Optimal Automatic Market Maker

Alexis Direr *

August 30, 2023

Abstract

This paper introduces discrete automatic market makers (DAMMs), a general family of automatic market makers characterized by a piecewise linear pricing curve which can approximate any continuous curves with arbitrary accuracy. Optimal sizes are derived from an infinite horizon utility-based convex problem. Our results back the common practice of investing liquidity where the market maker expects the price to fluctuate, with two qualifications. First, the degree of liquidity concentration heavily depends on market maker's degree of risk aversion and will be typically different from the price distribution. Second, a lower fraction of capital will be allocated towards price ranges far from current price as it takes time for the price to move to remote ranges and put the liquidity at work. The higher the opportunity costs of keeping liquidity idle, the higher the investment decay relative to current price.

J.E.L. codes: D8, E21

Keywords: Trading and Market Microstructure, Blockchain, Decentralized Finance

*Univ. Orléans, LÉO. Address: rue de Blois - BP 26739, 45067 Orléans Cedex 2 France. Email: alexis.direr@univ-orleans.fr. ORCiD: 0000-0002-4459-7780. The paper has benefited from early discussions with Vincent Danos and Hamza el Khalloufi. I thank for useful comments participants of the Laboratoire d'Économie d'Orléans and Mangrove's seminars.

Introduction

Decentralized finance is changing the way financial services are delivered by allowing users to access a wide range of financial products without the need for a third-party intermediary (Schär, 2021). One of the catalysts driving this new model of finance was the invention of automated market makers (AMMs), which provide an efficient and decentralized means of exchanging digital assets. An AMM is an algorithmic protocol that uses an automated pricing mechanism to facilitate liquidity for digital assets on public blockchains like Ethereum. They act like a two-sided platforms in which traders interact with liquidity providers (LPs) through a pricing curve which determines the price at which the two assets in the pool are exchanged. On one side, LPs provide capital by depositing the assets in a smart contract that implements the trading logic. They receive a share of trading fees proportional to their contribution to the pool of assets. On the other side, traders tap into the pool by exchanging one asset for the other according to an exchange rate which depends on the relative scarcity of the two assets in the pool.[1]

AMMs are particularly well fitted to public blockchains characterized by high latency and storage costs. The algorithmic determination of the exchange rate greatly simplifies the tasks of trading assets and market making the liquidity. AMMs eliminate the need for manual quote submission and provide instant liquidity with moderate transaction costs. They make it trivially easy for users to create new markets and provide liquidity to existing markets. They constitute nowadays the most liquid trading venues for a large number of digital assets. Uniswap, the leading AMM on the blockchain Ethereum, has been in 2022 the venue of 68M transactions totaling $620B in trading volume.

[1]Prominent AMMs include Curve (Egorov 2019), Balancer (Martinelli and Mushegian, 2019) and Uniswap (Adams, Zinsmeister and Robinson, 2020 and Adams et al., 2021).

Beyond the wide variety of pricing curves adopted, a key characteristic is path independence.[2] A pricing curve is path independent if each exchange rate is associated with an invariant quantity of the two assets available in the pool. It follows that every time the exchange rate returns to a previous level, the quantities of assets in the pool are also restored. This property is an important safeguard for AMMs which exchange rate is set without knowledge of the external market price. This extreme case of asymmetric information (Kyle, 1985) protects LPs from being out-traded by informed traders and arbitrageurs as long as the price goes back and forth.

Despite those benefits, the design of AMMs is still rudimentary. Pricing curves rely on ad-hoc functionals with little relation to optimization models. Most discussions and implementations have revolved around the curvature degree of the pricing curve and the spread of the liquidity in relation with price volatility. Quasi-linear pricing curves are adapted to low-volatility asset pairs (Angeris, Evans and Chitra, 2020), whereas more convex curves like Uniswap are best optimized for volatile and long-tail asset pairs (Angeris et al., 2019). Beyond the degree of curvature, other properties of the pricing curve are hardly studied by lack of a proper optimization framework. The rapid pace at which new AMMs go to market outlines the need of better theoretical frameworks to asses their utility and devise optimal strategies. The lag between practice and theory can be illustrated by the launch of Uniswap V3 in 2021 which provides LPs with the breakthrough option of concentrating their liquidity over narrow price ranges. While this gives LPs the opportunity to actively market make asset pairs, it is still unclear how to define an optimal market making policy in those markets.

[2]See Buterin (2017) for applying the concept to AMMs. Path independence is generally obtained if fees earned by the pool are excluded from the return (Angeris and Chitra, 2020). Adding fee revenue only strengthens the argument that LPs returns in path independent AMMs cannot be harmed by price volatility.

This paper proposes a micro-founded automated market making model which allows to derive the optimal pricing curve within the family of path independent AMMs. We introduce discrete AMMs characterized by a piecewise linear curve and which can approximate any continuous curves with arbitrary accuracy. Given market maker's price expectations, risk aversion and time preferences, order sizes maximize her intertemporal discounted utility.[3] Our results back the common practice of investing capital where the market maker expects the price to fluctuate, with two qualifications. First, the degree of liquidity concentration will heavily depend on the market maker's degree of risk aversion. For this reason, it will be typically different from the price distribution. On top of those factors, a lower fraction of the capital will be allocated towards price ranges far from current price as it takes time for the price to move to remote ranges and put the liquidity at work. The higher the opportunity costs of keeping liquidity idle, the higher the investment decay relative to current price.

The recent spectacular growth of AMMs on public blockchains has attracted academic attention leading to an emergent literature on the functioning of AMMs and their optimal design. Angeris et al. (2019) presents an early analysis of Uniswap, and more broadly of constant-product AMMs. Port and Tiruviluamala (2022) studies new methods to mix existing pricing curves. Wu and McTighe (2022) introduces a new type of pricing rule, the constant power root market marker. Those articles rely on specific functionals for the pricing rule. Capponi and Jia (2021) studies the optimal degree of convexity of the pricing curve which is modeled as a trade-off between protecting LPs from arbitrage losses and worsening the price impact, which deters volume and reduces trading fees. Neuder et al. (2021) investigates alternative strategies of liquidity concentration that LPs

[3]The term market maker may refer to the person who sets the pricing curve of the AMM at the smart contract level on behalf of liquidity providers, or the investor who market makes an asset pair for example by providing liquidity in an AMM with concentrated liquidity.

may follow in Uniswap V3 AMMs. They compute the expected utility derived from the stream of fees for each strategy and find that the uniform allocation is near-optimal for risk averse LPs. Fan et al. (2022) also proposes an optimization framework that translates LPs' beliefs about future asset valuations into an optimal pricing curve. In both articles, the scope of the results is limited to Uniswap V3 pricing rules. Goyal et al. (2022) study a more general optimal pricing curve conditional on a set of beliefs about the price dynamics. Their conclusions rely on the assumption of risk neutrality as LPs maximize transaction volumes settled by the AMM. Bergault et al. (2022) study efficient AMMs in the mean-variance framework when the pricing function can pull price information from a centralized exchange. They show that complementing the pricing function with a price oracle provides more efficient strategies.

The remaining of the paper is organized as follows. Section 2 presents the market making model and the constraints imposed by path independence. Section 3 compares the characteristics of the model with those of traditional continuous AMMs. Section 4 presents the optimality conditions. Section 5 analytically solves the market making problem with three execution prices. The general problem with an arbitrary number of prices is numerically solved in Section 6. Section 7 concludes.

1 Market making model

The strategy provides liquidity for a trading pair X/Y and posts bids of size δx_n^+ and asks of size δx_n^-, $n = 0, 1, ..., N$ over a predefined price grid $P = \{p_0, p_1, ..., p_N\}$ with $p_0 < p_1 < ... < p_N$, over which the price is expected to fluctuate. The AMM cannot short the assets. Asks are posted for every price belonging to the set P greater than the current price and bids for every price

inferior to the current price. No orders are posted where the last order has been filled. For example, at price p_n, the two closest orders are: (i) an ask at price p_{n+1} of size δx^-_{n+1} and (ii) a bid at price p_{n-1} of size δx^+_{n-1}. If the price goes up to p_{n+1}, the ask is taken first. The two closest orders are now: (i) a bid at price p_n of size δx^+_n and (ii) an ask at price p_{n+2} of size δx^+_{n+2}. If the price returns to p_n, the bid is taken first and so forth.

The strategy is path independent if a fixed relation exists between every price levels belonging to the set P and AMM's reserves in X and Y. Suppose that the AMM's exchange rate starts from p, then fluctuates when investors' and arbitrageurs' trade with the AMM's pool. If the exchange rate returns to p, AMMs reserves will return to their initial value as well.

A path independent strategy is defined over order pairs $(\delta x^+_n, \delta x^-_{n+1})$ composed of a bid at price p_n and an ask at a higher price p_{n+1}. They always produce a positive profit when they are sequentially filled provided they have the same size. A path independent strategy satisfies:

H1 $\delta x^+_n = \delta x^-_{n+1}$ for all $n \in \{0, 1, 2, ..., N-1\}$.

H2 Order pair's profit $b_n = -p_n \delta x^+_n + p_{n+1} \delta x^-_{n+1}$ is not reinvested in the pool.

The profit b_n associated with the filling of the order pair comes from selling at p_{n+1} and buying at the lower price p_n.

To prove path independence based on H1 and H2, let us start from price p_n, $n \in \{0, 1, ..., N-1\}$. If the price increases to p_{n+q}, $q = 1, ...N+1-n$, then

returns to p_n, AMM's new balances are:

$$\begin{aligned} x'_n &= x_n - \sum_{j=1}^{q} (\delta x^-_{n+j} + \delta x^+_{n+j-1}) \\ y'_n &= y_n + \sum_{j=1}^{q} (p_{n+j}\delta x^-_{n+j} - p_{n+j-1}\delta x^+_{n+j-1}) \\ &= y_n + \sum_{j=1}^{q} p_{n+j}(\delta x^-_{n+j} - \delta x^+_{n+j-1}) + \sum_{j=1}^{q} b_{n+j-1} \end{aligned}$$

The equalities $x'_n = x_n$ and $y'_n = y_n$ obtain if H1 and H2 are satisfied. The proof of the invariance when the price decreases before returning to p_n is symmetric.

The quantities held by the AMM for each execution price belonging to P can be expressed under H1 and H2 as:

$$\begin{aligned} x_n &= \sum_{i=n}^{N-1} \delta x_i^+ \\ y_n &= w_N - \sum_{i=n}^{N-1} p_i \delta x_i^+ \end{aligned} \tag{1}$$

In particular, AMM's budget is $w_N = y_N$ when the price reaches its upper bound, as all assets X are sold (recall that by assumption the price cannot increase anymore at this level).

2 Comparison with continuous AMMs

The strategy leads to a generalized family of path independent AMMs called discrete AMMs (DAMMs). In classical continuous AMMs, traders are allowed to exchange any quantities of the two assets insofar as a tight relationship between the quantities x and y held by the contract is preserved:

$$x \longrightarrow y = f(x)$$

with f the AMM's pricing curve, a continuously differentiable decreasing function. An example of curve is displayed in Fig. 1. An AMM following this rule is bound to sell or buy δx vs. δy in a way preserving the constraint:

$$x + \delta x \longrightarrow y - \delta y = f(x + \delta x)$$

<Insert Fig. 1>

If the AMM has trading fees, the relation is invariant if the fees are excluded from the pool, just like the benefits are extracted from the discrete AMM pool in H2. A popular pricing curve is defined by the constant product $xy = k$. Reserves increase or decrease in the pool in a way that keeps constant the parameter k or the geometric mean $\sqrt{xy}$ (Evans, 2020). The exchange rate is constantly realigned with the external market price thanks to arbitrageurs selling the expensive asset against the cheap one. Under a frictionless no-arbitrage condition, the AMM exchange rate is equal to the market price at the margin (see Fig. 2):

$$p = -\frac{dy}{dx} = -f'(x)$$

<Insert Fig. 2>

A strict relation between the quantities held in the two assets also exists in a DAMM for every price belonging to the set P:

$$x_n = \sum_{i=n}^{N-1} \delta x_i^+ \longrightarrow y_n = w_N - \sum_{i=n}^{N-1} p_i \delta x_i^+$$

where the relation between the two quantities is step-wise decreasing. The internal exchange rate is also aligned with market price but only intermittently. The relation between the quantities held by the DAMM is graphically represented with three prices in Fig. 3.

<Insert Fig. 3>

In Fig 4, the relation between x_n and y_n is composed of three linear segments. The slope of each segment corresponds to a price belonging to the set $P = \{p_0, p_1, p_2\}$. The total liquidity offered and demanded at each price is given by the length of the corresponding segment measured on the two axes.

<Insert Fig. 4>

The market price will typically be distinct from the prices included in the set P. This is the case of p in the figure, which lies between p_1 and p_2. If the price decreases below p_1, arbitrageurs buy the depreciating asset X. They take $y_2 - y_1$ from the AMM in exchange of $x_2 - x_1$.

A major difference with continuous AMMs is that the relation between the two assets is not constrained by a functional but by a step-wise relation which can approximate any continuous function with arbitrary accuracy provided the price grid is fine enough.[4]

3 Optimality conditions

This section presents the maximization problem faced by the market maker (MM) given the set of execution prices P and the price dynamics.

[4]Also, a DAMM earns the spread when the price goes back and forth as the step between any two prices is non-zero. Continuous AMMs apply a trading fee to earn a local spread.

3.1 Price dynamics

The price is assumed to follow an ergodic Markov chain over execution prices $P = \{p_1, ..., p_N\}$. Starting from $p_n \in P$, the probability of the price going down to p_{n-1} is π_n^-, up to p_{n+1} π_n^+ and staying at p_n $\pi^s = 1 - \pi_n^- - \pi_n^+$. The price cannot instantly jump to higher or lower price levels, so only the three main diagonals of the transition matrix are non-zero. Also p_0 cannot go down ($\pi_0^- = 0$) and p_N cannot go up ($\pi_N^+ = 0$).

Starting at p_n, the profit is extracted every time the price transits to p_{n+1} and eventually goes back to p_n, or moves to p_{n-1} and after an uncertain number of periods, goes back to p_n. The price sequence is assumed to start at the minimal price p_0. As a result, all price increases initiate or prolong a price loop whereas all price decreases partially or fully close a loop, locking in a benefit consumed by the MM.

Figure 5 shows possible price sequences. Starting at p_0, the price can only make clockwise price loops, either short ones ($p_0 \to p_1 \to p_0$) or longer ones ($p_0 \to p_1 \to p_2 \to ...$). Due to the ergodic naure of the Markov chain and the infinite horizon, every opening price loops will eventually close and lock in a profit for the MM.

<Insert Fig. 5>

3.2 Value functions

The MM derives utility $u(b_n)$ from benefit b_n, with u an increasing, twice continuously differentiable and concave function. Future utility of profit delayed by t periods is discounted by the factor $\beta \in (0, 1)$. Value functions V_n, $n \in \{0, 1, ..., N\}$, are intertemporal discounted sum of expected utility of profits. The maximiza-

tion problem with infinite horizon takes the iterative form:

$$\max_{\{x_n^+, n=0,\dots,N-1\}} V_n = \pi_n^- \beta\Big(V_{n-1} + u(b_{n-1})\Big) + \pi_n^+ \beta V_{n+1} + \pi_n^s \beta V_n \tag{2}$$

with $b_{n-1} = (p_n - p_{n-1})\delta x_{n-1}^+$. As previously noted, a price decrease is a necessary and sufficient condition for the MM to earn a profit. An analytical solution is possible in the special case in which the MM sets three execution prices, which is now reviewed before addressing the general case.

4 The three-price case

The price transits between three price levels: p_0, p_1 and p_2. To simplify further, we depart from the general model by assuming that the economy starts at state 1 (not 0) and that the price never stay at the same level more than one period. Starting at p_1, the MM either sells at p_2, buy back at p_1, or buys at p_0 and sells back at p_1. The state transition is presented in Fig. 6.

<Insert Fig. 6>

The profit obtained from buying at p_0 and selling at p_1 is $b_0 = (p_1 - p_0)\delta x_0^+$, whereas it is $b_1 = (p_2 - p_1)\delta x_1^+$ when the MM sells at p_2 and buys at p_1. MM's budget constraint expressed in state 2 is

$$p_1 \delta x_1^+ + p_0 \delta x_0^+ = w_2$$

The value function at p_1 is:

$$V_1 = \pi^+ \beta^2\Big(V_1 + u(b_1)\Big) + \pi^- \beta^2\Big(V_1 + u(b_0)\Big)$$

with π^+ the probability of a clockwise loop $p_1 \to p_2 \to p_1$ and $\pi^- = 1 - \pi^+$ the probability of an anticlockwise loop $p_1 \to p_0 \to p_1$. Regrouping the terms:

$$V_1 = \frac{\beta^2}{1-\beta^2}\Big(\pi^+ u\Big((p_2 - p_1)\delta x_1^+\Big) + \pi^- u\Big((p_1 - p_0)\delta x_0^+\Big)\Big)$$

and injecting the budget constraint $\delta x_0^+ = (w_2 - p_1 \delta x_1^+)/p_0$, the optimality condition w.r.t. x_1^+ equalizes the expected marginal utility of the capital used in the two price loops, weighted by their return rates:

$$\frac{p_1 - p_0}{p_0}\pi^- u'(b_1) = \frac{p_2 - p_1}{p_1}\pi^+ u'(b_2) \tag{3}$$

The MM concentrates the capital where it is the most profitable in expectation. If the probability π^- of an anticlockwise loop is high, all else equal $b_0 > b_1$ or $\delta x_0^+ > \delta x_1^+$, which means that the MM keeps a large fraction of her reserve in asset Y at price p_1. The MM also concentrates the capital *ceteris paribus* where the price return is the highest.

Concentrating capital where it is the most profitable is also riskier. The more risk averse the MM (the more concave the valuation function u), the more equal the reserves posted at the two ends of the market.

In the log case $u(b) = log(b)$, the share of capital kept for the two loops simplifies to:

$$p_0 \delta x_0^+ = \pi^+ w_2$$

$$p_1 \delta x_1^+ = (1 - \pi^+) w_2$$

In accordance with intuition, the higher the probability of a price decline, the larger the bid at p_0 and the smaller the bid at p_1.

5 Optimal order sizes with an arbitrary number of prices

This section shows how to compute and simulate optimal order sizes in the general case with $N+1$ price levels.

5.1 Model

Value functions in Eq. (2) can be represented in a matrix form:

$$V = \beta \Pi V + \beta \pi^{-} U$$

with Π the transition matrix of the Markov chain followed by the price and V and U the column vectors of the value functions V_n and utility functions $u(b_{n-1})$. Value functions can be expressed as linear combinations of period utilities: $V = \Gamma U$ with $\Gamma = (I - \beta\Pi)^{-1}\beta\pi^{-}$. The sums of utilities are weighted by the coefficients γ_n^t of the matrix Γ:

$$V_n^t = \sum_{n=1}^{N} \gamma_n^t u(b_{n-1})$$

with V_n^t the state $n = 0, ..., N$ value function evaluated at state $t = 0, ..., N$. The coefficients γ_n^t indicate the discounted probability of each state. The closer to 1, the more likely and temporally closer the profit. Next, the order size δx_0^+ is singled out from the MM's budget constraint

$$\delta x_0^+ = \frac{w_N}{p_0} - \sum_{n=1}^{N-1} \frac{p_n}{p_0} \delta x_n^+$$

and replaced in the value function:

$$V_n^t = \gamma_1^t u\Big(\frac{p_1 - p_0}{p_0} \sum_{n=1}^{N-1} \Big(w_N - p_n \delta x_n^+\Big)\Big) + \sum_{n=1}^{N-1} \gamma_{n+1}^t u\Big((p_{n+1} - p_n)\delta x_n^+\Big)$$

Optimality conditions w.r.t. δx_n^+, $n = 1, ..., N+1$ are:

$$\frac{p_n - p_{n-1}}{p_{n-1}} \gamma_n^t u'\Big((p_n - p_{n-1})\delta x_{n-1}^+\Big) = \frac{p_{n+1} - p_n}{p_n} \gamma_{n+1}^t u'\Big((p_{n+1} - p_n)\delta x_n^+\Big) \qquad (4)$$

which are equivalent to the optimality condition (3) encountered in the three-price problem, with similar interpretations.

The coefficients γ_n^t and therefore the optimality conditions depend on the current state t for two reasons. First the probability to go through a price loop around state n will depend on its position relative to state t. Second, the further the state n from current state t, the longer the price will take on average to go through a loop around state n. The utility of the associated benefit will be more discounted as a result. For those reasons, MM's preferences will be time inconsistent: order sizes which are optimal at state t will change after transiting to another state. Several concepts of rationality coexist in this case (Strotz, 1955). In accordance with the fact that the market making policy is usually encoded in an immutable smart contract, we assume that the MM can commit to a state-t optimal policy even when the current state is not t.

5.2 Numerical simulations

MM's preferences are represented by a constant relative risk aversion utility function:

$$u(b) = \frac{b^{1-\sigma}}{1-\sigma}$$

with $\sigma > 0$ the coefficient of relative risk aversion. The price grid over which orders are posted is composed of 46 prices, starting at $p_0 = 1000$ and following a geometrical sequence $p_{n+1} = 1.025 p_n$ with the maximum price p_{45} equal to 3038. The geometric distribution of prices allows each local price loop to give the MM the same price return $(p_{n+1} - p_n)/p_n$. Otherwise the MM would have an incentive to concentrate capital where the return is the highest.

Starting from n, the state can transit to $n-1$, n or $n+1$ one period later

but not to lower or higher states. As a result, the 46 $\times$ 46 transition matrix Π has zero values outside the three main diagonals. A symmetric matrix is used in numerical simulations. Given the probability $\theta \in (0,1)$, the transition probabilities for $n \in \{1,2,...,46\}$ are:

$$\Pi[n,n-1] = \Pi[n,n+1] = \theta \text{ and } \Pi[n,n] = 1-2\theta$$

except $\Pi[1,1] = \Pi[46,46] = 2\theta$.

The stationary price distribution associated with Π indicates where the MM expects the price to fluctuate on average. The use of a symmetric matrix results in a uniform stationary distribution with all states having the same stationary probability $1/(N+1)$, whatever the value taken by θ. This eliminates a second confounding with other factors of size concentration.

To numerically solve optimal order sizes, the value of δx_0^+ is arbitrarily fixed to compute state 0 marginal utility:

$$U_0' = \gamma_1^k \frac{p_1 - p_0}{p_0} \Big((p_1 - p_0)\delta x_0^+ \Big)^{-\sigma}$$

Then, all other order sizes can be derived by using the optimality condition (4):

$$\delta x_n^+ = \frac{1}{p_{n+1} - p_n} \left(\frac{p_{n+1} - p_n}{p_n} \frac{\gamma_{n+1}^k}{U_0'} \right)^{1/\sigma}$$

Last, order sizes are resized to fit the MM's budget $w_N = 1000$.

5.3 Results

The values taken by the preference parameters in the baseline simulations are $\sigma = 2$ and $\beta = 0.999$. Based on a weekly transition, the discount factor is consistent with a 5.3% annualized return. Value functions V_n^t for states $n = 0,...,N$ are

evaluated ate state $t = 0$. As previously noted, the value functions will be different if the reference state is changed. The transition probability between states is set to $\theta = 1/3$.

Figure 7 shows the distribution of order sizes $p_n \delta x_n^+$ at prices p_n and Fig. 8 the corresponding pricing curve. Order size is decreasing with the price despite a uniform stationary distribution. This is due to the interaction of two effects. First, the MM evaluating value functions at state 0, the higher the execution price the longer it takes for the market price to move, cross the execution price and eventually make profitable loops. Before this time, the capital remains idle, which has an economic opportunity cost captured by the discount factor β. Second, while concentrating capital over the closest prices is more profitable in expectation, risk aversion prevents the MM from over-concentrating over this range. Spreading capital over the full range also benefits from financial diversification. Since price loops open and close sequentially, the return from capital at a given price loop is negatively correlated to the return to other price loops.

Those interpretations are confirmed when parameters' values are changed. Fig. 9 shows the distribution of sizes for four different values of the relative risk aversion parameter σ in the baseline economy ($\beta = 0.999$, $t = 0$ and $\theta = 1/3$). The more risk averse the MM, the more evenly distributed the order sizes. For moderate risk aversion ($\sigma = 0.5$), the capital is strongly skewed in the neighborhood of the reference state t, where it is the most profitable in expectation.

Fig. 10 shows the size distribution when the discounting factor β takes the values 0.99, 0.999 (baseline), 0.9999 and 0.9999 (keeping $\sigma = 2$, $t = 0$ and $\theta = 1/3$ fixed). In accordance with previous intuitions, the higher the MM's opportunity cost of time, the more concentrated the price distribution around the current state. To the opposite, when the MM's discounting factor tends towards 1, the

size distribution converges to the stationary price distribution, which is uniform by assumption. The last configuration constitutes an important benchmark as it corresponds to the common practice in Uniswap V3-type AMMs to allocate the capital where the price is believed to fluctuate.

Fig. 11 shows the size distribution for two alternative reference states: $t = 23$ and $t = 46$ in the baseline scenario. In both cases, the closer other states are to the reference state, the higher the order sizes posted by the MM.

Last, Fig. 12 shows the size distribution when the probability of transiting to an adjacent state is lower ($\theta = 0.3$) and higher ($\theta = 0.5$) compared to the baseline model ($\theta = 1/3$). As expected, the size distribution is more concentrated around the reference state when the probability is lower. A lower probability means that a longer period is needed on average for the market price to move to remote execution prices, which reduces its profitability once discounted.

6 Conclusion

This paper proposes an optimal utility-based market making model for AMMs which can adapt to any expected price distribution. Instead of looking for the best functional, the model adopts a fully non-parametric approach by optimizing order sizes over a discrete price set. Any continuous pricing curves can then be approximated with arbitrary accuracy by narrowing the gap between prices at which orders are posted. The curve can be manually implemented by LPs in Uniswap-V3-type AMMs which let users concentrate their liquidity where they wish, or in programmable limit order books in which the strategy can be coded.

Our results back the common practice of allocating the capital proportionally to where the price is expected to fluctuate, with two caveats. First, how con-

centrated the capital should be compared to the price distribution is ultimately a risk-return trade-off. The model allows to relate in a quantitative way MM's risk preferences and the degree of concentration. Second, a lower fraction of the capital should be allocated over price ranges far from current price as it takes time for the price to move to remote ranges and put the liquidity at work. The higher the opportunity costs of keeping liquidity idle, the higher the investment decay relative to current price. Another lesson is that risk averse MM should allocate a positive amount of capital over the full price range. Paralleling the result of Samuelson (1969) for financial market, in absence of transaction costs, a small amount of capital is strictly preferred to zero capital everywhere the price has a non-zero probability to loop.

The model also allows to reverse engineer a pricing curve to retrieve the MM's price distribution given assumptions about time and risk preferences. As an illustration, Uniswap V2-type AMMs in which the capital is distributed evenly across all price ranges–from zero to infinity is consistent with an expected price uniformly distributed over the maximum range and a MM being extremely risk averse and having a time opportunity cost of zero. Those validity conditions are hard to justify on empirical grounds.

Although the distribution of order sizes is endogenous, the model takes as given the price set over which the MM expects the market price to fluctuate. The question of whether the strategy is best optimized for an arbitrary large number of execution prices or for a price set with a finite number of prices is not addressed by the paper. In the first case, the pricing curve would approach continuous curves found in traditional AMMs, which would open up the issue of the optimal fee rate that the AMM should apply to remain profitable. In the second case, the question remains of the optimal spacing of execution prices. Those unanswered questions are worth being pursued in future studies.

References

Adams H., Zinsmeister N. and D. Robinson (2020). Uniswap v2 core. URL: uniswap.org/whitepaper.pdf

Adams H., Zinsmeister N., Salem M., Keefer R. and D. Robinson (2021) Uniswap v3 core. Tech. rep.

Angeris G. and T. Chitra (2020) "Improved Price Oracles", arXiv preprint arXiv:2003.10001.

Angeris G., Evans A. and T. Chitra (2020) "When does the tail wag the dog? Curvature and market making", arXiv preprint arXiv:2012.08040.

Angeris G., Kao H.T., Chiang R., Noyes C. and T. Chitra (2019) "An analysis of Uniswap markets", arXiv preprint arXiv:1911.03380.

Bergault P., Bertucci L., Bouba D. and G. Olivier (2022) "Automated Market Makers: Mean-Variance Analysis of LPs Payoffs and Design of Pricing Functions", arXiv preprint arXiv:2212.00336.

Buterin V. (2017) "On Path Independence", post available on author's webpage at vitalik.ca/general/.

Capponi A. and R. Jia, R. (2021) "The adoption of blockchain-based decentralized exchanges", arXiv preprint arXiv:2103.08842.

Egorov M (2019) "Stableswap-efficient mechanism for stablecoin liquidity", available at curve.fi/.

Evans A. (2020) "Liquidity Provider Returns in Geometric Mean Markets", arXiv preprint arxiv.2006.08806.

Fan Z., Marmolejo-Cossío F., Altschuler B., Sun H., Wang X. and D. C. Parkes (2022) "Differential liquidity provision in Uniswap v3 and implications

for contract design", arXiv preprint arXiv:2204.00464.

Goyal M., and Ramseyer G., Goel A. and D. Mazières (2022) "Finding the Right Curve: Optimal Design of Constant Function Market Makers", arXiv preprint 2212.03340.

Kyle A. S. (1985) "Continuous Auctions and Insider Trading", *Econometrica*, 53 (6), 1315-1335.

Martinelli F. and N. Mushegian (2019) Balancer whitepaper Tech. rep. URL: balancer.fi/whitepaper.pdf

Wu M. and W. McTighe (2022) "Constant power root market makers", arXiv preprint arXiv:2205.07452.

Neuder M., Rao R., Moroz D.J. and D.C. Parkes (2021) "Strategic liquidity provision in uniswap v3", arXiv preprint arXiv:2106.12033.

Port A. and N. Tiruviluamala (2022) "Mixing Constant Sum and Constant Product Market Makers", arXiv preprint arXiv:2203.12123.

Samuelson P.A. (1969) "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *The Review of Economics and Statistics*, 51, 239-246.

Schär F. (2021) "Decentralized Finance: On Blockchain- and Smart Contract-Based Financial Markets", Federal Reserve Bank of St. Louis Review, Second Quarter, pp. 153-74.

Strotz, R. H. (1955) "Myopia and Inconsistency in Dynamic Utility Maximization", *The Review of Economic Studies*, 23 (3), 165–180.

Figure 1: *Continuous pricing curve*

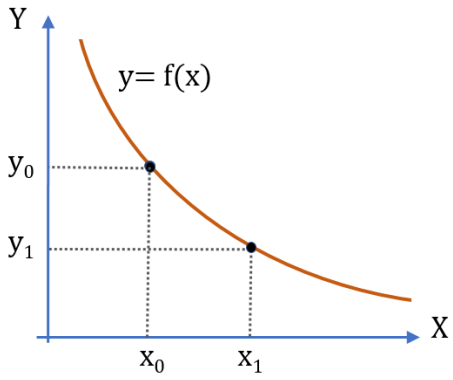

Note: The curve shows how much reserve in asset Y the AMM must hold given its reserve in asset X.

Figure 2: *The price as the marginal exchange rate*

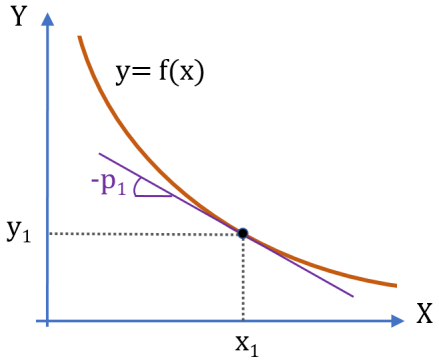

Note: The slope of the line tangent to the curve is the negative of the marginal exchange rate between the two assets.

Figure 3: *Piece-wise pricing curve*

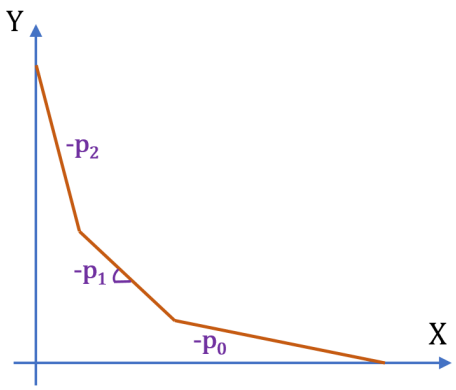

Note: The piece-wise curve indicates how much reserve in asset Y the AMM holds given its reserve in asset X

Figure 4: *Effects of a price decrease on the quantities held by the AMM*

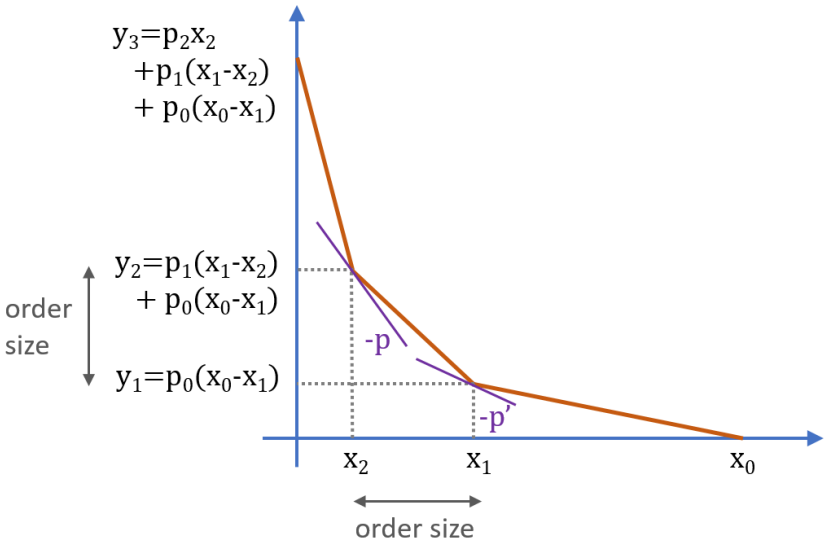

Note: The graphic indicates by how much the reserve in asset Y decreases and the reserve in asset X increases when the market price moves below p_1, from p to p'.

Figure 5: *Example of price sequences*

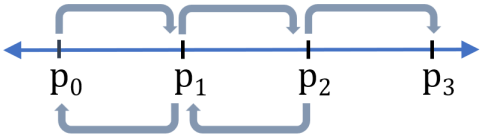

Figure 6: *Value functions with a three-price set*

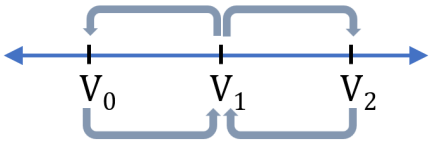

Figure 7: *Distribution of order sizes over the set of execution prices*

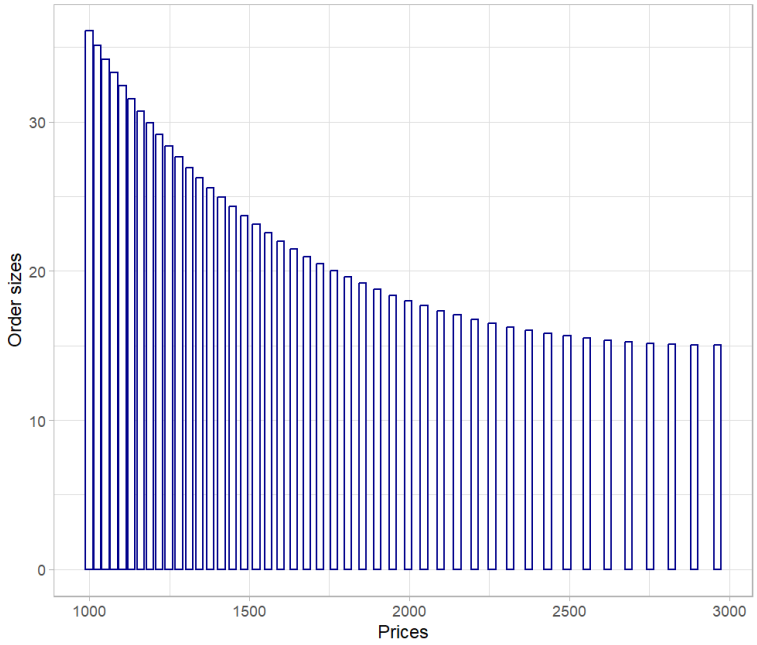

Parameter values (baseline): $\sigma = 2$, $\beta = 0.999$, $t = 0$ and $\theta = 1/3$.

Figure 8: *AMM's pricing curve (baseline)*

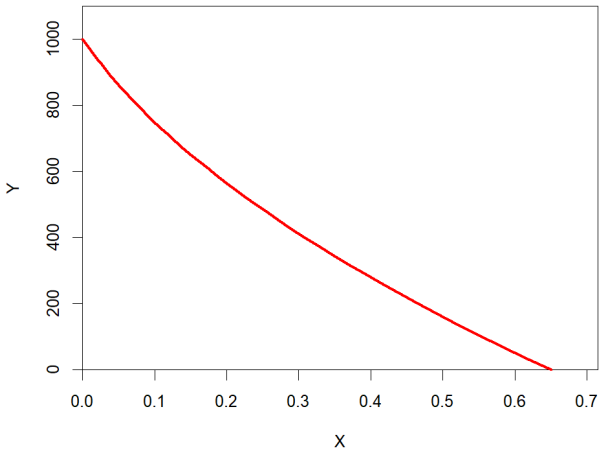

Notes. The pricing curve indicates how much the market maker holds asset Y for every quantity held in asset X when exchanging with traders.

Figure 9: Distribution of order sizes and relative risk aversion (RRA)

(a) *Low RRA* ($\sigma = 0.5$)

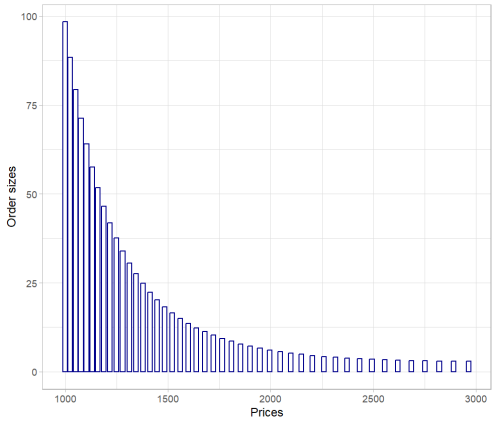

(b) *Baseline RRA* ($\sigma = 2$)

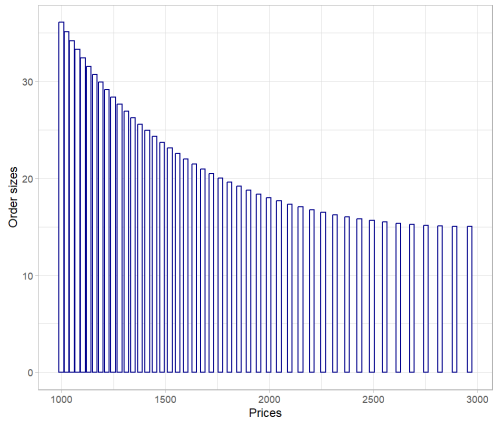

(c) *High RRA* ($\sigma = 5$)

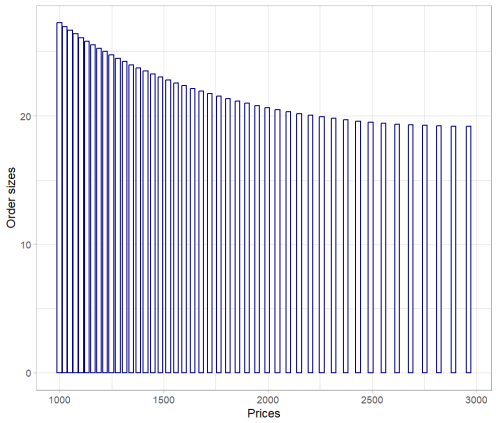

(d) *Very high RRA* ($\sigma = 20$)

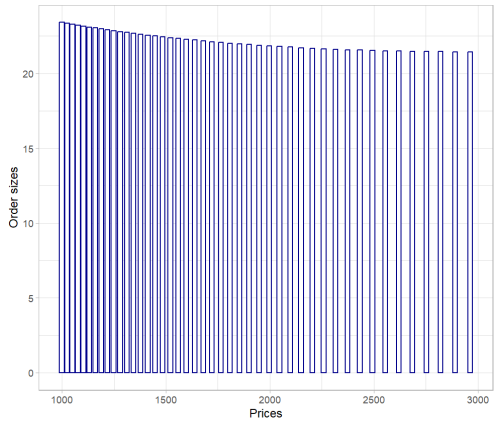

Parameter values (baseline): $\beta = 0.999$, $t = 0$ and $\theta = 1/3$.

Figure 10: Distribution of order sizes and time discounting (TD)

(a) *Extremely low TD* ($\beta = 0.99999$)

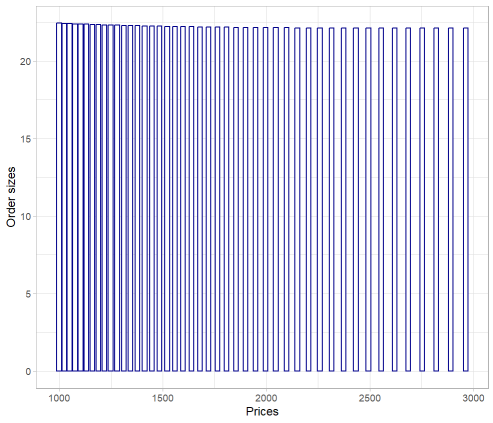

(b) *Low TD* ($\beta = 0.9999$)

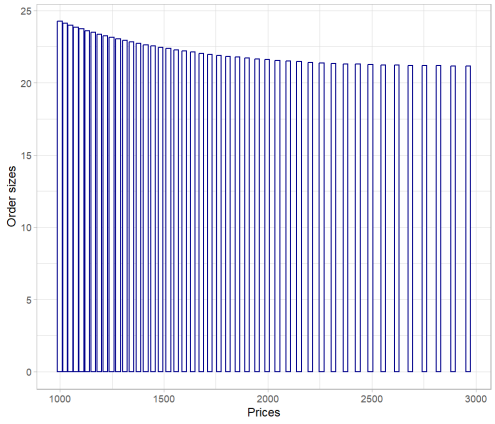

(c) *Baseline TD* ($\beta = 0.999$)

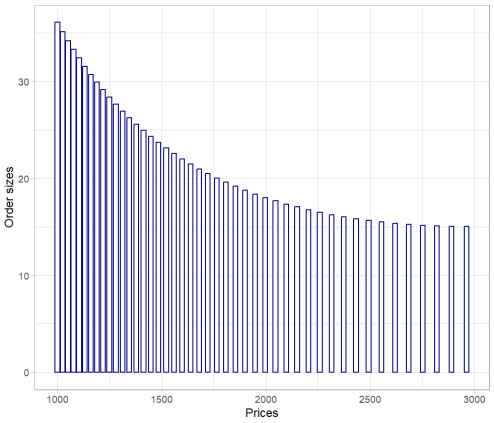

(d) *High TD* ($\beta = 0.99$)

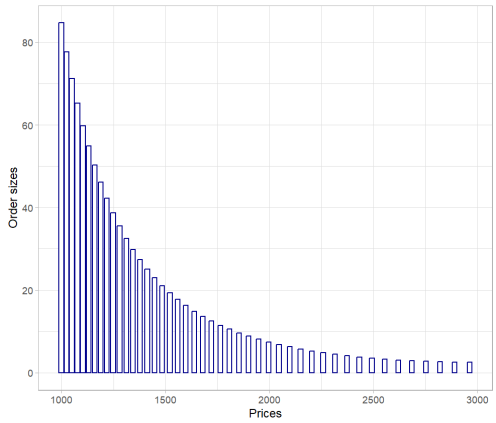

Parameter values (baseline): $\sigma = 2$, $t = 0$ and $\theta = 1/3$.

Figure 11: Distribution of order sizes and reference state

(a) *mid-state* ($t = 23$) (b) *highest state* ($t = 46$)

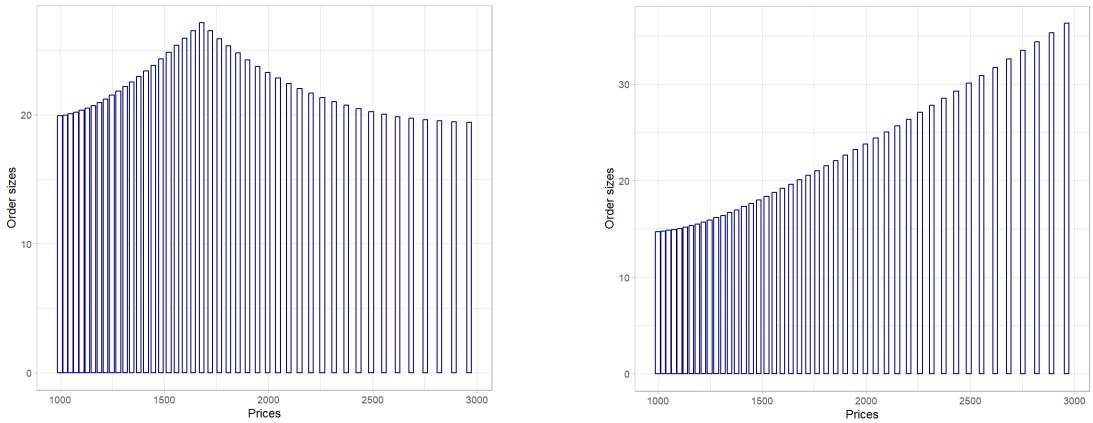

Note: The reference state is the state at which the value functions are maximized. Parameter values (baseline): $\sigma = 2$, $\beta = 0.999$ and $\theta = 1/3$.

Figure 12: Distribution of order sizes and probability of transiting to a different state

(a) *Low probability* ($\theta = 0.2$) (b) *High probability* ($\theta = 0.5$)

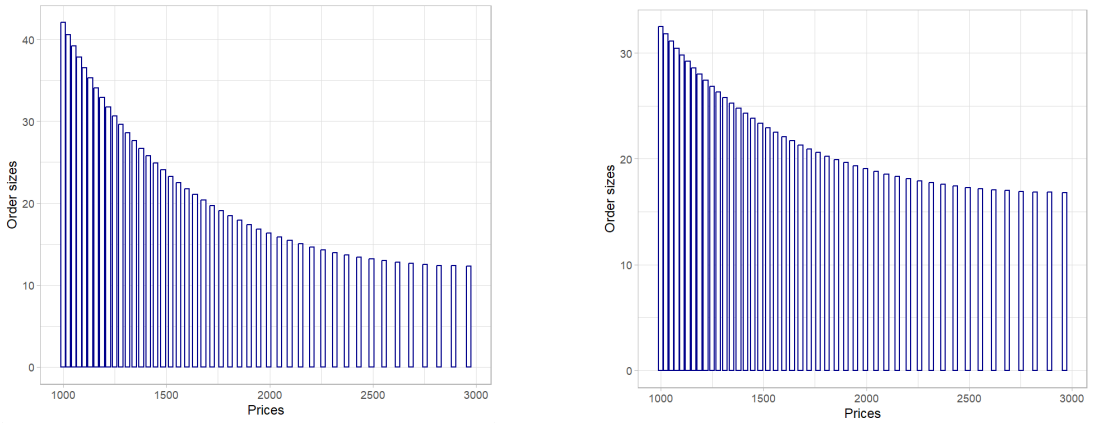

Note: The transition probability θ is the probability of transiting to a higher state ($\Pi[n, n-1]$) or lower state ($\Pi[n, n+1]$). Parameter values (baseline): $\sigma = 2$, $\beta = 0.999$ and $t = 0$.