

Adaptive Curves for Optimally Efficient Market Making

Viraj Nadkarni   

Princeton University, NJ, USA

Sanjeev Kulkarni   

Princeton University, NJ, USA

Pramod Viswanath   

Princeton University, NJ, USA

Abstract

Automated Market Makers (AMMs) are essential in Decentralized Finance (DeFi) as they match liquidity supply with demand. They function through liquidity providers (LPs) who deposit assets into liquidity pools. However, the asset trading prices in these pools often trail behind those in more dynamic, centralized exchanges, leading to potential arbitrage losses for LPs. This issue is tackled by adapting market maker bonding curves to trader behavior, based on the classical market microstructure model of Glosten and Milgrom. Our approach ensures a zero-profit condition for the market maker's prices. We derive the differential equation that an optimal adaptive curve should follow to minimize arbitrage losses while remaining competitive. Solutions to this optimality equation are obtained for standard Gaussian and Lognormal price models using Kalman filtering. A key feature of our method is its ability to estimate the external market price without relying on price or loss oracles. We also provide an equivalent differential equation for the implied dynamics of canonical static bonding curves and establish conditions for their optimality. Our algorithms demonstrate robustness to changing market conditions and adversarial perturbations, and we offer an on-chain implementation using Uniswap v4 alongside off-chain AI co-processors.

2012 ACM Subject Classification Theory of computation → Market equilibria

Keywords and phrases Automated market makers, Adaptive, Glosten-Milgrom, Decentralized Finance

Digital Object Identifier 10.4230/LIPIcs.AFT.2024.25

Related Version *Full Version:* <https://arxiv.org/pdf/2406.13794> [47]

Funding This work was supported by the National Science Foundation via grants CNS-2325477, the Army Research Office via grant W911NF2310147, a grant from C3.AI, a SEAS Innovation grant from Princeton University and a gift from XinFin Private Limited.

1 Introduction

Market making plays a crucial role in enhancing liquidity in financial systems. In traditional finance, effective market making strategies involve setting buy and sell prices (bid and ask quotes) as narrowly separated as possible, ensuring these quotes closely mirror the asset's true price on a limit order book. This strategy enables market makers to earn a marginal profit. Such efficacy in market making is driven by sophisticated models that analyze trader behavior [25, 34, 27]. These models have become fundamental in understanding the principles of microeconomics and the microstructure of markets.

Market Makers in DeFi. In the field of Decentralized Finance (DeFi), the concept of *automated* market making has gained prominence. DeFi employs Automated Market Makers (AMMs), particularly Constant Function Market Makers (CFMMs) [45][2], offering an



© Viraj Nadkarni, Sanjeev Kulkarni, and Pramod Viswanath;
licensed under Creative Commons License CC-BY 4.0

6th Conference on Advances in Financial Technologies (AFT 2024).

Editors: Rainer Böhme and Lucianna Kiffer; Article No. 25; pp. 25:1–25:22

Leibniz International Proceedings in Informatics



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

25:2 Adaptive Curves for Optimally Efficient Market Making

alternative to traditional limit order books. This approach reduces the computational effort needed to facilitate trades and ensures liquidity is available for tokens that are less frequently traded. Unlike conventional markets that primarily utilize limit order books for Peer-to-Peer transactions, decentralized markets implement a Peer-to-Pool-to-Peer structure. In this model, Liquidity Providers (LPs) aggregate their resources in a contract, which traders then utilize to meet their liquidity needs. Thus, any DeFi market needs to incentivize both the LPs and the traders to ensure a fair and efficient market is created.

Market depth and volatility. DeFi markets exhibit a range of distinct characteristics, primarily differentiated by market depth (or liquidity) and price volatility. Notably, the trading volume of stablecoins (approximately \$11.1 trillion) has recently exceeded the transaction volumes of centralized entities like MasterCard and PayPal [31]. Markets with significant liquidity, especially those trading stablecoins, are highlighted in this context [35]. Such markets typically experience minimal volatility, and their substantial depth minimizes the price impact of retail trades. Conversely, DeFi features hundreds of infrequently traded tokens, which suffer from a lack of liquidity, leading to high volatility and price sensitivity to even small-scale retail trades. These markets are also susceptible to swings in price caused by flash loan transactions [59]. This paper addresses the optimization of market making strategies for the latter (less liquid) kind of markets.

Incentives of LPs. A key challenge for Constant Function Market Makers (CFMMs) is motivating Liquidity Providers (LPs) to contribute their tokens to the pool. For this incentive to work, it is crucial for CFMMs to minimize the average losses on pooled assets. Yet, it is widely acknowledged that LPs often incur losses due to fluctuations in reserves [36] and a lack of market insights [43]. This paper concentrates on reducing the losses that stem from such informational deficiencies. Specifically, static curves in CFMMs frequently lead to LP losses as a result of arbitrage activities. These losses are intended to be offset by transaction fees, contrasting with centralized exchanges which benefit from higher liquidity and trading volumes but impose lower fees. For example, Binance, a centralized exchange, records a daily trading volume of approximately \$15 billion, significantly higher than Uniswap's \$1.1 billion [9], the largest decentralized exchange. The lower liquidity on platforms like Uniswap results in less current prices, making them more susceptible to arbitrage losses.

Arbitrage Loss. The specific type of arbitrage loss known as loss-versus-rebalancing (LVR) can be quantified in certain scenarios [43], and these losses continue to occur despite the implementation of trading fees [41]. In the case of a generic market maker who sets bid and ask prices for a volatile asset, arbitrage losses are defined relative to the asset's true market price. An arbitrageur engages in a buy transaction when the market price surpasses the ask price, and in a sell transaction when it drops below the bid price. The resultant loss for the market maker is calculated as the product of the price difference and the volume of the asset traded.

Trader behavior. In traditional financial systems, arbitrage-related losses are conceptualized as *adverse selection* costs, which arise from interactions with *informed traders* – those who are privy to the external market price, akin to arbitrageurs. A market maker achieves optimal operation by balancing these costs against the profits gained from *uninformed traders*, also known as *noise traders*. This balancing principle was first delineated by Glosten and Milgrom [25]. Within the Decentralized Finance (DeFi) ecosystem, trading parties

are differentiated into *toxic* and *non-toxic* order flows, which correspond to informed and uninformed traders, respectively [48, 15]. We extend this model to a more nuanced framework where traders are categorized along a continuous spectrum of information awareness, ranging from highly informed to completely uninformed, rather than being strictly classified as either *toxic* or *non-toxic* (Section 3).

CFMMs as prediction markets. For CFMMs, a significant portion of their losses also arises from the need to encourage traders to disclose their genuine price perceptions during transactions. In essence, CFMMs provide compensation to informed traders in exchange for their crucial market insights, which mirrors the principles of market scoring rules utilized in prediction markets to extract valuable information [24].

Conditions for optimality. A straightforward approach to reduce losses to arbitrage would be aligning the marginal price exactly with the external market price, which would require real-time data from a *price oracle* [16]. Yet, integrating oracles with market making strategies can lead to potential frontrunning risks [37] and necessitates reliance on centralized, potentially manipulable external entities [21, 67]. To circumvent these issues, our framework explicitly excludes the use of oracles. The objective is to deduce the hidden market price by analyzing trade history data, aiming for maximum efficiency in terms of data utilization. Further, the market maker uses this to adaptively set its bonding curve so that the loss to arbitrageurs is as close to zero as possible, ensuring an optimally efficient market. The market maker turning a profit would be undesirable since this would allow a competitor to undercut its prices and take away their order flow. In other words, it should quote an efficient and competitive market price, given only the information it has in form of the trading history. Keeping this objective in mind, we outline the key contributions of this work.

Our contributions

Optimal algorithms for adapting curves. (Section 4) We provide the differential equation that the demand curve of an optimally efficient market should follow (Theorem 1). When the statistics governing trader and price behavior are known and Gaussian/Lognormal, we show that this differential equation can be solved exactly using a dynamic bonding curve that changes its operating point using the Kalman Filter (Theorem 2 and Theorem 3).

Adapting to unknown market conditions. When the statistics governing trader and price behavior are unknown, we extend the previous approach by using an Adaptive Kalman Filter. We empirically show that both these approaches suffer significantly lower arbitrage losses compared to a static CFMM. (Section 5.2)

Robustness to adversarial manipulation. In presence of irrational traders that seek to make the price of the market maker deviate from the external market, we present a robust version of the adaptive curve algorithm that tolerates upto 50% of trader population being adversarial. (Section 5.3)

Comparisons with static curves. (Section 6) We provide theoretical comparisons of static and the proposed adaptive curve models by showing that the error in the adaptive AMM price, when viewed as an oracle, decays with more trades, while that of a static curve remains unchanged (Theorem 4).

Implied dynamics of static curves. (Section 6) We derive a differential equation (30) that governs the *implied* dynamical model given the static CFMM curves that are used in practice. We show that these CFMMs can only be optimally efficient in a model where inter-block time for the underlying blockchain vanishes, the CFMM has more liquidity than the external market, and that the trader and price behaviour is severely constrained. (Theorem 5)

On-chain Implementation. We specify the end-to-end system design for our proposed market maker. Furthermore, we provide an implementation of our algorithm using the recently released Uniswap v4 [57] platform and an off-chain machine learning co-processor Axiom [51]. This co-processor provides guarantees that the algorithms derived in this work are executed, and the result is put securely on-chain. (Section 8)

2 Related work

In this section, we reprise relevant literature surrounding the problem formulated in this paper. Although the motivation of the problem stems from literature studying AMMs in DeFi, our formulation derives heavily from classical works in market microstructure. The algorithms that we present in our work derive heavily from literature on control and robust filtering.

Automated Market Makers. Automated Market Makers (AMMs), particularly in the form of Constant Function Market Makers (CFMMs) [66, 45], are designed to incentivize trades that align prices with a more liquid external market [3]. However, this mechanism imposes costs on CFMM liquidity providers while generating profits for arbitrageurs [23, 29, 61]. This arbitrage profit, often quantified as “loss-versus-rebalancing” in scenarios where only arbitrageurs (informed traders) interact with the market maker, is proportional to external price volatility [43]. Various methods have been proposed to capture this loss, including on-chain auctions [39] and the application of auction theory to dynamically recommend ask and bid prices for an AMM [42]. Another recent study [26] suggests an optimal curve for a CFMM based on liquidity providers’ price beliefs. However, this study does not consider a dynamic model where traders react to market maker price settings. In [42], a dynamic trading model is examined, deriving optimal ask and bid prices. Yet, this approach necessitates the market maker’s knowledge of underlying model parameters, limiting adaptability to market conditions. Additionally, while [42] focuses on a monopolistic market maker, our work examines a *competitive* market maker. Our research aligns closely with [46], which addresses competitive market conditions but only within a liquid, non-volatile market with restricted trade sizes. Reinforcement learning algorithms to adapt CFMM bonding curves have been explored in [14], though the primary objectives there are fee revenue control and minimization of failed trades.

Optimal market making. The trader behavior model we employ is derived from the Glosten-Milgrom model [25], which is widely used in market microstructure literature. However, we modify it to incorporate a continuously changing external price. Several subsequent studies [17, 18] develop optimal market-making rules within a modified Glosten-Milgrom framework, but they assume that the underlying model parameters are known and that external price changes are communicated to the market maker. A more data-driven reinforcement learning approach is taken in [12], but their reward function presupposes direct information about the external hidden price, whereas we assume no access to a price oracle. Another aspect of

optimal market making in traditional market microstructure literature focuses on inventory management [30, 7], rather than the information asymmetry between traders and market makers. Our goal is to design a market maker that mitigates losses due to information asymmetry, similar to the Glosten-Milgrom model, without imposing inventory constraints. The Glosten-Milgrom model has been considered for AMMs in DeFi, though only for individual trades [5, 4].

Optimal filtering and control. The classical filtering and control literature underpins many contemporary automated systems, offering theoretical foundations and practical applications for dynamic system regulation. Kalman filtering, a robust statistical method, is widely used for estimating the state of a linear dynamic system from incomplete and noisy measurements [32]. This technique combines real-time measurements with prior estimates to produce updated predictions, proving crucial for systems requiring high accuracy and responsiveness, such as navigation systems, aerospace engineering, and automated trading. In control theory, concepts like optimal control and feedback mechanisms are fundamental for designing systems that maintain desired output levels despite external condition changes. These principles have been extensively explored in works such as [65], and further developed through modern control theories addressing non-linearities and uncertainties in system dynamics [68]. The integration of filtering and control methodologies has led to the development of Adaptive Kalman Filtering [40], which adjusts its parameters based on observed errors, enhancing performance in varying conditions. Adaptive Kalman Filtering has seen applications in diverse areas, including robotics [63], automotive systems [28], and finance [20], illustrating its versatility and robustness in handling dynamic, uncertain environments.

3 Preliminaries and model

We now describe the framework used for modeling trader behavior in response to the evolution of an external price process, that is hidden from the market maker, and the prices set by the market maker. We also state the objective that the market maker seeks to optimize, and provide the motivation behind it. The model and the objective are based on the canonical Glosten-Milgrom model [25] studied extensively in market microstructure literature. We assume that the market maker has access to an inventory of the asset and numeraire. The price of the asset is expressed in terms of the numeraire. We assume that time is discrete and indexed by t .

External price process. The external price process p_{ext}^t of a risky asset is assumed to follow a discrete time random walk, where the distribution of a price jump at any t is parametrized by σ . That is, we have

$$p_{ext}^{t+1} = p_{ext}^t + \Delta p_{ext}^t \quad (1)$$

where Δp_{ext}^t is i.i.d. $\sim \mathcal{D}_{\sigma, p_{ext}^t}$. Intuitively, the parameter σ can be thought of as a measure of volatility of the external price. Further, the random process p_{ext}^t can either represent the price of the asset in a larger and much more liquid exchange, or some underlying “true” value of the asset. In both cases, we assume that it is hidden from the market maker. To put it in DeFi terms, the market maker does not have an access to any “price oracle” that can tell it information about this external price.

Market Maker. The market maker consists of a pool containing an asset and a numeraire. It then publishes a demand curve $g_t(p)$ [42] before each trade. All traders have full access to the demand curve at any time. The demand curve specifies how the price of the asset changes based on the inventory available. In particular, $g(p_0^t)$ specifies the amount of asset when the initial operating point of the demand curve is p_0^t . That is, the price of an infinitesimal amount dx of asset is $p_0^t dx$. The amount of numeraire in the inventory at the same operating point is $-\int_0^{p_0^t} p dg(p)$. For an incentive compatible market maker, we constrain $g(\cdot)$ to be non-increasing (Proposition 2.1 in [42]). Equivalently, the market maker may be represented by the canonical *bonding curve* $\phi_t(x, y)$ [66] which is a function (of reserves x, y of the asset, numeraire respectively) that stays constant over any single trade.

Trader behavior. We assume that a trader appears at every time step t , and sees a noisy version of the external price in each time step, denoted by p_{trad}^t , with the noise distribution being parametrized by η .

$$p_{trad}^t = p_{ext}^t + \Delta p_{trad}^t \quad (2)$$

where Δp_{trad}^t is i.i.d. noise $\sim \mathcal{D}_{\eta, p_{ext}^t}$ sampled at each time step. Intuitively, the parameter η can be construed as measuring the level of “toxicity” or informed nature of the traders.

Trade actions. As mentioned before, the market maker publishes a demand curve $g_t(p)$ at each time step. The trader performs a trade that brings the operating point of the demand curve from p_0^t to p_{trad}^t . Thus, the trader behaves as if it is performing arbitrage between a market with price p_{trad}^t and the demand curve published. For instance, if the operating point p_0^t of the market maker is less than the trader price p_{trad}^t , then a trader would buy asset from the market maker to sell on the external market until the operating point shifts to p_{trad}^t .

Objective. Our objective is to design an algorithm to set ask and bid prices for the market maker, such that the expected loss with respect to the external market is minimized and the market maker stays competitive. Taking inspiration from [25], but extending to the modified trader behaviour, we get the following condition

$$p^t = E[p_{ext}^t | \mathcal{H}_{t-1}, p_{trad}^t] \quad \forall p_{trad}^t \quad (3)$$

where $\mathcal{H}_{t-1} = \langle (p_{trad}^\tau, g_\tau(\cdot)) \rangle_{\tau=0}^{t-1}$ is the history of trades and demand curves until time $t-1$, and p^t is the net price of the trade (the ratio of change in the reserves of the numeraire with the change in the reserves of the asset).

Interpreting the objective. Setting the prices as per the above objective makes the expected loss of the market makers to traders vanish, since

$$E[(p_{ext}^t - p^t) | \mathcal{H}_{t-1}] = 0 \quad (4)$$

Note that the expression above quantifies the expected loss for per unit trade of the asset.

The market maker can obtain a strictly positive profit by choosing a curve with slightly higher prices (than the ones obtained in (3)) on the ask side or slightly lower prices on the bid side (in other words, by choosing a CFMM with a higher curvature [4]). However, doing this would make it less competitive, since any other market maker with slightly greater bid or a slightly lesser ask would offer a better price and take away the trade volume. Although

we do not explicitly model other market makers, their presence is implicit in setting prices according to (3). This equation represents the ideal conditions for capital efficiency, where both the trader gets the best price possible while the market maker avoids a loss.

Also, note that the market maker incurs a loss in *every trade* made by an perfectly informed trader. Thus, to make the expected loss vanish, it should learn to set prices so that the loss to more informed traders is balanced by the profit obtained from less informed traders. The equation (3) can also be interpreted as striking this balance.

4 Differential Equation for the optimal curve

As mentioned in Section 3, we make use of the general *demand curve* formulation for market makers. Assume that the amount of asset in reserves at any price p_0^t is $g_t(p_0^t)$, while the amount of numeraire is $-\int_{p_0^t}^{p^t} p dg_t(p)$, where the function g_t needs to be non-increasing for the AMM to be incentive compatible [42]. Here, p_0^t is the initial operating point of the AMM at the beginning of the time slot t . Let $f_t(p)$ be the belief (probability distribution function) of the AMM over the price p of the asset at time t . Let $f_\eta(p)$ be the distribution of noise through which the price is observed by the traders. Let p_{trad}^t be the price that the trader observes. Then, the trader would make a trade with the AMM such that the marginal price of the asset just after the trade is p_{trad}^t . We now solve the optimal market conditions given by (3).

In this case, the effective price of a trade at time t is given by

$$p^t = \frac{-\int_{p_0^t}^{p_{trad}^t} p dg_t(p)}{g_t(p_0^t) - g_t(p_{trad}^t)}, \quad (5)$$

where p_0^t is the operating point of the AMM just before the trade. Further, by definition of conditional expectation of the external market price, we also have

$$E[p_{ext}^t | \mathcal{H}_{t-1}, p_{trad}^t] = \frac{\int_0^\infty p f_\eta(p_{trad}^t - p) f_t(p) dp}{\int_0^\infty f_\eta(p_{trad}^t - p) f_t(p) dp}. \quad (6)$$

since the noise at time t is independent of the external price at time t . Substituting (6) and (5) in (3) gives us

$$\frac{-\int_{p_0^t}^{p_{trad}^t} p dg_t(p)}{g_t(p_0^t) - g_t(p_{trad}^t)} = \frac{\int_0^\infty p f_\eta(p_{trad}^t - p) f_t(p) dp}{\int_0^\infty f_\eta(p_{trad}^t - p) f_t(p) dp} \quad \forall p_{trad}^t \quad (7)$$

Rearranging gives us

$$g_t(p_{trad}^t) = g_t(p_0^t) + \frac{\int_{p_0^t}^{p_{trad}^t} p g_t'(p) dp}{\beta_t(p_{trad}^t)} \quad \forall p_{trad}^t \quad (8)$$

where $\beta_t(p_{trad}^t) = \frac{\int_0^\infty p f_\eta(p_{trad}^t - p) f_t(p) dp}{\int_0^\infty f_\eta(p_{trad}^t - p) f_t(p) dp}$. On the other hand, if we let $p_{trad}^t \rightarrow p_0^t$ in (7), we get

$$p_0^t = \beta_t(p_0^t) \quad (9)$$

Thus, using (8) and (9), the demand curve of the optimal market maker can be given by the following theorem.

► **Theorem 1.** *Let the amounts of asset and numeraire in the reserves of a market maker be x_0^t, y_0^t . Then, the optimal demand curve $g_t(p)$ for a market maker obeys the following differential equation*

$$(\beta_t(p) - p)g_t'(p) + \beta_t'(p)g_t(p) - \beta_t'(p)x_0^t = 0 \quad (10)$$

with the constraints $\lim_{\delta \rightarrow 0^-} g(p_0^t + \delta) \geq x_0^t$ and $\lim_{\delta \rightarrow 0^-} -\int_0^{p_0^t + \delta} p dg(p) \geq y_0^t$. This equation can be solved separately for $p > p_0^t$ and $p < p_0^t$. Furthermore, the initial operating point p_0^t is a solution to the fixed point equation

$$p = \beta_t(p). \quad (11)$$

The differential equation (10) is obtained by simply applying the Leibniz rule of the derivative of a definite integral on (8). Note that we can rewrite (8) as

$$\int_{p_0^t}^{p_{trad}^t} \left(1 - \frac{p}{\beta_t(p_{trad}^t)}\right) dg_t(p) = 0 \quad (12)$$

We see that a discontinuity at p_0^t implies that $dg_t(p)$ is not a differential but is a negative real number. Since $p = \beta_t(p)$ at p_0^t , we have the term inside the integral $\rightarrow 0$ as $p_{trad}^t \rightarrow p_0^t$. This implies that a discontinuity in $g_t(p)$ is allowed at p_0^t . Further, since the (12) holds for $p_{trad}^t > p_0^t$ and for $p_{trad}^t < p_0^t$, the differential equation can be solved separately for $p > p_0^t$ and $p < p_0^t$.

Finally, assuming that the AMM finds solutions to (10) and (11), the trader is free to move the market to p_{trad}^t . The market maker then simply uses Bayes' rule to update its beliefs over the external prices as

$$f_{t+1}(p) = \frac{f_\eta(p_{trad}^t - p)f_t(p)}{\int_0^\infty f_\eta(p_{trad}^t - p)f_t(p)dp} \quad (13)$$

The updated belief $f_{t+1}(p)$ is now used to compute $\beta_{t+1}(p)$, thus completing the market making algorithm. Equations (10), (11) and (13) describe the complete dynamics of an optimally efficient market maker.

5 Optimal Solutions for special cases

In this section, we present solutions to the differential equation derived in Theorem 1 for some special cases: Gaussian and Lognormal price jumps and trader noises. That is, we assume that both Δp_{ext}^t and Δp_{trad}^t follow a Gaussian or Lognormal distribution. Under these assumptions, we get the following key results.

5.1 Kalman Filter algorithm for known market parameters

► **Theorem 2.** *If $\Delta p_{ext}^t \sim \mathcal{N}(0, \sigma^2)$ and $\Delta p_{trad}^t \sim \mathcal{N}(0, \eta^2)$ where σ, η are known to the market maker, then the fixed point equation $p_0^t = \beta_t(p_0^t)$ (11) has a unique solution given by the Kalman filter [32] estimate of p_{ext}^t , that is, we have*

$$p_0^t = E[p_{ext}^t | \mathcal{H}_{t-1}]. \quad (14)$$

Further, the differential equation (10) has a family of solutions, given by

$$g_t(p) = \begin{cases} x_0^t + \frac{y_0^t}{p_0^t} & \text{if } p \leq p_0^t \\ \max(0, \tilde{x}_0^t - C_t(p - p_0^t)^{\frac{K_t}{1-K_t}}) & \text{if } p > p_0^t \end{cases} \quad (15)$$

where K_t is the Kalman gain, and $C_t, \tilde{x}_0^t$ are non-negative constants, such that $\tilde{x}_0^t \leq x_0^t$.

■ **Algorithm 1** A Kalman filter based algorithm to adapt the AMM curve.

Require: Known η, σ , Reserves x_0, y_0

```

1:  $t \leftarrow 0$ 
2:  $T \leftarrow$  Number of total time slots
3: Initial price estimate  $p_{ext}^{0|0} \leftarrow p_{ext}^0$ 
4:  $P_{0|0} = 0$ 
5: while  $t \leq T$  do
6:    $\theta_t \leftarrow \frac{p_{ext}^{t-1|t-1} x_{t-1}}{p_{ext}^{t-1|t-1} x_{t-1} + y_{t-1}}$ 
7:   Publish bonding curve  $x^{\theta_t} y^{1-\theta_t} = x_{t-1}^{\theta_t} y_{t-1}^{1-\theta_t}$ 
8:   Observe trader action  $p_{trad}^t$ 
9:    $K_t \leftarrow \frac{P_{t-1|t-1} + \sigma^2}{P_{t-1|t-1} + \sigma^2 + \eta^2}$  ▷ Update Kalman gain
10:   $p_{ext}^{t|t} \leftarrow (1 - K_t)p_{ext}^{t-1|t-1} + K_t p_{trad}^t$  ▷ Update Kalman estimate of the external price
11:   $P_{t|t} = (1 - K_t)(P_{t-1|t-1} + \sigma^2)$  ▷ Update Kalman uncertainty
12: end while

```

The above solution corresponds to a CFMM with the initial slope of the bonding curve given by the Kalman estimate of p_{ext}^t by treating p_{trad}^t as noisy observations. Thus, we only need to calculate a single quantity $E[p_{ext}^t | \mathcal{H}_t]$ (see Algorithm 1), and set the other parameters of the demand curve according to (15). We get a similar solution if we assume that the external price follows geometric Brownian motion.

► **Theorem 3.** *If $\log \frac{p_{ext}^t}{p_{ext}^{t-1}} \sim \mathcal{N}(0, \sigma^2)$ and $\log \frac{p_{trad}^t}{p_{ext}^t} \sim \mathcal{N}(0, \eta^2)$ where σ, η are known to the market maker, then the fixed point equation $p_0^t = \beta_t(p_0^t)$ (11) has a unique solution which is a function of the Kalman filter [32] estimate of $\log p_{ext}^t$, that is, we have*

$$p_0^t = \exp \left(E[\log p_{ext}^t | \mathcal{H}_{t-1}] + \frac{P_{t|t}}{2(1 - K_t)} \right), \quad (16)$$

where K_t is the Kalman gain, and $P_{t|t}$ is the variance of the Kalman estimate of $\log p_{ext}^t$. Further, the differential equation (10) has a solution given by

$$g_t(p) = \begin{cases} x_0^t + \frac{y_0^t}{p_0^t} & \text{if } p \leq p_0^t \\ \max(0, \tilde{x}_0^t - C_t(p^{1-K_t} - \kappa_t)^{\frac{K_t}{1-K_t}} p^{K_t}) & \text{if } p > p_0^t \end{cases} \quad (17)$$

where $C_t, \kappa_t, \tilde{x}_0^t$ are non-negative constants such that $\tilde{x}_0^t \leq x_0^t$.

Family of optimal demand curves. We note that both Theorem 2 and Theorem 3 recommend a *family* of demand curves that satisfy the Glosten-Milgrom condition (3). The simplest curve in this family is the one where $\tilde{x}_0^t = C_t = 0$. This gives the simple demand curve which is constant except for a discontinuity at p_0^t . This corresponds to a constant sum market maker with bonding curve $y + p_0^t x = k$ with an adaptive slope given by the Kalman estimate of the external price conditioned on trader behavior.

An apparent contradiction, and a resolution. A feature of the optimal demand curves derived is the fact that they do not make it possible for a trader to express any price p_{trad}^t between 0 and ∞ . For instance, if we take the constant sum instance of the family of curves, it only allows the trader to express if the price of the trade is greater than or

25:10 Adaptive Curves for Optimally Efficient Market Making

less than p_0^t , but not the exact value of p_{trad}^t , which contradicts our assumption that the market maker can observe all of the history of trader prices $\langle p_{trad}^\tau \rangle_{\tau=1}^t$. We can resolve this contradiction by approximating the optimal demand curve as a sum of two demand curves $g_t(p) = g_t^{opt}(p) + g_t^{exp}(p)$, where the former curve is the *optimal* solution with most of the liquidity, and the latter curve with low liquidity to improve its price *expressiveness*. For instance, we can have $g_t(p) = (1 - \epsilon)g_t^{CSMM}(p) + \epsilon g_t^{CPMM}(p)$ where we have

$$g_t^{CSMM}(p) = \begin{cases} x_0^t + \frac{y_0^t}{p_0^t} & \text{if } p \leq p_0^t \\ 0 & \text{if } p > p_0^t \end{cases} \quad (18)$$

$$g_t^{CPMM}(p) = 1/\sqrt{p} \quad (19)$$

with $\epsilon \ll 1$. This combines the demand curves of the constant sum market maker (the optimal solution) and a constant product market maker (the expressive solution) as an approximation.

Other practical approximations. Another way we can parametrize the curve of the market maker is to choose it from a family of curves such that the initial marginal price matches p_0^t as prescribed by Theorem 2 and Theorem 3, and processing a trade on any curve in that family is computationally simple. To that end, we can use the Constant Mean Market Makers [22, 38] with its weighting factor as our variable parameter. Note that the CMMMs performs trades along the curve $x^\theta y^{1-\theta} = k$ where k is a constant, and x, y are the quantities of the asset and the numeraire in the market maker [22]. This ensures that no trade can exhaust either the asset or the numeraire from the AMM reserves. We know that the marginal price of the asset at any state of the reserves is given by $p = \frac{\theta y}{(1-\theta)x}$. Therefore, in our case, we set the value of parameter as $\theta_t = \frac{p_0^t x}{p_0^t x + y}$. This ensures that the starting price of any trade is p_0^t , and the market maker can only get a better price than that for a large trade.

5.2 Adaptive Kalman Filter algorithm for unknown market parameters

Need for more adaptivity. A major assumption while solving for the optimal demand curve in the Gaussian/Lognormal model was that the market parameters σ, η that control the variances of the price jump and noise were known to the AMM. These can indeed be obtained by analysing historical trading data in any market, and can be assumed to change slowly on the timescale that prices undergo changes. However, this assumption might not always hold for assets or tokens that are less well known or have no historical data. To deal with this case, we propose a modification to Algorithm 1. This ensures that we simultaneously estimate the parameters η, σ and hence help estimate the hidden external market price. To that end, we observe that we can write the likelihood function of all random variables and parameters of our model at time t as follows

$$L_t = L(\langle p_{trad}^\tau, p_{ext}^\tau \rangle_{\tau=1}^t, \eta, \sigma) = \prod_{\tau=1}^t \frac{1}{2\pi\sqrt{\eta\sigma}} \exp\left(-\frac{(p_{ext}^\tau - p_{ext}^{\tau-1})^2}{2\sigma^2} - \frac{(p_{trad}^\tau - p_{ext}^\tau)^2}{2\eta^2}\right), \quad (20)$$

which gives us the conditional log-likelihood given the trader actions as

$$E[\log L_t | \langle p_{trad}^\tau \rangle_{\tau=1}^t] = -\frac{t \log \sigma}{2} - \frac{t \log \eta}{2} - \frac{\sum_{\tau=1}^t A_\tau}{2\sigma^2} - \frac{\sum_{\tau=1}^t B_\tau}{2\eta^2}, \quad (21)$$

■ **Algorithm 2** Adaptive Kalman filter : EM algorithm estimating unknown parameters at time step t .

Require: Trader data $\langle p_{traded}^\tau \rangle_{\tau=1}^t$, Error Tolerance ϵ

- 1: $\sigma \leftarrow$ Initial guess σ_0
- 2: $\eta \leftarrow$ Initial guess η_0
- 3: $i \leftarrow 0$
- 4: $\text{LogLikelihood} = -\infty$
- 5: **while** $\text{LogLikelihood} < -\epsilon$ **do**
- 6: E step: Set A_τ, B_τ as per Equations (22) and (23), assuming guesses σ, η
- 7: M step: Set $\sigma = \sqrt{2 \sum_{\tau=1}^t A_\tau / t}, \eta = \sqrt{2 \sum_{\tau=1}^t B_\tau / t}$
- 8: $\text{LogLikelihood} \leftarrow$ Equation (21)
- 9: $i \leftarrow i + 1$
- 10: **end while**

where A_τ, B_τ are given by

$$A_\tau = E[(p_{ext}^\tau)^2 | \langle p_{traded}^\tau \rangle_{\tau=1}^t] + E[(p_{ext}^{\tau-1})^2 | \langle p_{traded}^\tau \rangle_{\tau=1}^t] + 2E[p_{ext}^\tau p_{ext}^{\tau-1} | \langle p_{traded}^\tau \rangle_{\tau=1}^t], \quad (22)$$

$$B_\tau = (p_{traded}^\tau)^2 + E[(p_{ext}^\tau)^2 | \langle p_{traded}^\tau \rangle_{\tau=1}^t] + 2p_{traded}^\tau E[p_{ext}^\tau | \langle p_{traded}^\tau \rangle_{\tau=1}^t]. \quad (23)$$

Estimating the unknowns. We now estimate the market parameters σ, η using the EM algorithm [19]. This can be done by first setting the terms A_τ, B_τ using (22) and (23) with the expectations on the RHS calculated via forward and backward runs of the Kalman filter assuming an initial guess estimate of σ, η . The forward runs of the algorithm involve computing $E[p_{ext}^\tau | \langle p_{traded}^\tau \rangle_{i=1}^\tau]$ for all $\tau = 1, \dots, t$. These are estimates of the external price given only the data in the past. This can be done using the Kalman filter updates given in Algorithm 1. Next, we use the Rauch-Tung-Striebel smoother [60], an essentially backward run of the Kalman filter algorithm given all the statistics obtained from the forward run. This computes statistics such as $E[p_{ext}^\tau | \langle p_{traded}^\tau \rangle_{i=1}^t]$, that is, the estimate of the external price in the past given *all* of the observations till the present time slot. This evaluates all of the terms in A_τ and B_τ , and completes the E-step of the EM algorithm.

After that, we use (21) to find values of σ, η that maximize the conditional log-likelihood function. This involves setting the gradient of the expected log likelihood function to zero

$$\nabla_\sigma E[\log L_t | \langle p_{traded}^\tau \rangle_{\tau=1}^t] = 0, \quad \nabla_\eta E[\log L_t | \langle p_{traded}^\tau \rangle_{\tau=1}^t] = 0. \quad (24)$$

While doing this, we ignore the dependence of A_τ, B_τ on η, σ to obtain

$$\sigma^* = \sqrt{2 \frac{\sum_{\tau=1}^t A_\tau}{t}}, \quad \eta^* = \sqrt{2 \frac{\sum_{\tau=1}^t B_\tau}{t}} \quad (25)$$

which completely specifies the M-step of the EM algorithm. This has been summarized in Algorithm 2.

Managing computation, and adapting to a non-stationary market. While Algorithm 2 added on top of our AMM helps us estimate the unknown market parameters, its computational complexity keeps growing linearly with each additional trade. This is because every trade adds another term in the series that computes the log-likelihood function, hence increasing the number of iterations in both the forward and backward runs of the Kalman

filter. Additionally, this algorithm also assumes there is no change in the market parameters η, σ with time, that is, the market conditions are stationary. We can make the algorithm less computationally heavy and adapt to non-stationarities by truncating the trading data to a recent history of p_{traded}^τ . This approach keeps track of all price estimates of the past, since they get refined with every backward run of the Adaptive Kalman algorithm, but only runs the algorithm on the recent history of price estimates and trading data.

5.3 Adversarial robustness

A prominent danger to the proper functioning of market making protocols is the presence of adversarial trader behavior. The model described so far and the optimal solutions presented remain valid only when traders interacting with the market are rational with respect to the external market. However, there is always the possibility that an AMM is being manipulated for profits extracted from other protocols (e.g. lenders[6], derivative markets, etc.) relying on the AMM as a price oracle [2]. Although any protocol using an AMM as a price oracle usually takes necessary precautions, such as ensuring a diverse portfolio of price signals, using outlier-robust statistics such as medians rather than means, etc. we show that our market making algorithms can be made robust to such market adversaries, when the proportion of such adversarial is less than half of all trading interactions.

Adversary model. The adversarial behavior we seek to guard against is the manipulation of p_{traded}^t that the AMM observes and not the external price p_{ext}^t . We assume that the external price is inferred from a deep market that is not easily manipulated. We further assume that a proportion α of the trader population is adversarial and the rest behave as per the rational model in Section 3. However, since the AMM does not know which trades are being manipulated by the adversary, we assign a sequence of learnable weights w_τ for all trade observations in the past $\langle p_{traded}^\tau \rangle_{\tau=1}^t$. To successfully manipulate the price, the adversary needs to push p_{traded}^t in a specific direction so as to induce a large discrepancy in the marginal prices of the AMM curve and p_{ext}^t .

Robust adaptive curve algorithm. A simple modification of the EM algorithm enables us to distinguish adversarial trades from honest trades [13, 64]. We first rewrite the log-likelihood function of the AMM as

$$E[\log L_t | \langle p_{traded}^\tau \rangle_{\tau=1}^t] = -\frac{t \log \sigma}{2} - \frac{t \log \eta}{2} + \sum_{\tau=1}^t \left(\frac{\log w_\tau}{4} - \frac{A_\tau}{2\sigma^2} - \frac{w_\tau B_\tau}{2\eta^2} \right), \quad (26)$$

which basically assumes that each datapoint has a different variance in noise η^2/w_τ . This implies that data with low weights have a higher variance, and are hence the adversarial outliers. We start our algorithm with an equal weight given to all datapoints, and then estimate A_τ, B_τ assuming those weights and running the forward and backward runs of the Kalman Filtering algorithm. After that, we set new weights by getting the critical points for the log-likelihood maximization using

$$\nabla_{w_\tau} E[\log L_t | \langle p_{traded}^\tau \rangle_{\tau=1}^t] = 0 \implies w_\tau^* = \frac{\eta^2}{2B_\tau}. \quad (27)$$

This completes the adversarially robust version of the Kalman filtering algorithm. We empirically demonstrate the effectiveness of the approach for $\alpha < 0.5$ compared to static curves, and a naive Kalman filtering approach (Section 7).

6 Implications for AMMs with static curves

If we view the market maker as a price oracle, then one can compare the performance of different algorithms based on how the mean squared error of the AMM changes with the incoming trades. In particular, one can compare how quickly the error goes down in the adaptive protocol proposed in Section 5.2 and a static curve such as Uniswap. This comparison has already been done for liquid markets [46], where we observe an exponential decay of the error for an adaptive protocol compared to a linear decay for a static one.

Error performance with trades. For our case, let us assume that we are comparing how Algorithm 1 performs in contrast to a static curve, if we assume that it is deployed on a blockchain with T trades in a block. Then, we can prove that the error decreases with number of trades for Algorithm 1 and stays constant for a static curve.

► **Theorem 4.** *Let there be T trades in a single block of transactions, with the external price at the creation of the block being p_{ext} . We denote by p_{KF}^T and p_{SC}^T the marginal prices of the Algorithm 1 and a static curve at the end of the block. Then, we have*

$$E[(p_{ext} - p_{KF}^T)^2] = \frac{\eta^2 \sigma^2}{T\sigma^2 + \eta^2} \quad (28)$$

$$E[(p_{ext} - p_{SC}^T)^2] = \eta^2 \quad (29)$$

Implied dynamics of static curves. We see that static curves are worse oracles because they do not use the realistic dynamical model to get the best estimate of the external price. The question then arises if there is any dynamical model that a particular static curve *is* optimal for. Note that the differential equation (10) can be viewed as an equation in $\beta_t(p)$ if the curve $g_t(p)$ is given. We now use this observation to work out the *implied* dynamical model underlying commonly used static curves. More formally, given the demand curve of a static AMM $g(p)$, we can find the corresponding $\beta_t(p)$ function by solving the following differential equation

$$\beta'_t(p)(g_t(p) - x_0^t) + \beta_t(p)g'_t(p) - pg'_t(p) = 0 \quad (30)$$

with the initial operating point p_0^t satisfying the constraint $p = \beta_t(p)$. We now solve this equation for some common CFMM curves to get the underlying implied price/trader dynamics.

Constant Sum Market Makers. The demand curve for a constant sum market maker is given by

$$g_t(p) = \begin{cases} x_0^t + y_0^t/p_0^t & \text{if } p \leq p_0^t \\ 0 & \text{if } p > p_0^t \end{cases} \quad (31)$$

where $p_0^t = p_0$ stays constant with time. Note that substituting the demand curve in (30) gives a trivial equation $\beta'_t(p) = 0$. This, coupled with the condition on the initial operating point implies that $\beta_t(p) = p_0$ for all t . Clearly, this implies a dynamical model where $\Delta p_{ext}^t = \Delta p_{trad}^t = 0$. In other words, a CSMM assumes that the external price stays constant with time and that the traders have a noiseless view of the price at all times.

Constant Mean Market Makers. The demand curve for Constant Mean Market Makers is given by

$$g_t(p) = \frac{c}{p^{1-\theta}} \left(\frac{\theta}{1-\theta} \right)^{1-\theta} \quad (32)$$

where the bonding curve for the constant mean market maker is of the form $x^\theta y^{1-\theta} = \text{constant}$, where x, y denote the reserves for the asset and numeraire respectively.

Substituting this in (30) and solving for $\beta_t(p)$ gives us the following function

$$\beta_t(p) = \frac{1-\theta}{\theta} p^\theta (p_0^t)^{1-\theta} \frac{1 - (p_0^t/p)^\theta}{1 - (p_0^t/p)^{1-\theta}} \quad (33)$$

which also satisfies $\beta_t(p_0^t) = p_0^t$. In particular, if we look at constant product market makers, we have $\theta = 1/2$, giving us $\beta_t(p) = \sqrt{p_0^t p}$. For this $\beta_t(\cdot)$ function, we have the following result.

► **Theorem 5.** *The following price and trader behavior dynamics yields $\beta_t(p)$ which obeys (33). Equivalently, the following model has a constant mean market maker with weight parameter θ as its optimally efficient solution,*

$$\log p_{ext}^t = \log p_{trad}^{t-1} + \epsilon_\sigma^t \quad (34)$$

$$\log p_{trad}^t = \log p_{ext}^t + \epsilon_\eta^t, \quad (35)$$

where $\epsilon_\sigma, \epsilon_\eta$ are independent Gaussian random variables with zero mean and variances σ, η respectively, where the variances satisfy the following conditions.

$$\sigma \ll 1 \quad (36)$$

$$\eta = \sigma \left(\sqrt{1/\theta - 1} \right) \quad (37)$$

The above theorem sheds light on why static curves fail to prevent arbitrage loss as effectively – the implied dynamic model that they assume is mismatched with more realistic trader behavior. This mismatch manifests itself in the following ways, as indicated by (36) and (37).

Low latency blockchain. Firstly, static curves assume that the price jumps between consecutive trades have a very small variance. This is equivalent to assuming that the inter-block times on the underlying blockchain go to 0. This is because the standard deviation of price jumps σ between blocks depends on the price volatility σ' as $\sigma = \sigma' \sqrt{\Delta t}$, where Δt is the inter-block arrival time. This observation confirms the conclusion reached in [41] and the broader DeFi community [62], where the arbitrage loss (or LVR) is calculated for the special case of all of the trader population being arbitrageurs and concludes that this loss indeed approaches 0 as the inter-block time goes to 0.


Constraints on noise traders. Secondly, the noise in the price that the traders see obeys a specific structure – $\eta^2 = \sigma^2(1/\theta - 1)$. For a constant product market maker, this specifically assumes that the variance of price jumps is exactly the same as the variance of the noise in trader beliefs about the price. This implies additional restrictions on price and trader behaviour that might not always be true in real markets. However, these constraints can potentially help decide how toxic and non-toxic trade flow is guided in “DEX-aggregator”

such as UniswapX [58], 1inch [1], CoWSwap [52], etc. to ensure that passive LPs (that invest in pools) get a fair price as defined by (3). The main threat that such aggregators pose to passive LPs is that much of the non-toxic (or uninformed, or noise) trades get satisfied internally without any trades flowing through the pools, while the surplus (usually toxic or informed) gets routed through the passive pools [8]. The equation (36) prescribes how much non-toxic flow should be “added in the mix” to ensure fairness to the passive pool LPs.

Dominance of DEXes. Thirdly, (34) highlights another key assumption – *the external market also reacts to the price on the AMM*. This is only true when the AMM has an amount of liquidity that is more than the external market, which is not true for most AMM pools or decentralized exchanges today [11]. This means that static curves are guaranteed to make a loss to arbitrageurs unless decentralized exchanges become the main sources of liquidity, and the inter-block arrival time on the underlying blockchain become negligible.

7 Empirical Results

In this section, we present the empirical performance of the algorithms discussed so far in this work¹.



(a) Varying price volatility σ .

(b) Varying trader noise η .

Figure 1 Percentage monetary loss per trade of our market making algorithms (Kalman Filtering and Adaptive Kalman Filtering) is much less than a static Uniswap curve for a Gaussian price jump and trader noise models.

Comparison with static curves. We simulate the model described in Section 3 and compare the performance of the algorithms we proposed (Figure [1]). We see that adapting the AMM curve according to algorithms 1 and 2 give a much lower monetary loss per trade than a static constant product curve that is used in Uniswap. Furthermore, the Adaptive Kalman Filter algorithm estimates the unknown market parameters correctly, leading to it achieving close performance with the optimal Kalman Filtering algorithm.

The same observation holds for prices that follow a geometric Brownian motion, as seen in Figure [2].

¹ Code for running all experiments has been shared [here](#)

25:16 Adaptive Curves for Optimally Efficient Market Making




Figure 2 Percentage monetary loss per trade of our making algorithms (Kalman Filtering and Adaptive Kalman Filtering) is much less than a static Uniswap curve for a Lognormal price jump and trader noise models.




Figure 3 If the AMM is treated as an oracle, then we can use the Robust Kalman Filtering to get a more accurate reading of the hidden external price.

Robustness to adversarial traders. In presence of adversarial traders, the Robust Kalman Filtering algorithm is used to change the curve of the AMM gives us a more accurate reading of the hidden external price in the presence of less than 50% of the population of traders being adversarial (Figure [3]). The adversarial traders, in this case, are assumed to perform large buy trades (with price belief p_{trad}^t about 5-7 standard deviations beyond normal trading size) to keep the AMM price above the external market. We also note that the monetary loss of the AMMs against the adversary stays in the profitable region for the adaptive curves, while the static curve suffers a loss to arbitrageurs even if the adversary is making trades that are irrational.

Robustness to non-stationary markets. Figure [4] tests the truncated version of the adaptive Kalman algorithm for changing market conditions. This algorithm is compared with the Kalman Filter algorithm, that knows the underlying market parameters exactly at every time step, and the static Uniswap curve. The monetary loss incurred by the AMMs is measured against changing variability of market conditions, measured by the volatility of




Figure 4 Percentage monetary loss per trade of our market making algorithms (Kalman Filtering and Adaptive Kalman Filtering) is much less than a static Uniswap curve for continuously changing market conditions.


price jump variance and trader noise (termed as volatility of volatility $\sigma_{\eta, \sigma}$). This metric measures the standard deviation of changing η, σ after every trade, with a sample path of these changing variables shown in Figure [4b] for $\sigma_{\eta, \sigma} = 0.04$. Recall that η, σ themselves govern how the external price is perceived by the traders and how it changes after every trade respectively. We see that the Adaptive Kalman Filter is able to perform almost as well as the Kalman Filter when the market changes are slow enough. However, the performance advantage over static curves vanishes as the changing market conditions become more erratic. This happens because the timescale over which market parameters suffer large changes becomes comparable to the timescale of the recent history considered by the truncated adaptive Kalman filter. This observation offers guidance to AMM designers on choosing the timescale over which adaptivity can offer an advantage over static markets.

8 On-chain System Implementation

Many prior works [14, 46] seek to implement adaptive market makers on a blockchain, where the adapting is done using machine learning algorithms that must be necessarily performed off-chain because of their computational load [53]. To that end, a group of protocol validators (separate from the validators of the underlying blockchain) are assumed, who run the bulk of the computation off-chain and post their results (such as satisfied orders, their prices, etc.) on-chain. However, recent developments in Layer 2 or rollup [33, 49] infrastructure, machine learning co-processors with zero-knowledge guarantees [51, 10], has given rise to several platforms that can be utilized directly to implement the adaptive market makers (or machine learning algorithms in general) we derived in the previous sections, without any need of additional validators. We draw upon these innovations for the blockchain implementation, and divide the approach into two parts. The overall design has been shown in Figure [5].

Hook Contract. The first part of such an implementation is the liquidity pool contract, which allows the canonical interactions with LPs and traders given a specific demand/bonding curve. In the Ethereum DeFi ecosystem, a recent proposal [57] by the Uniswap protocol (called Uniswap-v4) presents a highly customizable platform for adaptive market making. The main innovation is the introduction of a “hook” smart contract [54]. While prior versions

[55, 56] of the protocol presented market making in a canonical manner where the exact bonding curve followed during trade execution was static and only gave freedoms to LPs in terms of the distribution of liquidity along that fixed curve, the new platform allows the LP to specify changes to the bonding curve just before/after every single trade via the hook contract. We use this functionality to execute trades and add/remove liquidity according to the curve $x^{\theta_t}y^{1-\theta_t} = \text{constant}$. When traders/LPs put in canonical swap/liquidity addition transactions, they are first routed through the `PoolManager` contract. This maps a pool to a specific `Hook` contract that specifies all modifications to the curve before/after trade execution. All interactions with this contract are collected as input data by the second part of the system.



■ **Figure 5** System design for an on-chain implementation of our algorithms.

Off-Chain Co-processor. The second part runs the algorithm used to change the demand/-bonding curve given the history of trader interactions. This is done off-chain due to the high computational load of running the EM algorithm as part of Adaptive Kalman Filtering. For our implementation, we chose Axiom [51] as a platform to run this off-chain computation. The main part of this implementation is a typescript file `ZKcircuit` containing the details of Algorithm 2 implemented as an algebraic circuit, so that a zero-knowledge proof can be generated corresponding to the computation [50]. This file also verifiably collects data from the previous block, and runs the algorithm to come up with new estimates for the market parameters σ, η and hence the curve parameter θ . The off-chain Axiom client [50] posts this in the next block and is verified by the on-chain `Prover` contract of Axiom. This also invokes a callback to the `Hook` where the changes to the curve parameters are finally implemented. We provide open access to all files used in our implementation ².

² Code for the proof-of-concept implementation has been shared [here](#)

9 Conclusion and Future Work

Generalising to non-Gaussian and non-stationary behavior. Traders usually do not perfectly conform to the distributional assumptions we use to derive optimal solutions in Section 5.1. Our approach can be potentially be extended to such situations by the use of Neural Kalman Filters [44], which claim to work for non-Gaussian/non-stationary state space models.

Balancing toxic and non-toxic orderflow. The blockchain-level conditions for the optimality of static curves, as outlined in Section 6, provide guidance on how toxic/non-toxic orderflow, if discriminated correctly [62], should be allowed to use passive liquidity and still give LPs a fair price. Developing DEX aggregators that aware of these conditions would help limit the dangers to passive LPs in DeFi.

Extensions to other adaptive protocols. In this work, we have derived a correspondence between a dynamical model for prices and its optimal market making curve. This principle can be extended for stable control of other DeFi protocols, such as lending, that currently use static curves.

References

- 1 linch protocol. linch dex aggregator. <https://linch.io>. Accessed: 2023-05.
- 2 Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. Constant function market makers: Multi-asset trades via convex optimization, 2021. [arXiv:2107.12484](https://arxiv.org/abs/2107.12484).
- 3 Guillermo Angeris and Tarun Chitra. Improved price oracles. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. ACM, October 2020. doi:10.1145/3419614.3423251.
- 4 Guillermo Angeris, Alex Evans, and Tarun Chitra. When does the tail wag the dog? curvature and market making, 2020. [arXiv:2012.08040](https://arxiv.org/abs/2012.08040).
- 5 Jun Aoyagi. Liquidity provision by automated market makers, 2020. URL: <https://ssrn.com/abstract=3674178>.
- 6 Ayana T. Aspembitova and Michael A. Bentley. Oracles in decentralized finance: Attack costs, profits and mitigation measures. *Entropy*, 25(1), 2023. doi:10.3390/e25010060.
- 7 Marco Avellaneda and Sasha Stoikov. High frequency trading in a limit order book. *Quantitative Finance*, 8:217–224, April 2008. doi:10.1080/14697680701381228.
- 8 Ashwath Balakrishnan. Understanding uniswap’s new liquidity aggregator (tldr at the end). <https://members.delphidigital.io/feed/understanding-uniswaps-new-liquidity-aggregator-tldr-at-the-end>, 2023. [arXiv:2305.14604](https://arxiv.org/abs/2305.14604).
- 9 The Block. Block.co trade volume comparisons. <https://www.theblock.co/data/decentralized-finance/dex-non-custodial/uniswap-vs-coinbase-and-binance-trade-volume-7dma>. Accessed: 2023-09.
- 10 Brevis. Brevis website. <https://docs.brevis.network/>. Accessed: 2023-05.
- 11 centicio. Centralized exchange (cex) vs decentralized exchange (dex). which is the best crypto exchange? <https://medium.com/@centicio/centralized-exchange-cex-vs-decentralized-exchange-dex-which-is-the-best-crypto-exchange-148f48ea51c1>. Accessed: 2023-05.
- 12 Nicholas Chan and Christian Shelton. An electronic market-maker, January 2001.
- 13 Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. Kalman filtering with adversarial corruptions. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 832–845, 2022.
- 14 Dev Churiwala and Bhaskar Krishnamachari. Qlamm: A q-learning agent for optimizing fees on automated market making protocols, 2022. [arXiv:2211.14977](https://arxiv.org/abs/2211.14977).

- 15 CrocSwap. Discrimination of toxic flow in uniswap v3. <https://crocsnap.medium.com/discrimination-of-toxic-flow-in-uniswap-v3-part-1-fb5b6e01398b>. Accessed: 2023-09.
- 16 Diane Dai. Dodo integrates chainlink live on mainnet, kickstarts the on-chain liquidity revolution. <https://blog.dodoex.io/dodo-integrates-chainlink-live-on-mainnet-kickstarts-the-on-chain-liquidity-revolution-ee27e136e122>. Accessed: 2023-09.
- 17 Sanmay Das*. A learning market-maker in the glisten-milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.
- 18 Sanmay Das and Malik Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. *Advances in Neural Information Processing Systems*, 21, 2008.
- 19 A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977. URL: <http://www.jstor.org/stable/2984875>.
- 20 Graham Elliott. Forecasting with trending data. *Handbook of economic forecasting*, 1:555–604, 2006.
- 21 Shayan Eskandari, Mehdi Salehi, Wanyun Catherine Gu, and Jeremy Clark. SoK. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*. ACM, September 2021. doi:10.1145/3479722.3480994.
- 22 Alex Evans. Liquidity provider returns in geometric mean markets, 2020. [arXiv:2006.08806](https://arxiv.org/abs/2006.08806).
- 23 Alex Evans, Guillermo Angeris, and Tarun Chitra. Optimal fees for geometric mean market makers, 2021. [arXiv:2104.00446](https://arxiv.org/abs/2104.00446).
- 24 Rafael Frongillo, Maneesha Papireddygar, and Bo Waggoner. An axiomatic characterization of cfms and equivalence to prediction markets. *arXiv preprint*, 2023. [arXiv:2302.00196](https://arxiv.org/abs/2302.00196).
- 25 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. doi:10.1016/0304-405X(85)90044-3.
- 26 Mohak Goyal, Geoffrey Ramseier, Ashish Goel, and David Mazières. Finding the right curve: Optimal design of constant function market makers, 2023. [arXiv:2212.03340](https://arxiv.org/abs/2212.03340).
- 27 SANFORD J. GROSSMAN and MERTON H. MILLER. Liquidity and market structure. *The Journal of Finance*, 43(3):617–633, 1988. doi:10.1111/j.1540-6261.1988.tb04594.x.
- 28 F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. Particle filters for positioning, navigation, and tracking. *IEEE Transactions on Signal Processing*, 50(2):425–437, 2002. doi:10.1109/78.978396.
- 29 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*. ACM, September 2022. doi:10.1145/3558535.3559772.
- 30 Thomas S. Y. Ho and Hans R. Stoll. The dynamics of dealer markets under competition. *The Journal of Finance*, 38(4):1053–1074, 1983. URL: <http://www.jstor.org/stable/2328011>.
- 31 Peter Johnson and Sai Nimmagadda. The relentless rise of stablecoins. <https://digify.com/a/#/f/p/ef09be008ee64ab68bda4f0a558302a2>. Accessed: 2023-09.
- 32 Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- 33 Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, and Edward W. Felten. Arbitrum: Scalable, private smart contracts. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1353–1370, Baltimore, MD, August 2018. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity18/presentation/kalodner>.
- 34 Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985. URL: <http://www.jstor.org/stable/1913210>.
- 35 DeFi Llama. Uniswap-v3 tvl comparison for stable coins vs non-stablecoins. <https://defillama.com/protocol/uniswap-v3>. Accessed: 2023-09.
- 36 Stefan Loesch, Nate Hindman, Mark B Richardson, and Nicholas Welch. Impermanent loss in uniswap v3, 2021. [arXiv:2111.09192](https://arxiv.org/abs/2111.09192).

- 37 Christos Makidris. Front running, bots, slippage, oracle pricing errors: Amms are great, but there are problems. <https://cointelegraph.com/magazine/trouble-with-crypto-automated-market-makers/>. Accessed: 2023-09.
- 38 F. Martinelli. Balancer whitepaper. <https://balancer.fi/whitepaper.pdf>. Accessed: 2023-05.
- 39 Conor McMenamin, Vanesa Daza, and Bruno Mazorra. Diamonds are forever, loss-versus-rebalancing is not, 2022. [arXiv:2210.10601](https://arxiv.org/abs/2210.10601).
- 40 Raman K. Mehra. On the identification of variances and adaptive kalman filtering. *IEEE Transactions on Automatic Control*, 15:175–184, 1970. URL: <https://api.semanticscholar.org/CorpusID:238574860>.
- 41 Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. Automated market making and arbitrage profits in the presence of fees, 2023. [arXiv:2305.14604](https://arxiv.org/abs/2305.14604).
- 42 Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. A myersonian framework for optimal liquidity provision in automated market makers, 2023. [arXiv:2303.00208](https://arxiv.org/abs/2303.00208).
- 43 Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing, 2022. [arXiv:2208.06046](https://arxiv.org/abs/2208.06046).
- 44 Beren Millidge, Alexander Tschantz, Anil Seth, and Christopher Buckley. Neural kalman filtering, 2021. [arXiv:2102.10021](https://arxiv.org/abs/2102.10021).
- 45 Vijay Mohan. Automated market makers and decentralized exchanges: a defi primer, December 2020. doi:10.2139/ssrn.3722714.
- 46 Viraj Nadkarni, Jiachen Hu, Ranvir Rana, Chi Jin, Sanjeev Kulkarni, and Pramod Viswanath. Zeroswap: Data-driven optimal market making in defi. *arXiv preprint*, 2023. [arXiv:2310.09413](https://arxiv.org/abs/2310.09413).
- 47 Viraj Nadkarni, Sanjeev Kulkarni, and Pramod Viswanath. Adaptive curves for optimally efficient market making. *arXiv preprint*, 2024. [arXiv:2406.13794](https://arxiv.org/abs/2406.13794).
- 48 Alex Nezlobin. Order flow toxicity on dexes. <https://ethresear.ch/t/order-flow-toxicity-on-dexes/13177>. Accessed: 2023-09.
- 49 Optimism. Optimism docs. <https://community.optimism.io/>. Accessed: 2023-09.
- 50 Axiom Protocol. Axiom client circuit. <https://docs.axiom.xyz/docs/axiom-developer-flow/axiom-client-circuit>. Accessed: 2023-05.
- 51 Axiom Protocol. Axiom website. <https://docs.axiom.xyz/>. Accessed: 2023-05.
- 52 CoW protocol. Cowswap docs. <https://docs.cow.fi/overview/coincidence-of-wants>. Accessed: 2023-09.
- 53 Gate protocol. Offchain compute is all you need. <https://www.gate.io/learn/articles/off-chain-compute-is-all-you-need/1225>. Accessed: 2023-05.
- 54 Uniswap Protocol. Hooks on uniswap v4. <https://docs.uniswap.org/contracts/v4/concepts/hook-deployment>. Accessed: 2023-05.
- 55 Uniswap Protocol. Uniswap v2 core. <https://uniswap.org/whitepaper.pdf>. Accessed: 2023-09.
- 56 Uniswap Protocol. Uniswap v3 core. <https://uniswap.org/whitepaper-v3.pdf>. Accessed: 2023-09.
- 57 Uniswap Protocol. Uniswap v4 docs. <https://docs.uniswap.org/contracts/v4/concepts/intro-to-v4>. Accessed: 2023-09.
- 58 Uniswap Protocol. Uniswapx docs. <https://blog.uniswap.org/uniswapx-protocol>. Accessed: 2023-09.
- 59 Palamarchuk Roman. Flash loan attacks: Risks and prevention. <https://hacken.io/discover/flash-loan-attacks/>. Accessed: 2023-05.
- 60 Robert H Shumway and David S Stoffer. An approach to time series smoothing and forecasting using the em algorithm. *Journal of time series analysis*, 3(4):253–264, 1982.
- 61 Rohan Tangri, Peter Yatsyshin, Elisabeth A. Duijnste, and Danilo Mandic. Generalizing impermanent loss on decentralized exchanges with constant function market makers, 2023. [arXiv:2301.06831](https://arxiv.org/abs/2301.06831).

- 62 TheiaResearch. Better blockchains lead to more profitable liquidity providers. <https://twitter.com/TheiaResearch/status/1790717593440952757>. Accessed: 2024.
- 63 Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT Press, Cambridge, Mass., 2005. URL: <http://www.amazon.de/gp/product/0262201623/102-8479661-9831324?v=glance&n=283155&n=507846&s=books&v=glance>.
- 64 Jo-Anne Ting, Evangelos Theodorou, and Stefan Schaal. Learning an outlier-robust kalman filter. In Joost N. Kok, Jacek Koronacki, Raomon Lopez de Mantaras, Stan Matwin, Dunja Mladenič, and Andrzej Skowron, editors, *Machine Learning: ECML 2007*, pages 748–756, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- 65 Norbert Wiener. *Cybernetics: or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, 2 edition, 1948.
- 66 Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. Sok: Decentralized exchanges (dex) with automated market maker (amm) protocols. *ACM Comput. Surv.*, 55(11), February 2023. doi:10.1145/3570639.
- 67 Adelyn Zhou. Flash loans aren't the problem, centralized price oracles are. <https://www.coindesk.com/tech/2020/11/11/flash-loans-arent-the-problem-centralized-price-oracles-are/>. Accessed: 2023-09.
- 68 K. Zhou, J.C. Doyle, and K. Glover. *Robust and Optimal Control*. Feher/Prentice Hall Digital and. Prentice Hall, 1996. URL: <https://books.google.com/books?id=RPSOQgAACAAJ>.