Learning from DeFi: Would Automated Market Makers Improve Equity Trading?

Katya Malinova^{*} Andreas Park[†]

May 1, 2023 –Very preliminary, please do note cite–

Abstract

We investigate the potential for automated market makers (AMMs) to be economically viable in and improve traditional financial markets. AMMs are a new type of trading institution that have emerged in the world of crypto-assets and process a significant portion of global crypto trading volume. The current trend of tokenizing assets, the legitimization of crypto-token issuance via the EU's MiCA regulation, and the push by the S.E.C. to change the trading of retail orders presents an opportunity to consider AMMs for traditional markets. Our approach is to determine the parameters that would allow liquidity providers to profitably contribute to an AMM and calculate, based on U.S. equity trading data, if liquidity demanders would benefit from using the AMM for these parameters. Our analysis suggests that properly designed AMMs could save U.S. investors about 30% of annual transaction costs. The source for these savings is twofold: AMMs allow better risk sharing for liquidity providers and they use locked-up capital that otherwise sits idly at brokerages. The introduction of AMMs in traditional markets could particularly improve the liquidity and trading cost issues faced by small firms, allowing them to attract more investors and capital.

^{*}McMaster University malinovk@mcmaster.ca. Financial support from the Social Sciences and Humanities Research Council is gratefully acknowledged.

[†]University of Toronto andreas.park@rotman.utoronto.ca. Financial support from the *Global Risk Institute* and the *Social Sciences and Humanities Research Council* is gratefully acknowledged.

Automated market makers (AMMs) are a new type of trading institution that has no parallel in traditional financial markets and already processes 10-15% of worldwide trading volume in crypto-assets despite the learning curve required in using them. Although AMMs were designed for and run on blockchains like Ethereum, they could also be built on a traditional financial infrastructure. We ask a hypothetical question: can this type of system be economically viable and improve traditional markets?

Our approach is straightforward: we first determine conceptually how liquidity providers can profitably contribute to an AMM. We then take this concept to the data and estimate the applicable parameters using U.S. equity trading data. Then, we calculate if liquidity demanders would benefit from using the AMM for these parameters and compute the trading cost savings. Our analysis indicates that properly designed AMMs could save U.S. investors about 30% of annual transactions costs.

AMMs are a genuinely novel idea that have been developed in the wake of the proliferation of public blockchains. They allow continuous trading based on passive, pooled liquidity provision and shared adverse selection risk. By design, these systems are suitable for a broad swath of market participants to provide liquidity and require no specialized skills or equipment, allowing asset users to re-use the capital that is otherwise sitting idly in a brokerage account.

One reason for considering a new market mechanism like an Automated Market Maker (AMM) is that traditional public markets work well for large firms but not as well for smaller ones. Large firms have significant analyst followings, liquid markets, and low trading costs. Small firms often face the same regulatory and compliance costs as large firms, but have significantly worse liquidity. For instance, in our data of U.S. publicly listed firms, the smallest 10% firms have 59 times higher trading costs than the largest firms in 2019, making it difficult for them to attract investors and raise capital. Small firms do not even at all trade on 9% of trading days.

One of the key innovations of an AMM is the approach to liquidity provision, which is to attract existing asset owners as liquidity suppliers by providing them with a flow income from trading fees. An AMM is a liquidity pool into which owners deposit their assets. Other investors can trade against this pool, and risk and fee income are pro-rated among the depositors. When properly designed, an AMM deposit is an alternative to a buy-and-hold strategy that allows long-term investors, in particular, unsophisticated ones, to earn an income from trading fees while retaining exposure to their asset of choice.

The premise of all liquidity provision is that gains from balanced order flow outweigh losses from unbalanced flow. In a traditional market, intermediaries must constantly monitor the markets to make sure that they are not on the wrong side of the market, which requires massive investments in trading technology. Furthermore, existing owners play no role in liquidity provision because they only trade when they want to change their position. With an AMM, risk and income are shared, and risk sharing is a fundamental principle for increased economic efficiency. The question is: can this be made to work?

In this paper, we first derive closed-form formulae to express rewards for liquidity provision and profits for liquidity takers based on trading variables in Automated Market Maker (AMM) systems. For fixed trading volume and price change, a depositor's return depends on two key variables: the level of the fee, measured in basis points of the transaction value, and the size of the liquidity pool. To ensure a useful comparison across assets, we measure the size of an asset pool by a fraction of the asset's market capitalization.

Let α denote the fraction of an asset's market capitalization that is deposited as (passive) liquidity. There are two opposing forces: more deposited liquidity α makes it cheaper to trade but each liquidity provider gets less because the same fee gets distributed among more of them. Let $\overline{\alpha}$ denote the share of the deposited market capitalization such that liquidity providers break even, and let $\underline{\alpha}$ denote the share for which AMMs are cheaper than traditional markets for liquidity demanders. If $\overline{\alpha} > \underline{\alpha}$, then theoretically there exist an economically viable AMM configuration that improves welfare.

At the end of the day, however, it is an empirical question whether AMMs can be economically viable. We base our investigation on all U.S.-listed common stocks in CRSP and TAQ from 2014-2022, omitting the time period of the S.E.C. tick pilot because the design of the pilot might distort our findings (and likely favor AMMs). We study many combinations of AMM fees, assumptions on volume, trade sizes, and AMM income windows. We present our results for a reasonable set of assumptions on AMM usage.

Namely, we assume that, as is common practice today, there is an opening session prior to continuous trading that determines the opening price, which serves as the benchmark for AMM deposits. This process ensures that users are not exposed to overnight risk (with such risk, users may not be willing to provide liquidity to the pool). Therefore, the return risk that users are exposed to stems from post-open to close returns.¹ The pool sets the trading fee, and prospective liquidity providers make their determination based on what they would have earned on the previous trading day. They then compete to supply liquidity until the amount deposited reaches the break-even point based on the previous day's values.

The amount of liquidity provided is positively associated with the AMM fee. This creates two opposing forces for liquidity demanders: they benefit from a higher fee because liquidity provider supply more liquidity. And they suffer from paying a higher fee. Theoretically there is a fee that maximizes liquidity provider benefits (a point also made by Hasbrouck, Rivera, and Saleh (2022)). For small firms, the optimal fee is about 31 basis points, for large firms it is only 0.8 basis points. We perform most of our analysis assuming that this fee is deployed and updated daily based on the previous day's trading characteristics.

We then calculate the hypothetical contemporaneous benefit to liquidity demanders on the current day and compare it to the observed costs of the traditional market. This benefit can be positive or negative, for instance, depending on whether spreads are lower or higher than on the previous day.

Based on data from all U.S. publicly listed stocks, there is a strict welfare gain for 82% of day and stock observations. Aggregate benefits across all assets amount to more than 30% of annual trading costs, around \$6.5 billion for 2019, \$12.5 billion in 2020, and \$15.0 billion in 2021. At the optimal fee level, between 3-8% of a stock's market capitalization would need to be made available as liquidity, but we also compute that in practice, only 2-5% of this amount needs to be kept at the ready. The proportional

¹We have also performed the analysis using different time horizons and deposit horizons, and the results are similar.

benefit is largest for small firms, the dollar amounts (due to higher volumes) are highest for large firms. Overall, our analysis indicates that properly designed AMMs could improve the trading environment for investors, providing them with increased liquidity and substantially reduced trading costs.

The reader may wonder, however, about the source for the massive savings that we project. These are not a sack of magic beans, but have a simple explanation: shared risk and repurposing of idle capital. Currently, most shares sit idly at brokerages, except for the small fraction of investors that lend shares to short sellers. Even this lending activity is often bilateral, which makes it expensive and cumbersome to arrange. AMMs allow investors to systematically put their capital to use and earn extra income, while providing adequate risk sharing. A last point to mention is that our analysis is conservative in the sense that we are holding volumes fixed. Presumably, significantly lower costs of trading should attract further trading activity.

In conclusion, if properly implemented, an AMM trading arrangement has the potential to benefit market participants, particularly for smaller, less liquid securities.

Related Literature. There are several theoretical papers that study AMMs. Lehar and Parlour (2021) and Aoyagi and Ito (2021) compare AMMs with limit order books under asymmetric information. Lehar and Parlour (2021) study AMMs and limit order books in isolation and show that for many parametric configurations, investors prefer AMMs over the limit order market. Aoyagi and Ito (2021) model the co-existence of a centralized exchange and an automated market maker, and they study the venue choice of informed traders. Their main finding is that the informed traders react non-monotonically to changes in the risky asset's volatility, and that reaction causes non-monotonic shifts in liquidity supply. In related work, Aoyagi (2022) models the AMM liquidity provision decision under asymmetric information.

Capponi and Jia (2021) examine the impact of price volatility, caused by trading on centralized exchanges, on the welfare of liquidity providers in automated swap exchanges for a broad set of functions (twice-continuously differentiable, and convex). They identify the conditions for a breakdown of liquidity supply in the automated system, and they show that a pricing curve with larger convexity reduces arbitrage rents, but also decreases trading activities. Park (2021) studies properties of the standard AMM pricing function compared to a limit order book-type pricing.

Barbon and Ranaldo (2022) compare the liquidity for crypto-asset trading in decentralized and centralized exchanges empirically and argue that DEX prices are less efficient. Hasbrouck, Rivera, and Saleh (2022) study a model of DEX liquidity provision and the relationship of fees, liquidity, and volume. They show that an increase in trading fees may attract additional liquidity which in turn attracts more volume.

Our theoretical description is most closely related to Aoyagi and Ito (2021) and Park (2021), where we use the same model as the latter. The description of some features of liquidity provision is from Barbon and Ranaldo (2022).

I. Background on Automated Market Makers

An automated market maker establishes a liquidity pool to which tokens holders can contribute pairs of tokens, usually in return for fee payments. To remove one type of token from the pool, a liquidity demander has to deposit the other type of token. The exchange rate of tokens is determined by a so-called "bonding curve." A liquidity pool contains deposits of a units of the asset A-tokens and c units of cash. The pool's liquidity is defined by a function $L : \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}$. The rule is reminiscent of a Cobb-Douglas production function, where the good produced is liquidity and the inputs are the cash and asset. Combinations of a and c that describe the same value L(a, c) are on an iso-illiquidity curve. The ratio c/a is the price for a marginal unit of the asset.

There are two operations for liquidity pools. The first is the addition (or withdrawal) of liquidity. The rule for this operation is that the marginal price does not change. In a user adds/withdraws both cash and assets, $\Delta a, \Delta c$ with sign(Δa) = sign(Δc). There,

$$L(a,c) \leq L(a + \Delta a, c + \Delta c))$$
 for $\operatorname{sign}(\Delta a) = \operatorname{sign}(\Delta c) \geq 0$

That is, if a user wants to add/withdraw Δa units of liquidity they would need to add/withdraw Δc such that $\Delta c/\Delta a = c/a$. Functionally, when a user adds liquidity to a pool, they receive a receipt token that signifies their fractional ownership of the pool at the time of their deposit and that they can exchange at a future time when they reclaim their assets and cash from the pool.

The second operation are purchases of the asset from the pool (or sales of the asset to the pool) where $\operatorname{sign}(\Delta a) = \operatorname{sign}(\Delta c)$ These operations keep the liquidity constant. Specifically, to purchase quantity q of the asset A from the pool, a user deposits (or pays) a quantity $\Delta c(q)$ of cash. After the trade, the pool contains a - q of the asset and $c + \Delta c(q)$ cash. The *liquidity invariance condition* requires that the exchange amount $\Delta c(q|a,c)$ for function L(a,c) is such that for any feasible quantity q < a, $\Delta c(q|a,c)$ is such that

$$L(a,c) = L(a-q, c + \Delta c(q|a,c)).$$

Although there are several theoretically possible invariance rules, the most commonly used one for blockchain-based AMMs is the constant product pricing rule where $L(a, c) = a \cdot c$, and we will use this rule in the remainder of the paper. When trading quantity q, $\Delta c(q)$ satisfies $ac = (a-q)(c + \Delta c(q))$ for any q < a and $-\Delta c(q) < c$. Then

$$\Delta c(q|a,c) = q \cdot \frac{c}{a-q} \quad \text{and } p(q) = \frac{c}{a-q}, \tag{1}$$

so that $\Delta c(q)$ is the cost of the order and p(q) is the per unit price when trading q. Figure 1 illustrates the process of finding a price, Figure 2 plots the price function p(q) for an example with p(0) = 10 for various levels of liquidity. The latter figure shows that additions and withdrawals of liquidity change the curvature of the price curve for the same marginal price.

II. The Liquidity Supply and Demand Decision

A. Costs of Liquidity Provision

Liquidity supply is inherently risky because the fundamental value of an asset can change while the liquidity provider holds a position. The principle of liquidity provision is, therefore, that the liquidity supplier must earn more on trades that do not move the market than what they lose on those that move the market.

Very loosely, let the buyer-initiated volume be B and the seller initiated volume Sand, for the sake of the argument, assume B > S. Usually, when there are more buys than sales, the price rises, say, from p_0 to p_t , and we go with the assumption that after B buys and S sells, "the price is right." A liquidity providers earns the a fee on the total volume B + S, but loses on the part of the volume that moves the price, B - S.

In a traditional market, the market maker's fee is the bid-ask spread, and their loss is the price impact. In an AMM, the pricing function is continuous in the quantity. As Park (2021) shows, for reasonable assumptions on the pricing function (satisfied by the constant product function that we study here), prices are not regret free so that liquidity providers always lose when the fundamental of the asset moves. Neither do they allow liquidity providers to break even on average, e.g., by gaining from uninformed what lose by trading with informed traders. To compensate liquidity providers, therefore, the system needs to include an explicit compensation scheme. We assume here that compensation comes in the form of a fee F that is levied on the dollar volume of each trade. The fee is to be paid in in cash and saved in a separate account.

The specific revenue can be computed as follows. The price impact loss is the difference between the cash $\Delta c(B-S)$ that the liquidity provider receives for B-S, and the assessed value of the asset after the trade at the marginal price, $p_t \cdot (B-S)$. Consequently, for liquidity provision to be viable, it must hold that

$$F \times (B+S) + \Delta c(B-S) - p_t(B-S) \ge 0.$$
⁽²⁾

Note that the second part of the above, $\Delta c(B - S) - p_1(B - S)$, is the same as the net cost/benefit of providing liquidity relative to a buy-and-hold strategy. Namely, if a liquidity provider has a assets and c cash, these assets are worth $p_t a + c$ for buy-and-hold. When supplying liquidity, the provider has a - (B - S) assets and $c + \Delta c(B - S)$

cash, which is worth $p_1(a - (B - S)) + c + \Delta c(B - S)$. The difference of this amount and the value of buy-and-hold is the specified amount.

Crucially, in real markets, liquidity providers do not hold a position but rather aim for a net-zero position. They need to make investments in technology to accomplish this goal, and they need to borrow and lend assets and cash to cover short-term unmatched positions. In contrast, AMM liquidity providers need to own the risky asset before they can contribute it to a liquidity pool.² Therefore, for AMM liquidity pools, the suitable benchmark actually is a buy-and-hold strategy of the deposited assets. In what follows, we present a formal model based on Park (2021).

B. Model Primitives

There are two assets, risky A and safe C, the marginal price for A measured in Cat the beginning of trading is $p_0 = V_0$, where V_0 is the fundamental value of one unit of the asset. Liquidity providers make available their assets and cash for the liquidity pool based on the specifics of the pricing function, and they commit to provide their liquidity for a fixed time horizon. At the end of the time horizon, the fundamental of the asset is V_t , and we assume that arbitrageurs move the price so that $p_t = V_t$. The corresponding return $R = V_t/V_0 = p_t/p_0$ with $R \in [0, \infty)$ follows a distribution $\Phi : \mathbb{R}_+ \to [0, 1]$ with continuous density ϕ . Arbitrageurs will add or remove the asset in exchange for cash from the liquidity pool such that the marginal price at the end of the period is p_t . We use $p_t(R)$ to signify the marginal price that pertains for return R.

When trading with the liquidity pool, liquidity demanders pay fee which is an

²Lehar and Parlour (2021) document that withdrawals of liquidity are rare, and that liquidity providers appear to (buy-and) hold their assets in liquidity pools.

exogenous parameter of the AMM protocol. To simplify the exposition, we assume that fee F is collected on the dollar volume of transactions and collected separately from the liquidity pool.³ Therefore, when a liquidity demander buys q of the A-tokens, the liquidity suppliers receive a payment of $F \cdot \Delta c(q) + \Delta c(q)$ of cash.

Fees accrue for all trades, and to simplify the analysis, we assume that in addition to a price-moving arbitrage trade, during the liquidity provision horizon (say, a day), there is a volume V that is perfectly balanced.⁴ Since fees are collected based on the dollar volume, the order of trades affects the fee income. For instance, when a buyvolume of V/2 is followed by a sell-volume of V/2 assets, then the average price and thus average fees are higher than if the trades occurred in reverse order. Computing the expected income for arbitrary histories in continuous time is beyond the scope of this paper. Therefore, we simplify the analysis and assume that fees for volume V accrue at the marginal price p_0 . For an arbitrage trade of q^* , total fees collected are

$$F\Delta c(q^*) + F p_0 V. \tag{3}$$

³Many blockchain-based implementations of AMMs have a slightly more complex process that affects the value of the assets in the pool. That approach is best explained by example. To buy Q tokens, the buyer first deposits P(Q) into the contract and then extracts a quantity \tilde{Q} . The fees are subtracted from the amount P(Q) prior to their deposit so that the buyer formally receives $\tilde{Q} < Q$. Moreover, users always pay in the token that they deposit. Therefore, in practice, fees accrue separately per token. See Lehar and Parlour (2021) for details.

⁴This is a benign simplification because I assume that the arbitrageur trades such that the marginal price at the end of the period is the "correct" price p_t . Suppose noise volume from buying and selling differs. Let q^* denote the quantity that would move the marginal price from p_0 to p_t . Let noise buy volume be b and noise sell volume s and without loss of generality b > s and $p_t > p_0$. Then after observing b buys and s sells, the arbitrageur would trade $\tilde{q} = q^* - (b - s)$. After this trade, the marginal price is p_t . For the liquidity provider's fee income, it does not matter whether the trades stems from a noise traders or arbitrageurs. I therefore assume that in expectation, b = s = V/2 so that total expected volume is V.

The Role of Fees. As Park (2021) shows, the fee is, in fact, necessary because the constant product pricing function does not provide adequate compensation to liquidity providers. Compensation for liquidity providers in an AMM therefore differs from a limit order market. Suppose one trader adds q assets and another withdraws q. Since the assets in the pool after these two transactions are the same as at the beginning, liquidity invariance ensures that the same holds for cash. In other words, the cash flows off-set one another and liquidity providers earn no income from this off-setting order flow. This contrasts limit order books: suppose the midpoint of the order book at time t is m_t , the bid and ask prices are respectively $m_t - s$, $m_t + s$, and after a trade at t, the new midpoint is the last price. When there is a buy followed by a sale, the cash flows to liquidity providers are $m_t + s - (m_{t+1} - s) = m_t + s - (m_t + s - s) = s$. In other words, in a limit order book, the liquidity provider earns the spread.

C. The Decision Problem of Liquidity Suppliers

Liquidity providers deposit quantities a and c into the pool, with $c = ap_0$; the slope of the pricing function is determined by the constant product rule. We use $d = c + ap_0$ to describe the cash value of the initial deposit. Note that $d = 2c = 2ap_0$. Since liquidity supply is competitive, liquidity providers break-even in expectation. Computationally, it is easiest to determine the zero expected *return* condition

$$\frac{1}{d}\left(\int_0^\infty (\Delta c(q^*) - p^* p_t(R) + F \cdot \Delta c(q^*)) \ \phi(R)dR + F p_0 V\right) = 0. \tag{4}$$

The critical parameter for liquidity provision is the initial deposit d. To determine the liquidity provision decision, we first determine the loss that would result for a given change of the fundamental to a value V_t , where without loss of generality $V_t > V_0$. We can then compute the expected return over all possible value changes.⁵

When the change from V_0 to V_t occurs, arbitrageurs trade against the pool to buy the underpriced asset for as long as it's profitable, i.e., until the marginal price $p_t = V_t$. After the price moving trade of q^* , the pool contains $a_t = a_0 - q^*$ of the A token and $c + \Delta c(q^*)$ cash. For given $p_t = V_t$, the price moving quantity q^* thus satisfies

$$p_t = \frac{c + \Delta c(q^*)}{a - q^*}.$$
(5)

Constant product pricing is defined so that when trading q, quantity $\Delta c(q)$ must satisfy $ac = (a - q)(c + \Delta c(q))$. Therefore, the liquidity demander's cost of buying q is $\Delta c(q) = qc/(a - q)$. Using this functional form in (5) yields

$$p_t = \frac{c + q^* c / (a - q^*)}{a - q^*}.$$

Solving for q^*

$$q^* = a - \sqrt{\frac{ac}{p_t}} = a \left(1 - \sqrt{\frac{p_0}{p_t}}\right) = a \left(1 - \sqrt{R^{-1}}\right),\tag{6}$$

where the penultimate relation obtains because $c = ap_0$, and the last relation uses $R = p_t/p_0$. To compute the net payoff $\Delta c(q^*) - p_t q^*$, we express each component separately as follows. First, we divide the first term, $\Delta c(q^*)$, by 2c(=d) and simplify

⁵Barbon and Ranaldo (2022) use a similar formulation in their empirical work.

as follows:

$$\frac{\Delta c(q^*)}{2c} = \frac{1}{2c} \cdot \frac{q^*c}{a-q^*} = \frac{1}{2} \frac{q^*}{a-q^*} = \frac{1}{2} (\sqrt{R}-1), \tag{7}$$

where we substituted for q^* from (6) in the last step. Second, we divide $p_t q^*$ by $2ap_0(=d)$

$$-\frac{p_t q^*}{2ap_0} = -\frac{R}{2a} \cdot \left(a\left(1 - \sqrt{R^{-1}}\right)\right) = \frac{1}{2}(\sqrt{R} - R).$$
(8)

Combining the terms (7) and (8) yields the expression for the relative loss that results as a function of the gross return R

ILLRAS(R) =
$$\sqrt{R} - \frac{1}{2}(1+R)$$
, (9)

where we use the acronym ILLRAS as an abbreviation for the *Incremental Loss from* Long-Run Adverse Selection. Finally, we substitute (6) to compute the fee return

$$F \frac{\Delta c(q^*)}{d} + F \frac{p_0 V}{d} = \frac{F}{2} |\sqrt{R} - 1| + \frac{F V}{2a}, \tag{10}$$

where we use absolute values to account for sales. Combining the losses due to changes in the fundamental (9) and the fee return (10) and taking expectations, we can express (4) as follows

$$\int_0^\infty \left(\sqrt{R} - \frac{1}{2}(1+R) + \frac{F}{2}|\sqrt{R} - 1|\right) \phi(R)dR + F\frac{V}{2a}$$

To compute the equilibrium, one computes expectations with respect to R and solves

for the equilibrium deposit a^* :

$$a^* = \frac{FV}{2} \underbrace{\left(\int_0^\infty \left(-\frac{F}{2} |\sqrt{R} - 1| - \sqrt{R} + \frac{1}{2} (1+R) \right) \phi(R) dR \right)^{-1}}_{=:C^{\mathsf{CP}}(\phi, F)^{-1}}.$$
 (11)

As the derivation shows, the value a^* is unique. In our empirical analysis it will be useful to split $C^{\mathsf{CP}}(\phi, F)$ into two components:

$$C^{\mathsf{CP}}(\phi, F) = -F \times E[|\sqrt{R} - 1|/2] - E[\text{ILLRAS}].$$

A key component is the loss caused by changes in the fundamental value, expression (9), or ILLRAS(R). Figure 3 plots this expression as a function of the gross return. The function is weakly negative irrespective of whether the price rises or falls because the liquidity provider always has less of the desirable item (cash or the asset) for positive and negative returns. It is also important to emphasize this quantity is incremental to any returns from holding the asset (i.e., the depositor keep exposure to the asset's returns). For instance, if the gross return is 0.9 (the asset lost 10% of its value), the incremental loss is 13 basis points or 0.13%. This implies that the liquidity provider makes a loss of 10.13% when the price drops by 10% If the gross return is 1.1 (the asset gained 10% value), the depositor gains 12 basis points or 0.12% less. This implies that the liquidity provider gains 9.88% when the price increases by 10%.

Equation (11) is the amount of liquidity provided that causes liquidity providers to break even in expectation. The goal of this paper is to examine liquidity provision for listed equities. We therefore describe it as the fraction α of the shares outstanding Sthat are deposited in the pool. Since the asset's market capitalization is $M := p_0 S$, we can also write $d = 2\alpha \cdot M$ and $a = \alpha S$. The we can define $\overline{\alpha}$ as the maximum fraction of the shares outstanding that makes liquidity providers break even. Using (11) it is

$$\overline{\alpha} = \min\left\{1, \frac{FV}{2S} \frac{1}{C^{\mathsf{CP}}(\phi, F)}\right\},\tag{12}$$

where $\overline{\alpha}$ cannot exceed 1, hence the min function.

D. Costs for Liquidity Demanders

Liquidity demanders face the constant product price function based on initial deposits $c^* + p_0 a^* = 2p_0 a^* = d$. We now determine the trading costs per unit traded, normalized by price p_0 . These costs coincide with the price impact for a size q order. For unit quantity q = 1, these costs can also be thought of as the implicit (half) bid-ask spread. Formally, for a size q order, since $p_0 = c/a$, the per unit relative cost is

$$p \text{-} imp(q) = \frac{\frac{c}{a-q} - \frac{c}{a}}{\frac{c}{a}} = \frac{q}{a-q} = \frac{q}{\alpha S - q},$$
(13)

where the last expression stems from us expressing the supplied liquidity a as a fraction of the shares outstanding, $a = \alpha S$. When trading on a AMM, the liquidity demander has to pay this price impact plus the trading fee F.

We are specifically interested in the value of α such that a liquidity demander is indifferent between this cost and trading on the regular market where they pay the spread. Therefore, we seek the value of α that solves

$$\frac{q}{\alpha S - q} + F = \sigma. \tag{14}$$

The solution to this problem is the minimum viable deposit rate

$$\underline{\alpha} = \max\left\{0, \min\left\{1, \frac{q}{S} \frac{1+\sigma-F}{\sigma-F}\right\}\right\}.$$
(15)

The value $\underline{\alpha}$ cannot exceed 1, hence the upper limit of 1. Likewise, conceptually $\underline{\alpha}$ cannot be negative, hence the lower bound of 0. We note that there is no constraint that ensures that $\underline{\alpha} < \overline{\alpha}$; whether this holds is an empirical question.

We further need to highlight that we are computing $\underline{\alpha}$ for a specific q. To accommodate a large q, deposited liquidity has to be large to compete with the spread. However, for a large q, the spread may no longer be the right point of comparison either because there is not enough depth in the market to accommodate a large trade. In our empirical assessment we choose a "sufficiently large" quantity in our computations; in untabulated work, we also ran the analysis using average institutional orders (these are in excess of 20K shares) and average time-weighted market depth. The insights are usually similar. Ultimately, our view is that the AMM should be attractive to an average-sized, non-informed trade, and we believe the presented results reflect this.

E. Optimal Fees

Both thresholds for liquidity provision, $\underline{\alpha}$ and $\overline{\alpha}$ are functions of the AMM fee F and increase in F: liquidity suppliers earn more for higher fees and can therefore "tolerate" more sharing, and liquidity demanders require more liquidity to offset the fees.

The next question relates to the monetary benefit. When trading on traditional markets, liquidity demanders pay bid-ask spreads, which add up to double-digit billion dollar trading costs annually. Trading in an AMM costs money, too. To determine the relative benefit, we proceed as follows. We assume that fees are set before liquidity deposits are made, that all information is public, that beliefs and risk assessments are identical, and that liquidity providers are risk-neutral. Assuming each liquidity provider is a price taker and risk neutral, they keep contributing until the pool contains the break-even amount $\overline{\alpha}$ of liquidity. Accordingly, any surplus goes to liquidity takers.

For a trade of size q, the liquidity demander's contemporaneous benefit per dollar traded on an AMM instead of an open market is

$$\pi = \sigma - \frac{q}{\overline{\alpha}S - q} - F,\tag{16}$$

There are two opposing forces for the liquidity demanders's benefit from using the AMM with regards to the fee F: Higher fees make liquidity providers more willing to supply liquidity. This effect enters via $\overline{\alpha}$, which increases in the fee F. Higher fees also mean that liquidity demanders have to pay more. Hasbrouck, Rivera, and Saleh (2022) describe a similar mechanism in their paper.

The function π has a local maximum in fees at

$$F^{\pi} = \frac{1}{E[|\sqrt{R} - 1|/2] + V} \left(-2q \ E[\text{ILLRAS}] + \sqrt{-2qV \ E[\text{ILLRAS}]}\right).$$
(17)

Note that using the AMM is not necessarily an equilibrium because it is possible that for $F = F^{\pi}$, $\underline{\alpha} > \overline{\alpha}$. It is an empirical question whether AMM trading is viable.

III. Empirical Background

A. Assumptions for Computations

Our goal is to determine whether it is possible to implement AMMs profitably, using observations from existing markets. We do not claim that our approach is optimal — rather, it is a simply heuristic that we believe has intuitive appeal. If this simple approach delivers a strong indication that AMMs can be used profitably then this is a sufficient condition that a more complex optimization would deliver even stronger results.

First, we simplify $C^{\mathsf{CP}}(\phi, F)$ in two ways. First, we ignore the first term, $-F \times E[|\sqrt{R}-1|/2]$, i.e., we do not explicitly account for the fees that liquidity providers earn on the price moving trade q^* . Second, we do not compute the distribution and expectation but instead use the previous day's realized open-to-close return.⁶

The process we assume is as follows:

- 1. At the end of each day, all funds are withdrawn from the pool and added back to the liquidity providers' accounts.
- 2. The market has an opening auction. Its market clearing price determines the ratio of cash to assets for AMM deposits.⁷
- 3. The AMM operator computes, based on yesterday's deposits and volume the welfare optimizing fee F^{π} for the day.

 $^{^{6}}$ In as of now untabulated work that we will include in a future version of this paper, we also estimate the distribution of returns, and compute the respective expected values. The results are similar.

⁷This process ensures that users face no overnight risk from a liquidity pool deposit.

- 4. Liquidity demanders compute the previous day's open-to-close return, where we take the return from the first trade after the open to the last trade before the closing auctions. From this, they determine ILLRAS, i.e., they base their computation of ILLRAS on the point estimate of *R*.
- 5. Based on ILLRAS, F^{π} , and the previous day's dollar volume divided by factor two (see the discussion below), they determine $\overline{\alpha}$.
- 6. Liquidity providers also compute $\underline{\alpha}$ based on the previous day's optimal F^{π} , quoted spread σ , and average order size plus twice the standard deviation computed over 100 days.
- 7. If and only if $\overline{\alpha} > \underline{\alpha}$, they make a deposit of $2\overline{\alpha}M$ (cash and shares), where M is based on the opening price.

We are interested in the number of days for which the AMM is viable, $\overline{\alpha} > \underline{\alpha}$ and we are interested in the payoffs (sign and size) when the AMM is viable. Notably, on the "day of", the realization of π can be positive or negative (e.g., because spreads may be lower than on the previous day). We compute the payoffs to liquidity providers (who should break-even on average).

When looking at the data, we first perform the above process without optimizing over the fee (Step 3.) but for a range of fees F = 1, ..., 30. This will provide some insights into the importance of the fee for the variables. We then study what happens for the optimal fee F.

B. Data Sources and Sample

We focus on the universe of U.S.-listed common shares that are in the TAQ database, but we are excluding any symbol with a suffix (this therefore excludes preferred shares and dual-class shares). Our sample period is January 1, 2014 to November 30, 2022. This time horizon covers the S.E.C.'s "Tick Pilot" which ran from October 2016 to October 2018; the pilot changed the minimum price increment for a number of securities from 1 to 5 cents, thereby artificially increasing the tick size for a subset of securities. To ensure that our work is not affected by sample selection concerns, we exclude the time horizon of the tick pilot from our data.

We obtain information on shares outstanding at the daily level from CRSP. Data related to intra-day trading is from the WRDS computed statistics, which are based on TAQ data. We specifically use the intra-day trading volume, the number of transactions, the volume weighted average price, time-weighted bid-ask spreads, the opening price, the first price after the open and the last price before the close, and we compute from the data provided the average trade size across all orders, the average retail size, and average institutional size (for orders larger than 20K shares). We also obtain the cancel-to-trade ratio, where available, from the S.E.C.'s MIDAS database.

C. Trading Fees

In traditional markets, traders or their brokers have to pay a fee when submitting a marketable order. Our view is that the trading venue that operates an AMM would also charge a fee, and for simplicity we assume that it would be the same fee. The portion of the taker fee that does not go to exchange but to the liquidity provider is the maker fee/rebate. Therefore, the cost of a marketable order modulo the fee paid to the exchange is the spread plus the absolute value of the maker rebate. For the lower bound $\underline{\alpha}$, we would replace the spread σ with spread plus fees mr, $\sigma + mr$.

Choosing the right fee, however, is not trivial. Moreover, it is not clear to what extent this fee affect a trader's decision because this fee is often included in the commission. Another complication is that different brokers pay different fees, and that the marketplaces charge a large variety of different fees that usually depend on the submitting brokerage and may vary by security. On the NYSE, as of December 10, 2022, the "normal" maker rebate is \$0.0020 per share. On NASDAQ, the maker rebate varies widely and ranges from \$0.0020 to \$0.0030 per share. We want to be cautious in our analysis and not overstate our findings; therefore, we use the lower-end rebate of \$0.0020 per share.

D. Key Variables

Trading Volume and Intermediation. A key variable in our analysis is the amount of daily volume because it determines the fee income. In traditional markets, a portion of the volume is intermediated, because some traders interact with a professional liquidity provider who then offloads the position at a later stage. The extent of intermediation differs by security — usually larger firms have more intermediation than smaller ones. In an AMM, there is no intermediation. Ceteris paribus, volume in an AMM would therefore be weakly smaller than in an intermediated market. There are three volume scenarios:

1. If all trading is intermediated, the correct comparable volume would be half of

a day's volume.

- 2. If only the retail order flow is intermediated, then the correct comparable volume would be the daily volume minus the retail volume.
- 3. Many buy and sell retail orders offset one another. Most retail volume in U.S. markets flows through a small number of so-called wholesalers, and these have to offload only the unbalanced portion of the retail flow. Our third volume measure is therefore the daily volume minus the absolute volume of the difference between retail buy and sell volume.

In presenting our results, we focus on the first case which is most strongly biased *against* AMMs, i.e., we use *half the observed volume*. We also ran the analysis with the other volume measures and our results are qualitatively the same, albeit stronger in favor of AMMs.

WRDS provides several volume metrics to researchers that sometimes vary slightly: we use the daily maximum of these. We also use the WRDS-provided value for retail dollar-volume and institutional volume, where the latter is based on trades that are larger than 20K shares.

Order Size. Since there are sometimes very large orders that get processed in a specialized dark pool or OTC market, and AMM is unlikely to serve all order sizes better. At the same time, market orders in traditional markets may "walk the book" and trade at worse prices than the top-of-the book bid and ask prices; our data does not separately indicate the cost for such orders. A trader who splits a large order into

smaller pieces may be detected which increases the price of the total order, too. In an AMM, order splitting is not profitable by design.

Our approach is to base our analysis on averages and to allow for sizes larger than the average. Since the liquidity provision decision must be taken before trades arrive, we use long-run, 100-day averages and standard deviations thereof for trade sizes, where the average trade size is the ratio of volume to trades. As our decision order size, we use the average trade size plus two standard deviations of the 100-day average.

Trading Cost. Trades with an AMM are marketable orders. We therefore compare the AMM price impact in (13) plus fees to a measure of trading costs for a marketable order. The standard measures are the time-weighted bid-ask half-spread, the volume weighted effective half-spread, and the realized half-spread, which is the effective spread minus the price impact divided by factor 2. There are merits for each: the time weighted spread best describes the average costs that people face when they enter the market at a random time whereas the effective spread measures the costs at which people were willing to trade. In presenting our results, we focus on the quoted spread. Results with the effective spread are, however, similar.

The realized spread measures the take-home for intermediaries, subtracting the adverse selection caused by the trade. Since we measure adverse selection separately, the realized spread is not suitable for our purposes.

Horizons. AMMs have the premise of "stable" liquidity in the sense that depositors keep their assets in place for a long stretch of time to collect fees relative to a buy-and-hold strategy. Longer horizons usually mean higher variance in returns, but there are

also more fees that accumulate. Most importantly, there is overnight risk. The process described in Section A. eliminates this risk because it allows liquidity providers to reset the deposited quantities each day after the overnight information has been incorporated into the opening price. In practice, the process of liquidity provision could even be automated so that the liquidity provider would have to do any risk assessments and computations themselves. In untabulated results, we also performed the analysis for a four week (20 trading days) holding period. The findings are similar.

E. Key Metrics for Comparison

We are interested in three questions. First, we want to know whether an AMM is viable, that is, whether the minimum viable liquidity, $\underline{\alpha}$ is below the maximum viable liquidity $\overline{\alpha}$, $\underline{\alpha} < \overline{\alpha}$. In addition to the inequality, there are other "corner solutions" to consider: When $\underline{\alpha} = 0$, then liquidity demanders always prefer AMMs, when $\underline{\alpha} = 1$, no amount of AMM liquidity is sufficient for liquidity demanders to use the market. When $\overline{\alpha} = 0$, then liquidity suppliers would never be able to use an AMM profitably, when $\overline{\alpha} = 1$, then an AMM deposit always dominates a buy-and-hold strategy.

Our second variable of interest is the optimal fee, F^{π} . If on the previous day, $\underline{\alpha} < \overline{\alpha}$, then we assume liquidity providers will make a deposit, and we compute F^{π} , the optimal fee.

Third, the thresholds and the fee are computed based on past behavior. We compute the trading cost based on observed volume and we compute the trading costs advantage (or disadvantage) of an AMM relative to the regular market using observed volume and spreads. We thus determine the difference of the quoted spreads and the AMM price impact plus fees. Topics that we do not address. An implicit assumption in our analysis is that liquidity providers assume that, when introduced, the AMM provides strictly cheaper liquidity compared to the traditional market and would therefore attract all volume. There are three issues. First, some of the observed volume is intermediated. With an AMM, intermediation is most because AMMs rely on existing shareholders providing liquidity, and intermediaries are not long-term investors. Therefore, if markets would switch to AMMs entirely, intermediated volume would disappear.

The second issue is that AMMs would come into being only if they are cheaper. Usually, demand curves are downward sloping in trading costs and, therefore, lower costs should attract additional volume. As a check to see whether this is true, we compute the elasticity of volume relative to spreads, but we will not make assumptions regarding incremental volume that AMMs may "create."

The third issue is that some traders use limit orders to fill their positions, which allows, say, a buyer to buy at the bid rather than the ask. Arguably, for these traders, AMMs are more expensive, a caveat that we cannot avoid but that we share with most empirical analyses in the microstructure literature.

IV. Empirical Results

A. Animated Visuals

We begin by studying animated figures for fees $F = 1, ..., 30.^8$ All plots for variables of interest are against the stocks' market capitalizations. Figure 14 highlights the relationships among the cross-sectional averages for key variables. Larger stocks

 $^{^{8}\}mbox{Please}$ note that the animated graphics do not work for printouts — this is not Harry Potter.

are usually associated with larger volumes; small stocks have a higher concentration of trades among retail investors, and that large stocks have lower spreads. We also provide some info on the cancel-to-trade ratio, which indicates that it at best an imperfect measure of professional liquidity provision.

The level of the fee creates two opposing effects regarding the maximum and minimum viable liquidity thresholds, $\underline{\alpha}$ and $\overline{\alpha}$. First, a higher fee increases the income per dollar traded, which lowers $\overline{\alpha}$. Second, a higher fee makes it harder for the AMM to be attractive for liquidity demanders relative to the traditional market, which raises $\underline{\alpha}$.

As a first step, we demonstrate the dynamic behavior of AMM benefits, based on the annual level for the different fees. Figure 4 plots the fraction of days for which an AMM would yield liquidity suppliers a positive payoff π . Note that a necessary condition for a positive π is that $\overline{\alpha} - \underline{\alpha} > 0$ for the preceding day. We observe two opposing effects. For low market capitalization stocks, the fraction of days with AMM viability increases in fees whereas for high market cap stocks, it is the reverse. This illustrates the opposing forces: low market cap stocks have low volume, and holding volume constant, these markets must provide liquidity providers with higher fees to entice them to contribute to an AMM. These stocks often trade at high spreads and therefore the liquidity demanders "easily" receive better execution in an AMM. High market cap stocks, usually trade at low spreads, and as the fee rises, it becomes simply impossible for an AMM to beat the bid-ask spread.

Figure 5 plots the difference $\overline{\alpha} - \underline{\alpha}$ based on annual averages for $\overline{\alpha}$ and $\underline{\alpha}$ against average market capitalizations, animated for fees $F = 1, \ldots, 30$. This graphic shows that as the fee rises, on average $\overline{\alpha} - \underline{\alpha}$ declines, except for the smallest stocks. Figure 6 plots an animation of the average daily dollar benefits (truncated at 100K for better visibility) from AMMs against market cap as the fee rises from 1 bps to 30bps. The figure illustrates the trade-off from the fees: as the fee rises, the benefit of AMMs over the limit order book dwindles. This effect is due to two reasons. First, there are fewer occasions and stocks for which AMMs are viable at all. And, second, when fees are high, the benefits relative to a traditional limit order market are low.

B. Optimal Fee Visuals

We begin by exploring the data visually based on annual averages for $\underline{\alpha}$, $\overline{\alpha}$, benefits B, and the benefit maximizing fee F^{π} . Furthermore, we also compute the average number of days for which $\pi > 0$ for various fees and the aggregated annual average π · dollar volume. To ensure that our observations are not driven by extreme outlyers, we winsorize the benefits π at the 0.1% level. As outlined above, we display results for q =average size + 2 × standard deviation (average size), with the benchmark volume equal to half the daily dollar volume, and costs in traditional markets as the quoted half-spread.

In the following figures there is group of mid-size companies for which AMMs do not work. The biggest difference between these firms and those for which AMMs work is that the former on average have only about 20% of the latter's daily volume. For low volume stocks, it is challenging for liquidity providers to earn a sufficient income.

Figure 7 plots the benefit-maximizing fees (in basis points) F^{π} against a firm's market cap. Since larger firms usually trade at lower market caps, as expected, the maximum possible fee declines in the market cap. There is much dispersion for low market cap stocks: these stocks often have volatile and heterogeneous volumes and

bid-ask spreads.

Figure 8 plots $\overline{\alpha}$, the maximal amount of deposited liquidity that liquidity providers can tolerate before they lose money. This required deposit rate is declining in the market cap.

Figure 9 plots the fraction of trading days for which AMMs can be operated profitably for liquidity demanders. The figure shows that for the vast majority of stock-year averages, AMMs provide lower fees on more than 80% of trading days, except for the above-identified, low-volume stocks.

Figure 10 plots the average liquidity demander benefit $\pi(F^{\pi})$ in basis points. The figure shows that liquidity demanders save the largest amount proportional cost for small stocks. Notably, spreads for stocks are high so there is more to save.

Figure 11 plots the fraction of total transactions costs that liquidity demanders have saved, computed as $\pi(F^{\pi})/\sigma$. As the figure shows, the average avings are on the order of 30%.

Figure 12 plots the average daily dollar amount that liquidity demanders have saved, computed as $\pi(F^{\pi}) \times \text{dollar}$ volume. The largest dollar savings pertain to large stocks because of the higher volume.

Finally, Figure 13 plots the average return for liquidity providers in basis points. If we did our design correctly, this amount should be on average zero because liquidity providers choose to provide liquidity that makes them break even. As the figure shows, liquidity provider benefits indeed appear to be close to zero.

C. Tabulated Results

Our next step is to examine the summary statistics for our sample. We compute summary statistics for three splits of the data. in the first we average and aggregate across the entire sample. In the second, we average and aggregate by year, where we note that the year 2016 has 9 months of data (Jan-Sept, until the beginning of the tick pilot), 2018 has 3 months (Oct-Dec, after the conclusion of the tick pilot), and 2022 has 3 months (Jan-March, restrictions due to data availability). In the third, we average and sum across market capitalization deciles, where we determine these deciles based on annual averages.

Table I presents general summary statistics for trading variables. We use simple averages and do not filter stocks based on price. Aggregated across all stock, bidask half-spreads are around 40-50 bps, which strongly suggests that an AMM that charges 1-10 bps and that has sufficient liquidity may be economically beneficial. Across market-cap deciles, the average quoted spread falls from over 100 to single digits. In turn, this indicates that for large stocks, only very liquid AMMs with low fees will be viable. We note that trading volume in 2020 and 2021 was significantly higher than in the preceding years, probably as a by-product of the COVID pandemic lockdowns and shift to work-from-home for people who had money to invest. Moreover, trading volume concentrates in the 20% largest firms.

Examining $\underline{\alpha}$ and $\overline{\alpha}$, based on various fees, displayed in Table II, we observe that for a 1 bps fee, $\overline{\alpha}$ is around $10 \pm 2\%$ across the years and market capitalization thresholds. The average value for $\overline{\alpha}$ is usually larger for larger firms. This indicates that when volume is low, the collected AMM fees cannot be distributed among too many providers. For $\underline{\alpha}$ the situation is the reverse, where the effect here is driven by the smaller spread in larger firms. Here, liquidity demanders are only better off if AMMs are sufficiently liquid. In combination, as already indicated by the figures, AMMs are viable for various levels of fees for small firms but only for small fees (if at all) for large firms.

The most important results are in Table III which computes behavior and outcome variables when the used fee is F^{π} (applied and computed per stock and day). The first column lists the optimal fee itself. As indicated by the plots, smaller firms have larger optimal fees compared to large firms. This observation is also visible in the numbers: the average optimal fee for the smallest decile is about 31bps whereas for the large firms, it's less than 1bps. The second column shows for what fraction of days AMM are better than the traditional market whereas the third column shows the number of days for which AMMs are used (i.e., $\overline{\alpha} > \underline{\alpha}$). Notably, AMMs are not better whenever they are used, e.g., because spreads on the traditional market may be lower. Overall, the data show that AMMs are used and better than the traditional market between, roughly, 80-90% of days. The third column displays the average savings per stock and day and the third column lists the savings across the entire sample. The data show that although the proportional savings in terms of relief from the bid-ask spread are largest for small stocks, the dollar-amounts are larger for large stocks, simply because they have higher so much more volume. The last column shows that across the board, AMMs would save liquidity demanders over 30% of transactions costs.

V. Capital Requirements

Our empirical results indicate that liquidity providers would need to deposit on average between 3% and 8% of a firm's shares outstanding to an AMM. For most (though not all) investors, shares are unused capital, and therefore such a deposit under the right conditions is unlikely to be burdensome. However, liquidity provision requires an offsetting amount of cash, a potentially very large sum that is not free because unlike a securities deposit, cash earns interest or needs to be borrowed. During much of our sample period, interest rates were close to zero but for a 6% annual rate, the daily required payments to liquidity providers would increase by around 2bps. These costs need to be recovered via AMM fees, and that may make AMMs not viable for high marketcap stocks where optimal fees are on average below 2bps.

However, in this section we argue that in practice, the actual cash requirement will be much lower than the equivalent of the market cap deposit.

The constant product function that we present implicitly allows liquidity demanders to extract all c cash and a assets. In these extreme cases, the price would go to zero and infinity respectively, and unlikely scenario. Moreover, many second generation AMMs such as UniSwap V3 have a refined pricing function where liquidity providers choose a price band to supply liquidity, and not the full range. In practice AMMs would likely have such a functionality, if only because of single-stock circuit breakers. Therefore, the full amount of cash and shares would never be used in practice.

To better quantify the amount of actual cash needed, we proceed as follows. Assume that liquidity providers agree to supply liquidity as long as the price for a stock is in a band $p_0 \times R \in (\underline{R}, \overline{R})$, where p_0 is the price at which users made their deposit. A price rise coincides with traders removing assets and adding cash, a price decrease coincides with a removal of cash. Since we are concerned with cash liquidity, only price decreases are relevant. For the price to move by R, by equation (6) there is quantity q(R) that has been traded with the pool such that

$$q(R) = a(1 - \sqrt{R^{-1}}),$$

where $a = \alpha S$, the initial deposit of assets, measured as a fraction of shares outstanding. Since R < 1, note that this quantity is negative and thus signifies a sale. We seek to compute how much needs to be available to have liquidity available for prices in $(p_0 \cdot \underline{R}, p_0 \cdot \overline{R})$, we divide both sides by $a = \alpha S$. This is the fraction of the shares outstanding that would need to be made available to accommodate a price drop of R

maxshares(
$$\underline{R}$$
) $\equiv 1 - \sqrt{\underline{R}^{-1}}$. (18)

To obtain the amount of cash to accommodate a drop of R, we substitute back into the price equation and simply

$$\Delta c(R) = c \cdot \frac{a(1 - \sqrt{R^{-1}})}{a - a(1 - \sqrt{R^{-1}})},$$

which simplifies to the fraction of the cash deposit for $R = \underline{R}$

$$\operatorname{maxcash}(\underline{R}) \equiv \left| \frac{1 - \sqrt{\underline{R}^{-1}}}{\sqrt{\underline{R}^{-1}}} \right|, \tag{19}$$

where we use the absolute value for greater clarity (for R < 1, the value will otherwise

be negative).

What would be a reasonable value for <u>R</u>? One benchmark emerges from circuit breakers: U.S. markets have single-stock circuit breakers that are triggered, loosely, if a stock price drops by 10% during the continuous session. After a circuit breaker has been triggered, trading would be halted and would restart with a price-determining auction. Presumably, this would mean that AMMs would be halted, too. For R = 0.9, maxcash(0.9) $\approx 5.1\%$, i.e., about 5.1% of the required AMM deposit in stocks is necessary in cash.

Another option is to derive this number empirically based on past return profiles. One example is to compute averages and standard deviations or rolling intra-day returns. A reasonable approximation is to set \underline{R} as the average minus two or three standard deviations to ensure prices are supported for 95% or 99% of the price band respectively. Another approach would be to consider only the distribution of price decreases or lower partial moments.

Figure 15 displays the year-stock averages of $maxcash(\underline{R})$ based on a 20-day (4week) rolling average of the intra-day return R minus three standard deviations. The average is 2.4%, the maximum 46%. Most of the larger values stem from the year 2020, when markets swung wildly at the start of the COVID pandemic. The implication is that the funding costs for an AMM deposit on a day-to-day basis will be on the order of a fraction of a basis point because the required (and accessible) cash deposits are only a fraction of the deposited capital.

In summary, although cash deposits are not free, the required amounts are but a fraction of the "on-paper" liquidity.

VI. Discussion and Conclusion

The principle behind the superiority of AMMs is not unlike Budish, Cramton, and Shim (2015)'s proposal for frequent batch auctions. Namely, following Foucault (1999), Budish, Cramton, and Shim (2015) show that limit order books expose individuals to adverse selection risk, which inevitably creates costs. They propose periodic auctions as a solution because auctions pool multiple liquidity providers and traders, which eliminates the "sniping" risk and which may lead to better risk sharing. AMMs have this feature, too. Yet frequent auctions have conceptual and organizational shortcomings. A short list is that they require ongoing liquidity provision so as to avoid price dislocations, it is unclear who would supply liquidity, they need to be integrated with the remainder of financial markets which operate in continuous time, their implementation and usage will be technologically challenging, it's not clear how competing auction systems would interact (do we need to create a monopoly?), and so on. AMMs do not have these problems: liquidity provision is passive and in expectation worth the liquidity providers' while; trading is continuous; prices are predictable; different AMMs can run parallel because their liquidity aggregates; and they are easy to use.

Despite the technological challenges, the Securities and Exchange Commission (SEC) has proposed to mandate an auction system for retail orders in its December 2022 roadmap. We believe it is important to have a conversation about this proposal and consider alternative market structures. Glosten (1994) predicted that the electronic limit order book would become the standard in open, public markets. This prediction has proven to be true, as OTC and dealer markets for equities have largely disappeared in favor of the limit order book. Automated Market Makers (AMMs) have only recently

emerged as a concept, albeit for crypto-assets only. Our work provides an empirical case for the benefits of AMMs and argues that they are often better. Notably, they work for *both* passive liquidity providers *and* active liquidity demanders.

However, our approach to introducing AMMs may be overly optimistic; for instance, we assume that they would process all non-intermediated volume. Additionally, our study uses a rudimentary AMM, and there are ways in which they could be improved. For instance, users could set ranges for which they provide liquidity, which would reduce liquidity providers' exposure (UniSwap V3 has such a provision) and lowers capital requirements.

Another criticism of our analysis is that we do not consider the role of people who want to trade with limit orders. Namely, in practice not all trades have intermediaries on the passive side whereas our model implicitly requires that all traders use market orders only. However, a UniSwap V3-type setting can somewhat accommodate the equivalent of trading with limit order trading. In the V3 setting, users can specify a price range for which they supply liquidity. Someone who wants to sell can therefore specify a price range with the lower price at or above the opening price. If the price increases during the day, they would sell the asset. Of course, there are limit to this approach because of pro-rating. In fairness, few models and analyses consider the welfare of limit order traders.

There are also many regulatory obstacles and challenges that would need to be addressed.

First, AMMs require not only assets but also cash, and therefore there would have to be a significant short-term cash-lending market, possibly with off-setting collateral arrangements. Setting up such a market is not a simple task. There are technological ways to improve capital usage, e.g., the Balancer protocol allows liquidity pools that consist of multiple assets and only one of them needs to be cash. This allows traders to change positions in assets without having to convert them to cash, for example, trading Facebook stocks against Google/Alphabet stocks directly.

Second, taxation may complicate the operation and use of AMMs for liquidity providers. If broker-dealers were to operate an AMM, depending on the jurisdiction, each trade may create a taxable event for every AMM liquidity provider. Tracking these events will be challenging and it will affect the profitability of liquidity provision. Furthermore, the number of events may lead tax authorities to classify AMM liquidity provision as a professional activity, in which case, any capital gains may be classified as ordinary income.

Third, there are multiple models of AMMs and AMM pricing, and some of them may conflict with securities regulation. For example, when assets are in a pool, it raises questions about who the beneficiary owner/holder of records is and how dividends, splits, reverse splits, and votes would be handled. Additionally, depositing assets into an AMM may conflict with investment managers' directives and obligations. This means that regulations may need to be tweaked in order to accommodate the presence of such pools. Deposits and AMMs themselves could also be viewed as a securities offering, which would significantly complicate their deployment. Would each pool require a separate prospectus? Would this allow cross-pool trading? Regulation here could well be a poster-case for the creation of barriers that generate economic harm.

Fourth, it is important to consider who would run AMM pools. In the blockchain

world, AMMs can be established "on-chain" because users can self-custody and control their assets. However, exchange operators in their current form cannot take on this role because they do not have custody of the assets. They also stand to lose most of their data income. Broker-dealers also face a conflict, as an AMM pool would negatively impact their securities lending business and would incur additional costs for recordkeeping. Unless volume increases substantially, it is unclear whether broker-dealers would benefit from the introduction of this new tool. As we show, AMM liquidity is additive, so it is unnecessary to establish a single entity to run an AMM pool. Having multiple pools is fine, provided they are accessible and linkable. However, if brokerdealers run AMMs, there would need to be a system to access the different pools, and that requires the linking of disjoint systems. On the other hand, public blockchains do not have this problem because all AMMs run on the same infrastructure. AMMs that co-exist with limit order books would face the problem that their functionality may be incompatible with Reg NMS. Again, regulation could be a barrier to significant cost savings.

However, the world is changing: public blockchains are enabling the trading of tokenized assets, traditional financial institutions are developing tokenizations and may soon run trading platforms like AMMs, and regulatory changes such as the European Union's "Markets in Crypto Assets" (MiCA) or the UK's crypto regulation roadmap pave the way for asset tokenizations with regulatory certainty. Another fair question to ask is: if AMMs work so well, why don't they process 100% of trading volume in crypto-assets? Our answer is that blockchain-based AMMs have a steep learning curve and are still not user friendly: Although the interfaces for AMMs like UniSwap are straightforward to use, there are significant complications. For instance, traders need to learn and be comfortable with using so called self-custody wallets. They also (almost) always need two tokens to trade: one is the native cryptocurrency that they need to pay blockchain validators and the other is the asset of choice that they want to trade. Note though that an AMMs as we treat them here do not require a blockchain — an exchange or broker-dealer could organize trading in an AMM, and they could provide a seamless user experience.

Finally, to be transparent: our analysis is optimistic as we assume that a very significant fraction of all trading volume would move to AMMs if they were available. Arguably, however, our analysis is also conservative because it does not take into account potential increases in trading volume that may occur as a result of decreased trading costs.

REFERENCES

- Aoyagi, Jun, 2022, Liquidity provision by automated market makers with asymmetric information, Working paper Hong Kong University of Science and Technology.
- ———, and Yuki Ito, 2021, Liquidity implications of constant product market makers, Working paper UC Berkeley https://ssrn.com/abstract=3808755.
- Barbon, Andrea, and Angelo Ranaldo, 2022, On the quality of cryptocurrency markets: Centralized versus decentralized exchanges, Swiss Finance Institute Research Paper No. 22-38 University of St.Gallen School of Finance https://ssrn.com/abstract=3984897.

- Boehmer, Ekkehart, Charles M Jones, Xiaoyan Zhang, and Xinran Zhang, 2021, Tracking retail investor activity, *The Journal of Finance* 76, 2249–2305.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *The Quarterly Journal* of Economics 130, 1547–1621.
- Capponi, Agostino, and Ruizhe Jia, 2021, The adoption of blockchainbased decentralized exchanges, working paper Columbia University https://ssrn.com/abstract=3805095.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market1, Journal of Financial Markets 2, 99–134.
- Glosten, Lawrence R., 1994, Is the electronic open limit order book inevitable?, The Journal of Finance 49, 1127–1161.
- Hasbrouck, Joel, Thomas Rivera, and Fahad Saleh, 2022, The need for fees at a dex: How increases in fees can increase dex trading volume, Working paper New York University https://ssrn.com/abstract=4192925.
- Lehar, Alfred, and Christine Parlour, 2021, Decentralized exchanges, Working paper University of Calgary link.
- Park, Andreas, 2021, The conceptual flaws of automated market making, Working paper University of Toronto https://papers.ssrn.com/abstract=3805750.

A The Origin of AMMs

Although public blockchains such as Ethereum are general purpose value management system, a blockchain network does not make a market. Instead, until mid-2020, users could trade blockchain-based items or tokens only on centralized, "off-chain" exchanges. This changed by with the release of Automated Market Maker (AMM) systems such as UniSwap or SushiSwap, which have seen tremendous user uptake, processing billions of dollars worth of transactions every day, often more than the largest centralized exchanges such as Binance and Coinbase.

The nexus of the development of AMMs were discussions around the pricing rule to determine exchanges for tokens using blockchain-based systems. The first mention of automated decentralized market making is in a 2016 Reddit post by Vitalik Buterin. Martin Koppelmann proposed the first constant product pricing scheme.

Why was there no blockchain-based trading earlier? In principle, it is possible to organize crypto-asset trading directly on a blockchain by registering limit orders as "smart contracts." However, this is not practical because each new limit order costs a validation fee and because unmatched orders would waste resources as they need to be processed by all 10,000+ network nodes.

There are numerous refinements of the constant product rule, and there are pools that administer not just pairs but arbitrary numbers of assets (e.g. Balancer). Most of the refinements lead to even lower trading costs while offering liquidity providers protection against large price movements. The goal of this paper is to establish a simple baseline against which one can assess the merits of AMMs for traditional markets in terms of improving welfare. Most likely, innovation would further increase welfare if



Figure 1 Constant Product Prices in Automated Market Makers

9

The figure illustrates how the exchange rate of stocks to cash is determine in an automated market maker. The initial deposit in the contract is for a assets and c cash, so that the pool's liquidity is L(a, c). A trader wants to buy (i.e., remove) quantity q of the assets. The price $\Delta c(q|a, c)$ is such that after the purchase, the amount of cash in contract is such that the point $(a - q, c + \Delta c(q|a, c))$ is on the $L(\cdot, \cdot)$ curve.

the concept itself is implemented.



The Per-Unit Price Function for Constant Product AMMs

The figure shows the per-unit price function for constant product market making for an asset that has a marginal price of \$10 for various levels of deposited liquidity. The marginal price is $p_0 = c/a$; the per unit price is p(q) = 10a/(a-q). We plot this function for a = 1, ..., 5.



Figure 3 The Incremental Loss from Long-Run Adverse Selection (ILLRAS) for Constant Product Pricing

File size restrictions prevent us from including this file directly in the document. Please click here for a GIF of the figure and here for a pdf with the figure included.

Figure 4 Animation for the Fraction of Days that AMMs dominate by Fee

File size restrictions prevent us from including this file directly in the document. Please click here for a GIF of the figure and here for a pdf with the figure included.

Figure 5 Animation for the Difference between the Minimum and Maximum Viable Liquidity

File size restrictions prevent us from including this file directly in the document. Please click here for a GIF of the figure and here for a pdf with the figure included.

Figure 6 Animation for the Daily Benefits by Fee



Figure 7 Benefit-Maximizing AMM Fees



Figure 8 Fraction of MarketCap Deposited



Figure 9 Percent Days that AMM is better



Figure 10 Cost Advantage AMM vs. Traditional (measured as difference in relative transaction costs)



Figure 11 Fraction of Transaction Costs that AMM saves per day and stock



Figure 12 Dollar Amount of Transaction Costs that AMM saves per day and stock

50



Figure 13 Liquidity Provider Daily Cost/Benefit





Panel A: % retail vs. Market Cap

Panel B: Quoted Spread vs. Market Cap



Panel C: Cancel-to-Trade vs. Market Cap Panel D: Cancel-to-Trade vs. Quoted Spread

Figure 14 Relationships among the Explanatory Variables



Figure 15 Required Cash Requirements as a Fraction of Asset Deposits

Table ISummary Statistics Traditional Trading Variables

This table presents general summary statistics for trading variables. We use simple averages and do not filter stocks based on price. We present averaged and aggregated across the entire sample, split by years (where 2016 has 9 months of data (Jan-Sept, until the beginning of the tick pilot), 2018 has 3 months (Oct-Dec, after the conclusion of the tick pilot), and 2022 has 3 months (Jan-March, restrictions by data availability)), and split by market capitalization deciles, where we determine these deciles based on annual averages. The first column contains the average, time-weighted bid ask spread, the second the fraction of retail trading as provide by WRDS and computed according to Boehmer, Jones, Zhang, and Zhang (2021), the third column is the fraction of institutional trading based on trade sizes above 20K, the fourth column is total dollar volume in trillions of USD, and the last column is a rough approximation of trading costs computed as the produce of the bid-ask spread and dollar volume.

	quoted spread in bps	%retail	%institutional	dollar volume (in trillions)	transaction costs (in billions)
full sample	49.3	9.8	13.8	391.0	181.0
2014	44.5	9.5	13.6	43.2	16.4
2015	51.4	9.6	12.8	45.5	17.3
2016 (9m)	54.4	10.7	12.1	32.9	11.9
2018 (3m)	64.6	10.2	12.7	15.2	8.5
2019	52.0	10.3	13.2	53.0	21.1
2020	50.8	10.2	14.3	79.3	43.7
2021	35.8	9.0	17.2	94.8	48.8
2022 (3m)	41.1	9.2	14.7	27.1	13.5
lowest mcap quintile	153.6	24.9	3.4	0.8	6.6
2	108.6	18.6	7.2	0.9	4.6
3	74.6	12.8	10.3	0.9	4.0
4	50.4	9.1	12.6	1.8	5.3
5	36.1	7.4	12.8	3.5	7.4
6	24.1	5.9	13.1	6.4	10.8
7	16.2	5.1	13.8	13.3	15.2
8	11.2	4.4	15.8	24.4	20.9
9	7.3	4.4	20.2	55.6	33.1
highest mcap quintile	4.0	5.1	30.3	283.5	73.2

Table IISummary Statistics for Viability Thresholds $\underline{\alpha}$ and $\bar{\alpha}$

This table presents general averages for the fraction of market cap provided as liquidity that makes liquidity demanders indifferent between using an AMM and a traditional market, $\underline{\alpha}$, and the threshold for liquidity suppliers to break even, $\bar{\alpha}$. The rows are structured as in Table I, the columns are for various fee levels, where F^{π} is the fee that maximizes liquidity demander benefits.

			\bar{lpha}					<u>α</u>					$\bar{\alpha} - \underline{\alpha}$		
fee	1bps	5bps	10bps	30bps	F^{π}	1bps	5bps	10bps	30bps	F^{π}	1bps	5bps	10bps	30bps	F^{π}
full sample	8%	17%	23%	35%	6%	3%	19%	34%	62%	3%	5%	-2%	-12%	-28%	3%
2014	9%	19%	26%	40%	6%	2%	23%	39%	66%	2%	7%	-4%	-13%	-27%	4%
2015	9%	19%	25%	38%	6%	2%	22%	38%	66%	2%	7%	-3%	-13%	-28%	4%
2016	9%	18%	24%	37%	5%	2%	21%	38%	64%	2%	7%	-3%	-13%	-27%	3%
2018	7%	14%	20%	31%	5%	2%	17%	32%	59%	2%	5%	-2%	-12%	-28%	3%
2019	9%	18%	24%	37%	6%	3%	21%	38%	62%	3%	6%	-3%	-14%	-25%	3%
2020	7%	15%	20%	32%	5%	2%	14%	29%	58%	2%	4%	0%	-9%	-26%	3%
2021	8%	18%	24%	36%	7%	4%	16%	32%	63%	4%	4%	1%	-8%	-27%	3%
2022	7%	15%	20%	30%	6%	6%	17%	31%	62%	5%	1%	-2%	-12%	-32%	1%
Quintile															
1	4%	9%	12%	20%	8%	2%	2%	2%	4%	3%	2%	7%	10%	16%	4%
2	4%	9%	12%	19%	7%	5%	5%	6%	12%	6%	-1%	3%	6%	7%	1%
3	5%	11%	15%	23%	7%	6%	7%	9%	29%	7%	-1%	4%	5%	-6%	1%
4	6%	13%	18%	28%	8%	7%	9%	13%	50%	7%	-1%	5%	5%	-22%	1%
5	7%	15%	21%	32%	6%	3%	4%	12%	66%	3%	4%	11%	9%	-34%	3%
6	8%	17%	23%	36%	5%	1%	3%	22%	79%	1%	7%	14%	1%	-43%	4%
7	9%	20%	26%	41%	5%	1%	7%	42%	90%	1%	9%	12%	-16%	-50%	4%
8	11%	23%	30%	46%	5%	1%	22%	63%	96%	1%	10%	1%	-33%	-50%	4%
9	13%	26%	34%	51%	4%	1%	48%	80%	99%	1%	12%	-22%	-46%	-48%	4%
10	13%	27%	36%	54%	3%	3%	81%	96%	99%	0%	10%	-54%	-59%	-45%	2%

Table III Summary Statistics of Based on Optimal Fees

This table presents summary statistics for benefits and outcomes when the fee of the AMM is F^{π} for each stock. These fees (as averages) are in the first column. The rows are structured as in Table I.

	average optimal fee (in bps)	Days AMMs are better	Days AMM are used	average daily savings per stock & day	aggregate savings (in B\$) all stocks	aggregate transactions costs (in B\$) all stocks	% saved
all	10.6	82%	86%	\$8.303	\$54.3	\$167.4	32%
2014	9.7	79%	83%	\$4,472	\$4.8	\$14.6	33%
2015	11.4	79%	83%	\$4.712	\$5.2	\$15.6	33%
2016	12.6	79%	84%	\$4.883	\$3.9	\$11.3	34%
2018	12.4	81%	86%	\$8,802	\$2.3	\$7.3	31%
2019	11.0	83%	87%	\$6,081	\$6.5	\$19.1	34%
2020	10.9	85%	89%	\$11,705	\$12.5	\$39.8	31%
2021	8.1	85%	89%	\$11,425	\$15.0	\$46.4	32%
2022	9.4	83%	87%	\$12,498	\$4.2	\$13.3	32%
Quintile							
1	30.8	78%	84%	\$1,906	\$1.2	\$3.9	30%
2	24.1	72%	80%	\$1,513	\$1.0	\$3.2	33%
3	17.6	72%	80%	\$1,420	\$0.9	\$3.0	32%
4	12.7	74%	81%	\$1,945	\$1.3	\$4.2	31%
5	8.0	82%	87%	\$3,015	\$2.0	\$6.3	32%
6	4.9	87%	90%	\$4,773	\$3.1	\$9.4	33%
7	3.3	88%	90%	\$7,066	\$4.6	\$14.0	33%
8	2.2	88%	91%	\$9,983	\$6.5	\$19.5	33%
9	1.5	88%	91%	\$15,787	\$10.2	\$31.5	32%
10	0.8	89%	91%	\$35,654	\$23.4	\$72.4	32%

Table IV Summary Statistics of for Cash Requirements

This table presents the estimated amount of shares and cash that liquidity providers would need to produce to be liquidity providers. The number is computed assuming that the AMM shares the optimal fee F^{π} . The rows are structured as in Table I.

	fraction of mcap in liquidity pool	% cash required	average per day cash per stock (in millions of USD)
full sample	6%	2%	\$172.1
2014	6%	3%	\$155.3
2015	6%	2%	\$105.0
2016	5%	2%	\$108.9
2018	5%	3%	\$194.6
2019	6%	2%	\$146.6
2020	5%	3%	\$220.4
2021	7%	2%	\$213.8
2022	6%	2%	\$220.1
Quintile			
1	8%	4%	\$1.0
2	7%	3%	\$2.2
3	7%	3%	\$4.2
4	8%	3%	\$6.9
5	6%	3%	\$12.3
6	5%	3%	\$21.4
7	5%	2%	\$35.4
8	5%	2%	\$58.5
9	4%	2%	\$118.5
10	3%	1%	\$752.8

57