the paper is that synchronization timing may itself generate economically meaningful benchmark effects. The resulting decomposition separates benchmark underperformance into two components: a volatility component generated by the curvature of the liquidity-pool value function together with quadratic variation in the reference price process, and a directional timing component generated by synchronization under nonmartingale price dynamics. This decomposition clarifies that classical LVR captures only part of the benchmark effects associated with liquidity provision.

The reserve-matching framework also leads naturally to a synchronization timing problem. When reference prices are martingales, synchronization timing has no expected value and all admissible reserve-matching policies achieve the same expected benchmark performance. Under submartingale and supermartingale dynamics, however, synchronization timing becomes economically valuable, and optimal policies are characterized by synchronization at running minima or running maxima of the reference price process.

More broadly, the paper highlights the importance of global properties of price dynamics for AMM benchmark analysis. The bubble example shows that strict local martingale behavior may generate synchronization timing value even when the stochastic differential representation of prices contains no explicit drift term. Thus benchmark performance may depend not only on local volatility structure, but also on whether the reference price process is a true martingale.

The framework developed here suggests several directions for future research. One important extension is to study synchronization timing in settings where arbitrage activity and reserve updates are jointly endogenous rather than exogenously specified through benchmark policies. Other directions include equilibrium interactions among multiple agents, transaction costs and execution frictions.

Overall, the paper suggests that benchmark losses in AMMs are shaped not only by volatility and finite liquidity, but also by the timing structure through which reserves are synchronized with external prices.

References

[Adams et al.(2025)] Austin Adams, Ciamac C Moallemi, Sara Reynolds, and Dan Robinson. 2025. am-amm: An auction-managed automated market maker. In *International Conference on Financial Cryptography and Data Security*. Springer, 93–108.

[Angeris et al.(2019)] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2019. An analysis of Uniswap markets. *Cryptoeconomic Systems* 1, 1 (2019). Article 3. arXiv:1911.03380.

[Baggiani et al.(2025)] Leonardo Baggiani, Martin Herdegen, and Leandro Sánchez-Betancourt. 2025. Optimal dynamic fees in automated market makers. *arXiv preprint arXiv:2506.02869* (2025).

[Capponi and Zhu(2024)] Agostino Capponi and Brian Zhu. 2024. Optimal Exiting for Liquidity Provision in Constant Function Market Makers. (2024). Working paper.

[Cartea et al.(2023)] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. 2023. Predictable losses of liquidity provision in constant function markets and concentrated liquidity markets. *Applied Mathematical Finance* 30, 2 (2023), 69–93.

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

[Deng et al.(2023)] Jun Deng, Hua Zong, and Yun Wang. 2023. Static replication of impermanent loss for concentrated liquidity provision in decentralised markets. *Operations Research Letters* 51, 3 (2023), 206–211.

[Drossos et al.(2025)] Thanos Drossos, Daniel Kirste, Niclas Kannengießer, and Ali Sunyaev. 2025. Automated Market Makers: Toward More Profitable Liquidity Provisioning Strategies. In *The 40th ACM/SIGAPP Symposium on Applied Computing (SAC '25)*. ACM.

[Fukasawa et al.(2023)] Masaaki Fukasawa, Basile Maire, and Marcus Wunsch. 2023. Weighted variance swaps hedge against impermanent loss. *Quantitative Finance* 23, 6 (2023), 901–911.

[Goyal et al.(2023)] Mohak Goyal, Geoffrey Ramseyer, Ashish Goel, and David Mazières. 2023. Finding the Right Curve: Optimal Design of Constant Function Market Makers. In *EC '23: Proceedings of the 24th ACM Conference on Economics and Computation*. 783–812.

[Hasbrouck et al.(2026)] Joel Hasbrouck, Thomas J Rivera, and Fahad Saleh. 2026. The need for fees at a dex: How increases in fees can increase dex trading volume. *Management Science* (2026).

[Loesch et al.(2021)] Stefan Loesch, Nate Hindman, Mark B. Richardson, and Nicholas Welch. 2021. *Impermanent Loss in Uniswap v3*. Technical Report. Bancor. arXiv:2111.09192.

[Maire and Wunsch(2024)] Basile Maire and Marcus Wunsch. 2024. Market neutral liquidity provision. Preprint.

[Milionis et al.(2022)] Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046* (2022).

[Milionis et al.(2023)] Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2023. Loss-versus-rebalancing under fee models with Poisson block times. arXiv:2305.14604.

[Nezlobin and Tassy(2025)] Alexander Nezlobin and Martin Tassy. 2025. Loss-Versus-Rebalancing Under Deterministic and Generalized Block-Times. arXiv:2505.05113.

[Pintail(2019)] Pintail. 2019. Uniswap: A good deal for liquidity providers? https://pintail.medium.com/uniswap-a-good-deal-for-liquidity-providers-104c0b6816f2