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

Katya Malinova[†] Andreas Park[‡]

August 4, 2023

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. We develop a theoretical approach for AMM liquidity provision in equities and then calibrate the model parameters, based on U.S. equity trading data. Our analysis suggests that an optimally designed AMMs could save U.S. investors billions in transaction costs each year. 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 savings arise across the board and would particularly improve the liquidity and trading costs for small firms, allowing them to attract more investors and capital.

[*]Financial support from the *Social Sciences and Humanities Research Council*, *Mackenzie Investments Chair*, *Global Risk Institute*, and *The Canadian Securities Institute Research Foundation* is gratefully acknowledged. We thank seminar and conference participants at the 2023 Women in Microstructure Workshop, the Swiss National Bank Conference on Digital Assets, the Oxford-Man Institute, UC Santa Barbara DeFi Workshop, VGSF, and NHH Bergen as well as Alfred Lehar, Angelo Ranaldo, Julia Reynolds, and Alexander Wehrli.

[†]McMaster University `malinovk@mcmaster.ca`.

[‡]University of Toronto `andreas.park@rotman.utoronto.ca`.

Automated market makers (AMMs) are a new type of trading institution that has no parallel in traditional financial markets and already processes 15-20% 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?

We first characterize the liquidity provision decision for an AMM, building on Lehar and Parlour (2021). One of our contributions is to describe liquidity as a function of a firm's equity. We then derive the AMM parameters that maximize trader welfare and calibrate the model using U.S. equity trading data. Our analysis suggests that an optimally designed AMMs would save U.S. investors up to 30% of annual transactions costs and lead to a substantial welfare improvement.

AMMs are a genuinely novel market institution 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 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.

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. Arguably, AMMs are based on a mindset different from that typically associated with liquidity provision. An AMM represents a liquidity pool in which owners deposit their assets. It provides them with the opportunity to use their existing capital to earn income from trading fees while retaining exposure to their chosen asset. Risk and fee income are pro-rated among depositors. Consequently, liquidity providers don't have to participate in an arms race for superior speed and market monitoring technology. There's also no pressure for liquidity providers to maintain a zero overnight position in an asset; in fact, they desire exposure to the asset. This approach is especially beneficial for long-term investors, including those who may lack sophisticated trading knowledge. Through better risk sharing, AMMs would increase economic efficiency.

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 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. 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. Liquidity demander prefer to pay lower fees: let $\underline{\alpha}$ denote the share for which AMMs are cheaper than traditional markets for liquidity demanders. If $\overline{\alpha} > \underline{\alpha}$, theoretically there exist an economically viable, welfare-improving AMM configuration.

Liquidity suppliers are intuitively willing to commit more capital when the trading fee is higher (all else equal). The higher the liquidity, the lower is the price impact of each trade, and, therefore, liquidity demanders benefit (somewhat) from higher fees. Assuming that liquidity providers compete and submit liquidity so that they break even, $\alpha = \overline{\alpha}$, as we show, there is non-zero fee that maximizes liquidity demander welfare (a similar result is in Hasbrouck, Rivera, and Saleh (2022)).

At the end of the day, however, it is an empirical question whether AMMs can be economically viable and whether the alleged cost savings would materialize. 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 (likely favor AMMs). We then calibrate the key model parameters —optimal fees and amount of liquidity provided— to observations, of volumes, trade sizes, return characteristics, and transactions costs. For small firms (which have low volume), the optimal fee is about 31 basis points, for large firms (which have higher volume) it is only 0.8 basis points. For these fees, liquidity providers would break even if they provide no more than 8% of shares outstanding for small firms and 3% for the largest firms as liquidity.

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. Average benefits across all firms are around 30% of annual trading costs, around \$6.5 billion for 2019, \$12.5 billion in 2020, and \$15.0 billion in 2021. The proportional benefit is largest for small firms, the dollar amounts (due to higher volumes) are highest for large firms.

Although it may be (financially) straightforward to provide the shares for liquidity, AMMs require an offsetting amount of cash as liquidity and this amount is not small. We show, however, that in practice, only 2-5% of the theoretical amount of cash (i.e., 3-32 basis points) need to be kept at the ready. Overall, our analysis indicates that optimally 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. There is, therefore, an entirely different interpretation of our findings. Namely, our analysis gives an indication of hitherto unmeasured excess trading costs, based on the current market behavior of market participants and stock characteristics. 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.

Liquidity providers in Aoyagi and Ito (2021) derive a positional benefit from trading with uninformed traders by withdrawing and reposting liquidity after each trade. Differently, in Lehar and Parlour (2021) liquidity providers only withdraw liquidity infrequently. In their model, arbitrageurs approximately revert uninformed transactions, which eliminates the positional gain for liquidity suppliers but generates fee income. We follow Lehar and Parlour (2021) and assume the presence of a balanced order flow, which is similar to an arbitrary uninformed flow that gets reversed by arbitrageurs.

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 (2023) 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 builds on Lehar and Parlour (2021) and Park (2023). The description of some features of liquidity provision follows 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. If a user adds/withdraws both cash and assets, $\Delta a, \Delta c$ with $\text{sign}(\Delta a) = \text{sign}(\Delta c)$ and

$$L(a, c) \lesseqgtr L(a + \Delta a, c + \Delta c)) \quad \text{for } \text{sign}(\Delta a) = \text{sign}(\Delta c) \gtrless 0.$$

If a user wants to add/withdraw Δa units of liquidity they would need to add/withdraw Δc such that $\Delta c / \Delta a = c/a$.

The second operation are purchases of the asset from the pool (or sales of the asset to the pool) where $\text{sign}(\Delta a) = -\text{sign}(\Delta c)$. 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.

The important implication of the above setup is that in an AMM, prices move only through trades whereas changes in liquidity are price neutral.

Traditional stock markets are almost always organized as limit order books in which limit orders are ordered by price. The cost for a trade of size q is then the sum of the costs for the different orders that the order trades against. For instance, a buy order for 1,000 shares may trade against 100 shares for \$10.01, 500 for \$ 10.02, and 400 for \$10.05 as the order "walks the book," and its cost is $100 \times \$10.01 + 500 \times \$10.02 + 400 \times \$10.05$. Mathematically, the cost of an order can then be approximated by an integral over a marginal pricing function $\gamma(s)$ such that $\Delta c(q) = \int_0^q \gamma(s)\ ds$. Consequently, the constant product rule is equivalent to a limit order market with marginal price $\gamma(s) = ca/(a - s)^2$.

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 any 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.

Many microstructure models take a short-term perspective: liquidity providers incur a positional loss against informed traders and recover these losses by a positional gain from trading with uninformed traders. Such an approach implicitly requires that a market maker can liquidate a position at an inefficient price.

In contrast, the premise of an AMM is that existing asset holders make longer-term deposits and provide liquidity for a meaningful amount of time. We therefore assume that deposits and withdrawals occur at efficient prices, so that liquidity providers share trading-related risks. Lehar and Parlour (2021)'s approach is related to ours: they assume that any uninformed trade is (approximately) reverted by arbitrageurs. Therefore, liquidity providers have no positional gain from trading against uninformed traders so that an AMM needs to include an explicit compensation scheme. We assume here, as is common in practice, that compensation comes in the form of a fee F that is levied on the dollar volume of each trade.

Very loosely, let the buyer-initiated volume be B and the seller initiated volume S and, 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 assume that prices are efficient, i.e., that after B buys and S sells, "the price is right." A liquidity provider earns the fee on the total volume $B + S$, but incurs a positional loss because of the price-moving volume, $B - S$.

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

$$\text{fee income} + \Delta c(B - S) - p_t(B - S) \geq 0. \tag{2}$$

Note that the second part of the above, $\Delta c(B-S)-p_t(B-S)$, is 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 $p_t(a - (B - S)) + c + \Delta c(B - S)$. The difference of this amount and the value of buy-and-hold is what is specified in (2).

Crucially, in traditional 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.[1] Therefore, for AMM liquidity pools, the suitable benchmark actually is a buy-and-hold strategy of the deposited

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

assets. In what follows, we present a formal model based on Park (2023).

B. *Model Primitives*

There are two assets, risky A and safe C, the marginal price for A measured in C at 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 value 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 a fee which is an exogenous parameter of the AMM protocol. To simplify the exposition, we assume that fee F accrues based on the dollar volume of transactions and is pro-ratedly paid directly to liquidity providers.[2] Therefore, when a liquidity demander buys q units 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. We assume that in addition to a price-moving trade, during the liquidity provision horizon (say, a day), there is a volume V that is perfectly

[2] Many blockchain-based implementations of AMMs have a 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.

balanced.[3] Since fees are collected based on the dollar volume, the order of trades affects the fee income. For instance, when a buy-volume of $V/2$ assets is followed by a sell-volume of $V/2$, 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 V accrue at the marginal price p_0. Let q^* denote the price-moving trade; then total fees are

$$F\Delta c(q^*) + Fp_0V. \tag{3}$$

The Role of Fees. As Park (2023) shows, the fee is necessary because prices are not regret-free and the constant product pricing function does not provide adequate compensation to liquidity providers. Moreover, 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

[3] This is a benign simplification because we 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. We therefore assume that in expectation, $b = s = V/2$ so that total expected volume is V.

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 $d = 2c = 2ap_0$.

Suppose that the fundamental changed from V_0 to V_t. We assume that arbitrageurs trade against the pool to buy the underpriced asset for as long as it's profitable, i.e., until the marginal price $p_t = V_t$.

Lemma 1: *Let $R = p_t/p_0$. The quantity that moves the price from p_0 to p_t is $q^*(R) = a\left(1 - \sqrt{R^{-1}}\right)$.*

Proof. 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

$$p_t = \frac{c + \Delta c(q^*)}{a - q^*}. \tag{4}$$

Under constant product pricing, 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 (4) yields

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

Solving for the price moving quantity q^*

$$q^*(R) = a - \sqrt{\frac{ac}{p_t}} = a\left(1 - \sqrt{\frac{p_0}{p_t}}\right) = a\left(1 - \sqrt{R^{-1}}\right), \tag{5}$$

where the penultimate relation obtains because $c = ap_0$, and the last relation uses $R = p_t/p_0$ for the gross return, where also $R = V_t/V_0$. □

We can establish the following result:

Lemma 2: *The expected return on deposit d from providing liquidity is*

$$\int_0^\infty \left(\sqrt{R} - \frac{1}{2}(1+R) + \frac{F}{2}|\sqrt{R}-1|\right)\ \phi(R)dR + F\frac{V}{2a}. \tag{6}$$

Proof. Liquidity providers receive a net payoff $\Delta c(q^*(R)) - p_t q^*(R)$. Collecting terms, for balanced volume V, the expected return on deposit d from providing liquidity is

$$\frac{1}{d}\left(\int_0^\infty (\Delta c(q^*) - p^* p_t(R) + F \cdot \Delta c(q^*))\ \phi(R)dR + Fp_0V\right). \tag{7}$$

To compute the net positional payoff from trading q^*, $\Delta c(q^*) - p_t q^*$, we express each component separately as follows.[4] First, we divide the first term, $\Delta c(q^*)$, by $2c (= d)$ and simplify as follows:

$$\frac{\Delta c(q^*)}{2c} = \frac{1}{2c} \cdot \frac{q^* c}{a - q^*} = \frac{1}{2}\frac{q^*}{a - q^*} = \frac{1}{2}(\sqrt{R} - 1), \tag{8}$$

where we substituted for q^* from (5) in the last step. Second, we divide $p_t q^*$ by

[4] Barbon and Ranaldo (2022) use a similar formulation in their empirical work.

$2ap_0(= d)$

$$-\frac{p_t q^*}{2ap_0} = -\frac{R}{2a} \cdot \left(a\left(1 - \sqrt{R^{-1}}\right)\right) = \frac{1}{2}(\sqrt{R} - R). \tag{9}$$

Combining the terms (8) and (9) yields the expression for the relative loss that results as a function of the gross return R

$$\text{ILLRAS}(R) = \sqrt{R} - \frac{1}{2}(1 + R), \tag{10}$$

where we use the acronym ILLRAS as an abbrevaition for the *Incremental Loss from Long-Run Adverse Selection*. Finally, we substitute (5) to compute the fee return

$$F\,\frac{\Delta c(q^*)}{d} + F\,\frac{p_0 V}{d} = \frac{F}{2}|\sqrt{R} - 1| + \frac{F\ V}{2a}, \tag{11}$$

where we use absolute values to account for sales. Combining the losses due to changes in the fundamental (10) and the fee return (11) and taking expectations, we can express the expected return to liquidity provision (7) as follows

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

$\square$

The critical parameter for liquidity provision is the initial deposit. We can now establish the following expression for the break-even liquidity deposit.

Proposition 1: *The break-even deposit is* $a^* = FV/2 \cdot C^{\mathsf{CP}}(\phi, F)$*, where* $C^{\mathsf{CP}}(\phi, F) = -F \times E[|\sqrt{R} - 1|/2] - E[ILLRAS]$.

Proof. With competitive liquidity supply, the liquidity providers would break-even in expectation, (7) $= 0$ Rearranging terms, the break-even deposit a^* can be expressed as follows

$$a^* = \frac{FV}{2}\left(\int_0^\infty \left(-\frac{F}{2}|\sqrt{R}-1| - \sqrt{R} + \frac{1}{2}(1+R)\right)\ \phi(R)dR\right)^{-1}. \tag{12}$$

The term in parentheses can be split into two components

$$\begin{aligned} & \int_0^\infty \left(-\frac{F}{2}|\sqrt{R}-1| - \sqrt{R} + \frac{1}{2}(1+R)\right)\ \phi(R)dR \\ = & \ -\int_0^\infty \frac{F}{2}|\sqrt{R}-1|\ \phi(R)dR - \int_0^\infty (\sqrt{R} - \frac{1}{2}(1+R))\ \phi(R)dR \\ = & \ -F \times E[|\sqrt{R}-1|/2] - E[\text{ILLRAS}] = C^{\mathsf{CP}}(\phi, F). \end{aligned}$$

As the derivation shows, the value a^* is unique. ☐

It is important to emphasize that (6) is the return for liquidity provision, not the return for holding the asset. When (6) $= 0$ holds, the liquidity provider breaks even for liquidity provision but earns a return on holding the asset itself.

A key component that drives a^* is the loss caused by changes in the fundamental value when holding the asset in a liquidity pool, expression (10), 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 assets) for positive and negative returns.

The amount ILLRAS(R) is incremental to any returns from holding the asset (i.e., depositors retain exposure to the asset's returns). To get a sense of the magnitude of this loss consider two scenarios: if the gross return is 0.9 (the asset lost 10% of its value), the incremental loss relative to buy-and-hold 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 compared to buy-and-hold. This implies that the liquidity provider gains 9.88% when the price increases by 10%.

Equation (12) is the maximum amount of liquidity that can be provided so that liquidity providers experience no loss. The goal of this paper is to examine liquidity provision for listed equities. A very useful feature of our formulation is that it allows us to describe liquidity by the fraction of shares outstanding for a listed asset. Namely, we can describe liquidity by the fraction α of the shares outstanding S that 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 $a = \alpha S$, and then let $\overline{\alpha}$ be the break-even fraction of the shares outstanding for liquidity providers.

Corollary 1: *The maximal liquidity deposit (i.e., the deposit for which liquidity providers break even in expectation) measured in the fraction of shares outstanding is given by*

$$\overline{\alpha} = \frac{FV}{2S} \frac{1}{C^{\mathsf{CP}}(\phi, F)}. \tag{13}$$

Note that we implicitly require that $\overline{\alpha} \in [0, 1]$; i.e., if the right hand side expression in (13) is larger than 1, then any liquidity provided is feasible for liquidity providers and we set $\overline{\alpha} = 1$. Likewise, if the right hand side expression in (13) is negative, then no amount of liquidity no matter how small would allow the liquidity providers to break even and we set $\overline{\alpha} = 0$.

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 the 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

$$p\text{-}imp(q) = \frac{\frac{c}{a-q} - \frac{c}{a}}{\frac{c}{a}} = \frac{q}{a-q} = \frac{q}{\alpha S - q}, \tag{14}$$

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 whether liquidity demanders prefer the AMM to the traditional market where they pay the spread σ, i.e.,

$$\frac{q}{\alpha S - q} + F \leq \sigma. \tag{15}$$

The left hand side of the above is decreasing in the liquidity variable α. For the liquidity demander to be willing to use the AMM, the above inequality needs to be satisfied. We can then define the minimal amount of liquidity that ensures this to be the case.

Lemma 3: *The smallest amount of liquidity that would need to be deposited so that a* *liquidity demander who wants to trade q is willing to use the AMM is defined by*

$$\underline{\alpha} = \frac{q}{S}\frac{1+\sigma-F}{\sigma-F}. \tag{16}$$

Proof. The amount $\underline{\alpha}$ is the value of α such that (15) is solved with equality. $\square$

As with $\overline{\alpha}$, we require implicitly that $\underline{\alpha} \in [0,1]$. If the right hand side expression of (16) is negative, then any amount of AMM liquidity would dominate the traditional market, and we set $\underline{\alpha} = 0$. Conversely, if the right hand side expression of (16) exceeds 1, then no amount of AMM liquidity would allow the AMM to be preferred to the traditional market and we set $\underline{\alpha} = 1$. 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 limit order 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. 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

$$\pi(q) = \sigma - \frac{q}{\overline{\alpha}S - q} - F, \tag{17}$$

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. We can show the following.

Proposition 2: *Assume that liquidity provision is competitive so that the deposited liq-*

uidity is $\alpha = \overline{\alpha}$. Then for given q, there exists a fee F^π that minimizes the AMM trading costs for the liquidity demander.

Proof. Liquidity demanders submit liquidity until they break even with $\overline{\alpha} = \frac{FV}{2S}\frac{1}{C^{\mathsf{CP}}(\phi,F)}$. Their liquidity is a function of the fee F. The trading cost for liquidity demanders for trading q is $q/(\overline{\alpha}S - q) + F$. Substituting $\overline{\alpha}$, we compute that local minimum for the cost function

$$F^\pi = \frac{1}{E[|\sqrt{R}-1|/2]+V}\left(-2q\ E[\text{ILLRAS}] + \sqrt{-2qV\ E[\text{ILLRAS}]}\right). \qquad (18)$$

□

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.

F. *Assessing Welfare Improvement: Spread and Depth*

In assessing the liquidity improvement, we compute two measures: first, the transaction cost savings, encapsulated in the value of (17) for the optimal fee F^π. Second, for given F^π, we compute the maximum quantity q^π that a trader can trade such that the cost in the AMM is the same as in the traditional market. This quantity solves $\pi(q) = 0$, and it is

$$q^\pi = \frac{S\alpha^\pi(\sigma - F^\pi)}{1+\sigma-F^\pi}. \qquad (19)$$

The reference point for this quantity is the depth in the traditional market because trades larger than the depth would lead to a cost that exceeds the spread.

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 simple heuristic that we believe has intuitive appeal. If this simple approach delivers a strong indication that AMMs can be used profitably then a more complex optimization would likely deliver even stronger results.

An important feature of an AMM is that liquidity providers keep their deposits in the liquidity pool for a period of time. The longer they keep them, the more fees they earn, but also the larger is the risk of adverse price movements. For risk sharing, it is further important that liquidity providers don't differentially withdraw their funds right at the first sign of an adverse price movement. More broadly, we envision that these liquidity providers are not like professional intermediaries that hold positions purely for trade facilitation but, rather, these are long-term investors who simply seek to put their capital to use.

With this in mind, we focus our analysis on a day-by-day process, that is, we compute within-day returns and assume that liquidity deposits are locked for the duration of the trading day. The specific process that we assume is as follows:

1. The market opens with an auction. Its market clearing price determines the ratio of cash to assets for AMM deposits.[5]

2. The AMM operator computes the welfare optimizing fee F^{π}.

[5]This process ensures that users face no overnight risk from a liquidity pool deposit.

3. Liquidity providers submit liquidity up to $\overline{\alpha}$, provided that $\overline{\alpha} > \underline{\alpha}$, using the previous day's volume and the estimated expected return.

4. Trading commences and liquidity deposits are locked up.

5. At the end of each day, all funds are withdrawn from the pool and added back to the liquidity providers' accounts.

B. Estimating Returns

A critical component in the liquidity deposit decision is the asset's return distribution because we need to determine $C^{\mathsf{CP}}(\phi, F)$. We follow two approaches. In the first, ad hoc formulation, liquidity providers look back one day and base their decision on the previous day's return. In the second, we assume liquidity providers compute intra-day returns over a long stretches of time and base their decision on this distribution.

Concretely, for the ad hoc approach 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^*, and we use the previous day's realized open-to-close return as the "best guess" for the coming day's return.

In our second approach, we stay close to the model and estimate $E[|\sqrt{R} - 1|/2]$ and $E[\text{ILLRAS}]$ parametrically based on realized returns. Namely, we estimate $\phi(R)$ using the Gamma distribution with its scale and slope parameters, where Gamma distributed returns follow density $f(R) = \frac{1}{b\Gamma(c)} \left(\frac{R}{b}\right)^{c-1} \cdot e^{-\frac{R}{b}}$ with scale and shaped parameters b, c, mean $E[R] = bc$, and variance $Var(R) = b^2 c$.[6] We use the observed open-to-close returns R for our sample from January 2014 to end-March 2021, we ignore days with

[6]In untabulated work, we also estimate the distribution non-parametrically; the results are similar.

returns smaller than .5 (the price halved) and larger than 2 (the price doubled), and we require there to be at least 100 return observations. Figure 4 provides two visible examples for realized and estimated returns: one for Tesla Inc, the other for Microsoft.

Using the return distribution, we can compute the two distribution related components of $C^{\mathsf{CP}}(\phi, F)$, $E[|\sqrt{R}-1|/2]$ and $E[\text{ILLRAS}]$. We then apply these values in the computation for the optimal fees for the remainder of the sample, from April 2021 to sample-end March 2022.

We believe that there is merit for both approaches: the first would be easy to implement in practice, the second is closer to the model. In presenting both, in particular the ad hoc, first one, we therefore show that an AMM can be implemented straightforwardly even when sacrificing 1:1 matching of the model setup.

C. Variables of Interest

We are interested in the optimal fee F^{π}, the number of days for which the AMM leads to a welfare gain $\pi > 0$, the size of the average welfare gain π and the aggregate dollar amount saved, and the equivalent quantity q^{π}. 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). Finally, we also compute the payoffs to liquidity providers who in our approach should break-even on average.

D. 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.

E. 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. We make the (strong) assumption that every share trading is intermediated once and that no intermediation takes place in the AMM. Therefore, we approximate the balanced volume as half of the previous day's volume. We believe that this approach is reasonably cautious: Lehar and Parlour (2021) assume that uninformed trades would be (approximately) reversed by an arbitrageur and therefore, if we followed this view, uninformed, non-intermediated volume would be doubled, or, put differently, intermediated volume would be replaced with arbitrageurs' volume.

WRDS provides several volume metrics to researchers that sometimes vary slightly: we use the daily maximum of these.

Order Size. Since there are sometimes very large orders that get processed in a specialized dark pool or OTC market, and an 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 per se, but by splitting orders buying and selling traders may take turns which would lower both their costs.

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 (14) plus fees to a measure of trading costs for a marketable order. We compare this cost to the half of the time-weighted quoted bid-ask spread; the effective spread may be biased when traders time their trades to take advantage of liquidity patterns.

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.[7] The process described above 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.

F. 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}.

[7] For instance, as the *Financial Times* (July 5, 2023) recently observed, cumulative overnight returns for the S&P500 ETF SPY since its inception in 1993 lead to a 10 fold price increase whereas intra-day returns are essentially "flat."

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.

Finally, we also need to compute whether empirically liquidity providers break even.

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 moot 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. Optimal Fee Visuals

We begin by exploring the data visually. All figures plot the variable of interest against the firm's log-market cap. Larger firms commonly have more volume and lower bid-ask spreads, and this sorting therefore provides a useful reference point. Each point in a graph represents an annual average of the variable of interest, and we use various colors to differentiate the years. One salient feature of several plots is that there are some mid-size firms for which the AMM does not do well. This is a group of roughly 200 firms that trade very infrequently, i.e., there is not enough volume to generate a sufficient fee income to offset adverse selection losses.

All figures that we discuss in what follows have two panels: Panels A on the left show datapoints for our ad hoc approach, where we approximate today's expected return by yesterday's; this data shows annual averages. Panels B on the right panel plot datapoints for the approach where we first estimate the return distribution parametrically. A datapoint here signifies the average for the last year of our data, from April 1, 2021 to March 31, March 2022.

Figure 5 displays the optimal fee. These fees are larger for smaller firms, chiefly because they have less volume and liquidity providers need to earn more per trade to break even.

Figure 6 shows the payoff to liquidity providers, where we use red values for losses in Panel B. For our modelling approach to be valid, the payoffs to liquidity providers should be close to zero — which they are. For instance, for estimated returns in

Panel B, average liquidity provider returns are -0.2bps.

Figure 7 plots the deposited amounts $\overline{\alpha}$, which would allow liquidity providers to break even. These are generally decreasing in size, albeit weakly.

Figure 8 plots the number of days for which an AMM is indeed feasible, $\underline{\alpha} < \overline{\alpha}$, i.e., days for which liquidity providers are willing to provide amounts of liquidity that would entice liquidity demanders to use the AMM.

Figure 9 is the key finding of our paper: it displays the average benefit to liquidity providers π of using an AMM. The figure shows that almost always liquidity demanders would benefit from the presence of an AMM, with higher relative gains for small firms. Notably, our analysis makes no assumptions about behavioral changes that would result from lower transactions costs. It is conceivable that stocks with strong costs declines would increase their trading activity, which would further magnify the benefits.

Figure 10 displays the average daily savings from using an AMM. Here, gains are larger for large firms because dollar savings compound over many transactions.

Finally, Figure 11 displays the implied excess depth of an AMM, that is, the amount that a trader can trade on an AMM to have the same cost as on a traditional markets minus the depth of the traditional market. As can be seen, the implied depth of AMMs almost always exceed that of the traditional market. This finding is particularly noteworthy. AMM pricing is convex and larger quantities cost more whereas on traditional markets, traders can trade the full depth at a constant spread. Since our AMM calibration is based on the average order size, one may worry that costs quickly escalate when traders submit larger orders; notably, posted depth often exceeds that size. However, we show that liquidity demanders can conceivably trade much more than the average trade size on an AMM before they have the same cost as the traditional market. In other words, the AMM also has better depth than the traditional market.

B. 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, 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, bid-ask half-spreads are around 40-50 bps. 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. Moreover, trading volume concentrates in the 20% largest firms.

The most important results are in Table II 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 feasible. 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.[8]

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, an 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 returns is in a band $R \in (\underline{R}, \bar{R})$. 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 (5) 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

[8] During much of the sample, borrowing rates were close to zero, and presumably, when capital costs rise, so would the bid-ask spreads in the competing traditional market.

to compute how much needs to be available to have liquidity available for prices in $(p_0 \cdot \underline{R}, p_0 \cdot \bar{R})$, 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

$$\text{maxshares}(\underline{R}) \equiv 1 - \sqrt{\underline{R}^{-1}}. \tag{20}$$

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 $R = \underline{R}$

$$\text{maxcash}(\underline{R}) \equiv \left| \frac{1 - \sqrt{\underline{R}^{-1}}}{\sqrt{\underline{R}^{-1}}} \right|, \tag{21}$$

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 $\underline{R}$? 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 $R = 0.9$, $\text{maxcash}(0.9) \approx 5.1\%$, i.e., about 5.1% of the required AMM deposit in stocks is necessary in cash. Thus, if $\overline{\alpha} = 5\%$, one would need 25 basis points of the market cap in cash for the AMM to work.

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 13 displays the year-stock averages of maxcash($\underline{R}$) based on a 20-day (4-week) 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 costly, the required amounts are but a fraction of the "on-paper" liquidity.

VI. Discussion and Conclusion

The principle behind the superiority of AMMs is somewhat similar to 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 costless for the liquidity providers; 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. 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.

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 is 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, current tax rules 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 they 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 adapt for the investors that they protect to enjoy the benefits. 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 broker-dealers 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 blockchain-based 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, 2023, The conceptual flaws of decentralized automated market making, *Management Science* forthcoming.

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 the concept itself is implemented.

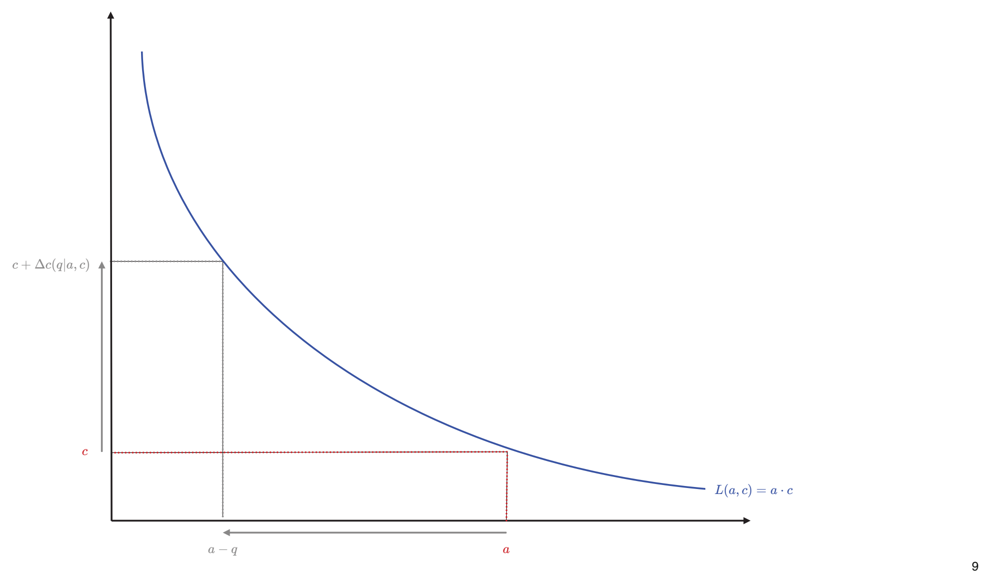

Figure 1
Constant Product Prices in Automated Market Makers

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.

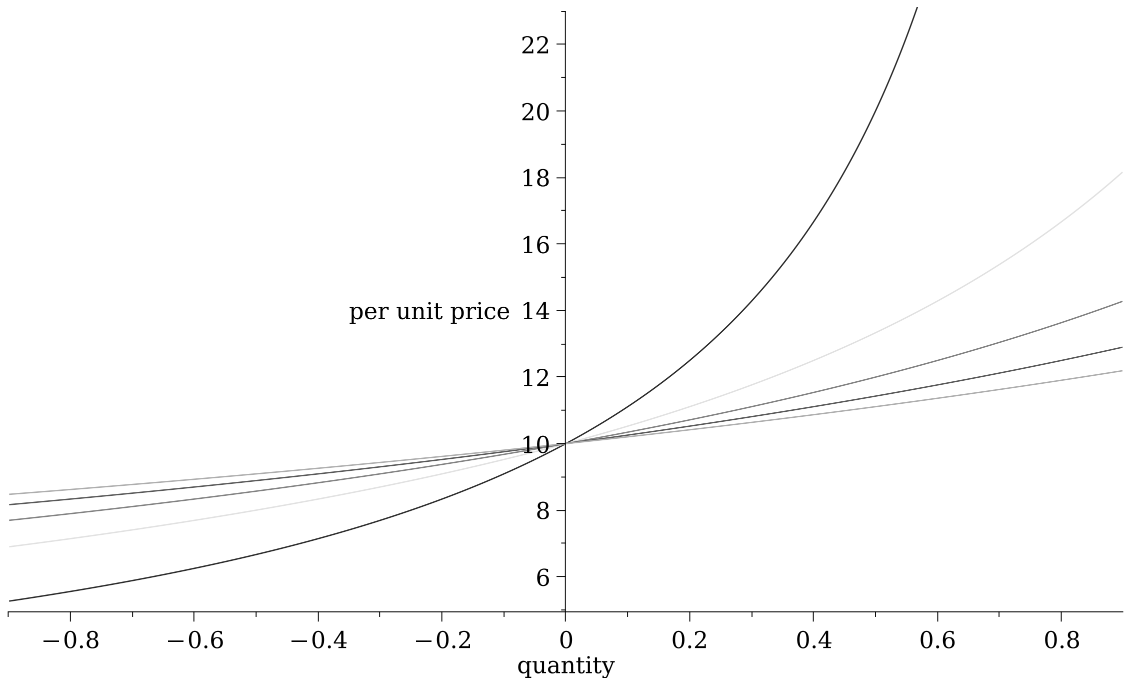

Figure 2
The Per-Unit Price Function for Constant Product AMMs

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 $a = 1, \ldots, 5$.

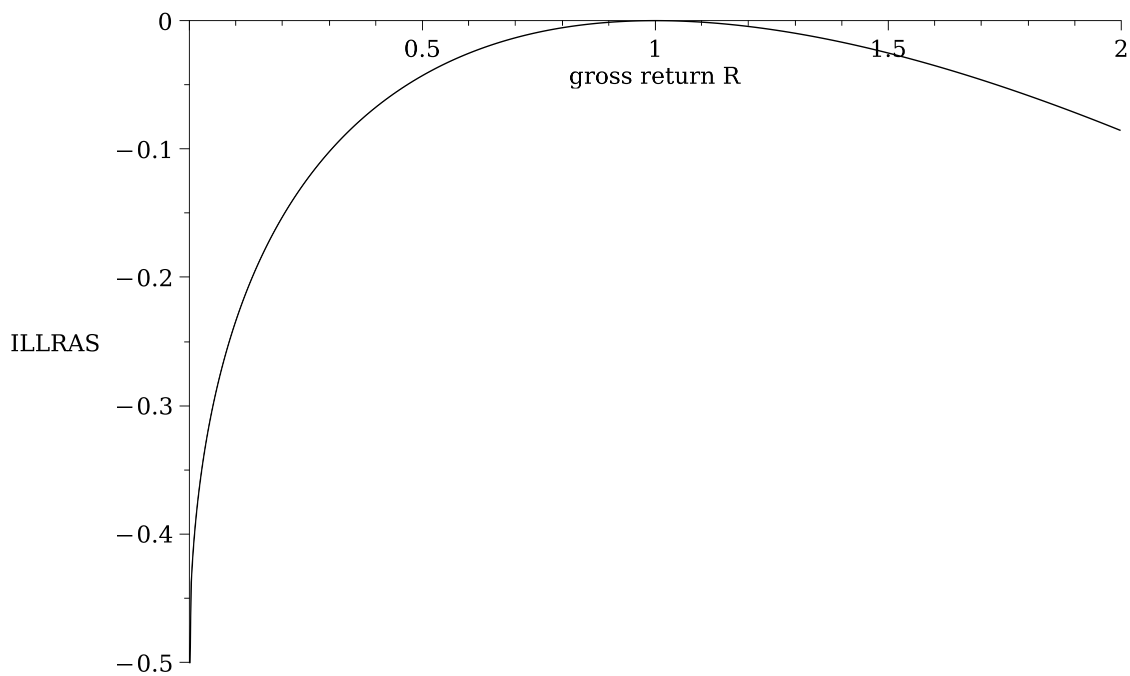

Figure 3
The Incremental Loss from Long-Run Adverse Selection (ILLRAS)
for Constant Product Pricing

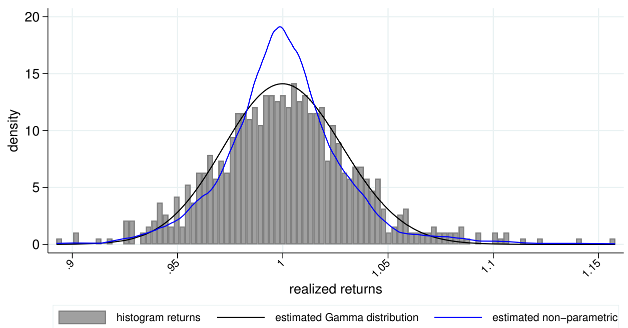

Panel A: TSLA

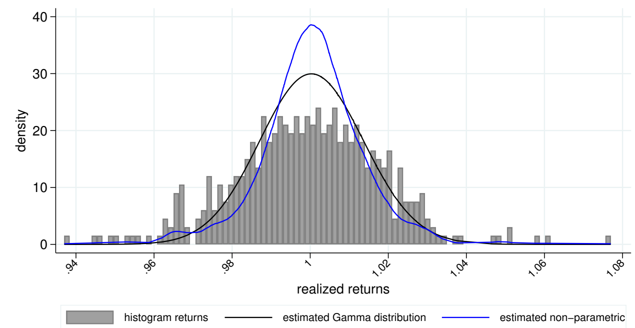

Panel B: MSFT

Figure 4

Realized and Estimated Return Examples

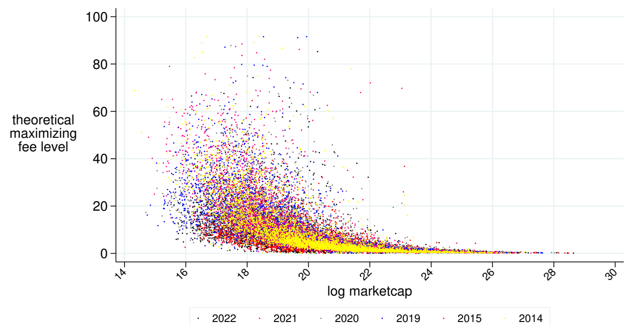

Panel A: based on ad hoc returns

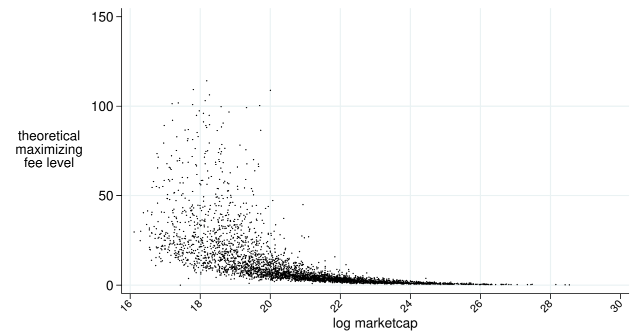

Panel B: based on estimated returns

Figure 5
Benefit-Maximizing AMM Fees

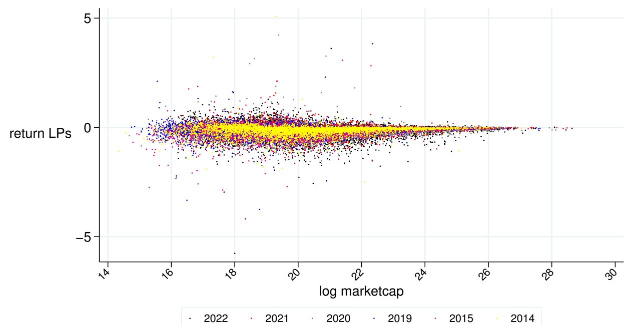

Panel A: based on ad hoc returns

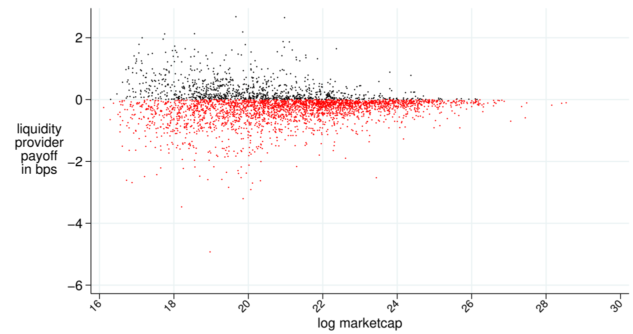

Panel B: based on estimated returns

Figure 6
Costs/Benefits to Liquidity Providers

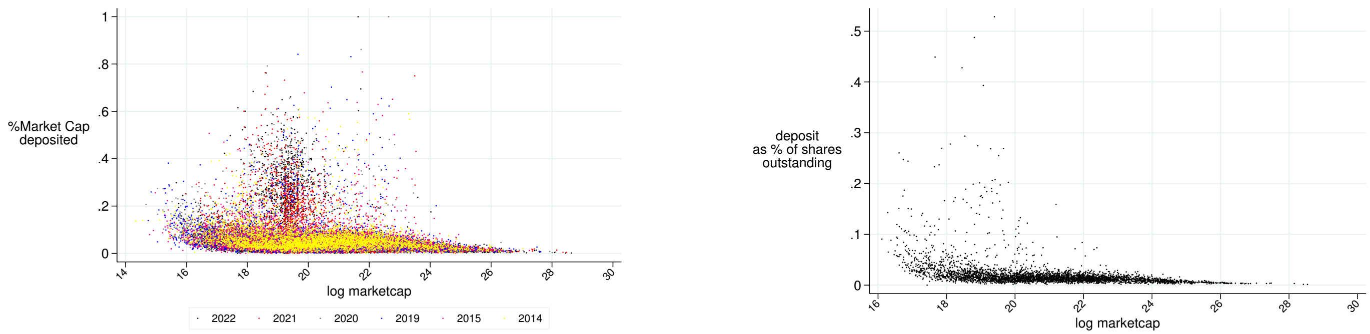

Panel A: based on ad hoc returns

Panel B: based on estimated returns

Figure 7
Fraction of MarketCap Deposited

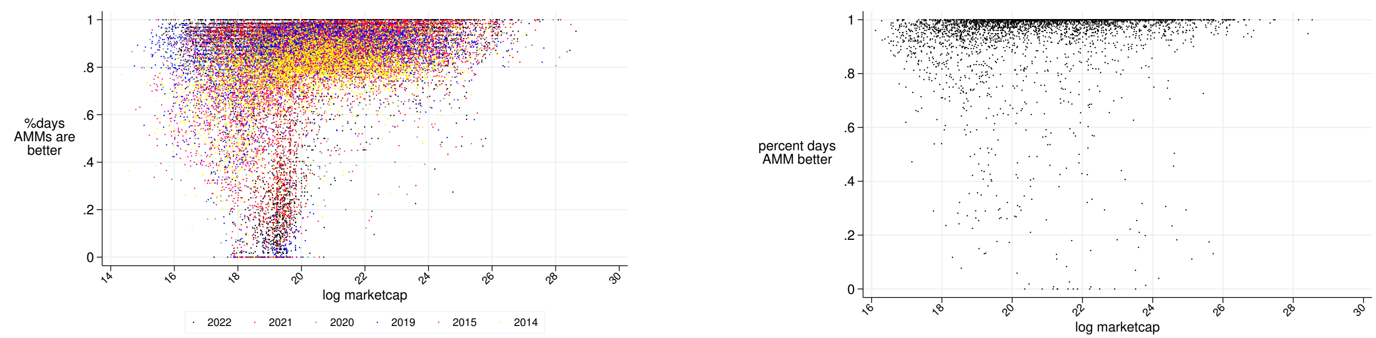

Panel A: based on ad hoc returns

Panel B: based on estimated returns

Figure 8
Percent Days that AMM is better

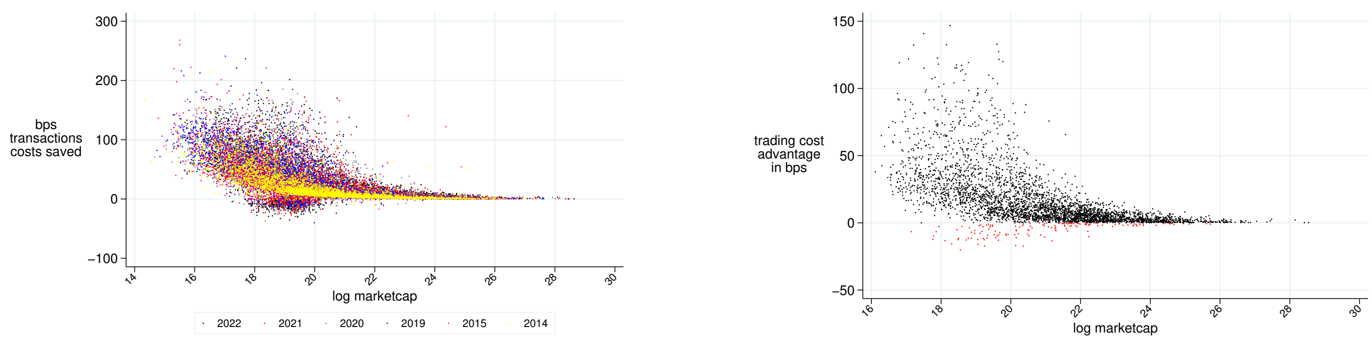

Panel A: based on ad hoc returns *Panel B: based on estimated returns*

Figure 9
Cost Advantage AMM vs. Traditional (measured as difference in relative transaction costs)

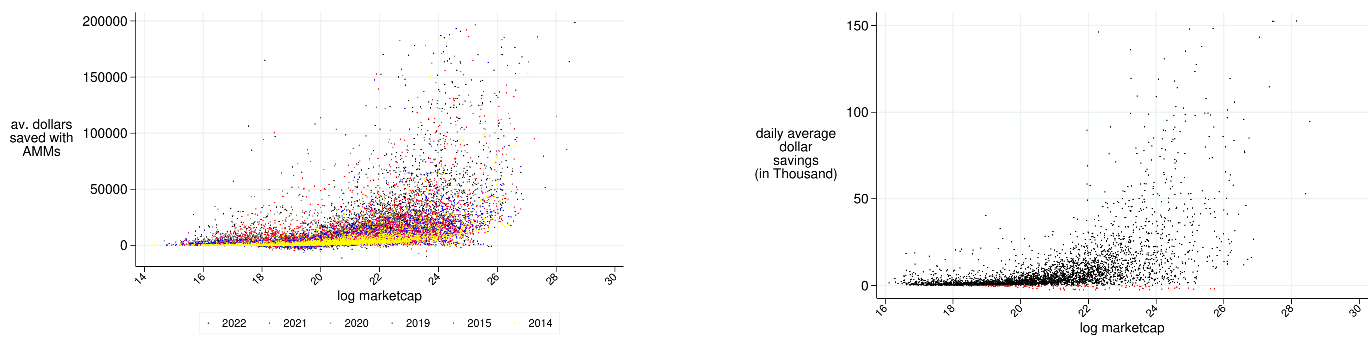

Panel A: based on ad hoc returns *Panel B: based on estimated returns*

Figure 10

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

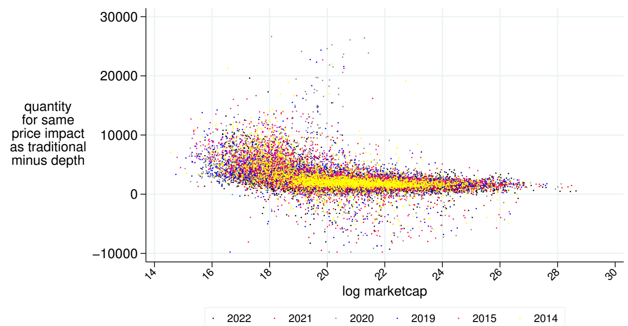

Panel A: based on ad hoc returns

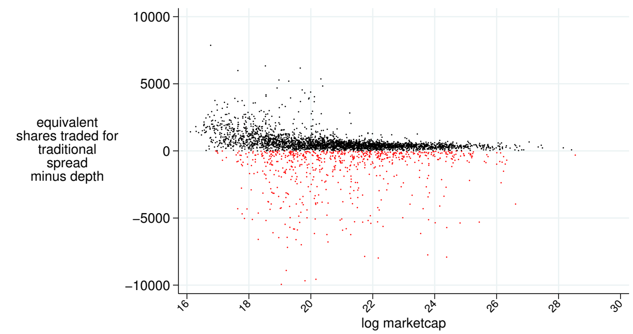

Panel B: based on estimated returns

Figure 11
Equivalent Depth for AMM Trading

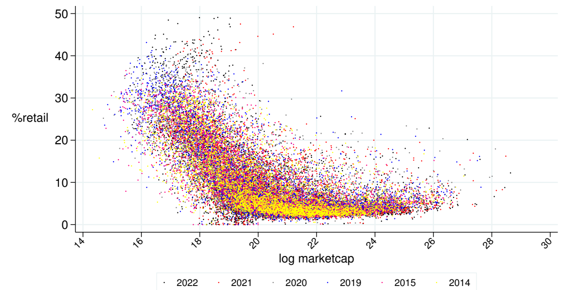

Panel A: % retail vs. Market Cap

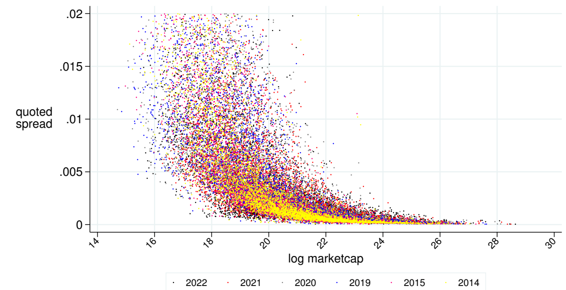

Panel B: Quoted Spread vs. Market Cap

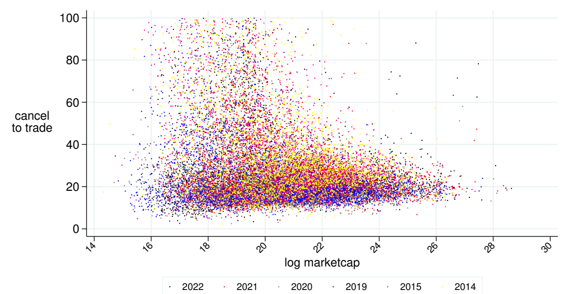

Panel C: Cancel-to-Trade vs. Market Cap

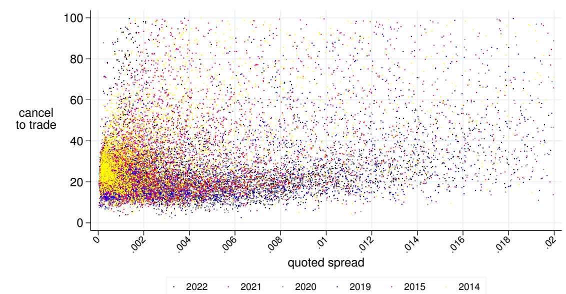

Panel D: Cancel-to-Trade vs. Quoted Spread

Figure 12
Relationships among the Explanatory Variables

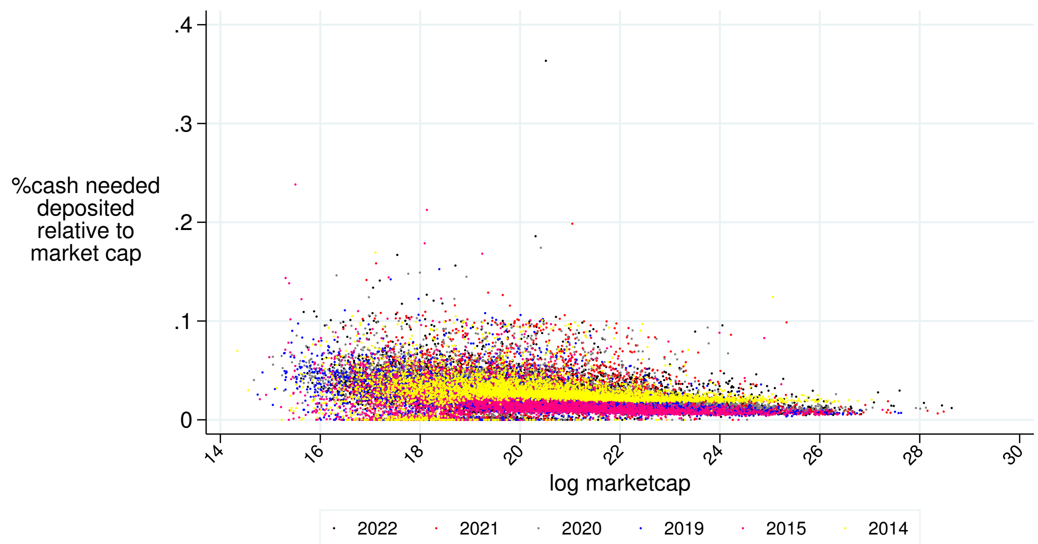

Figure 13
Required Cash Requirements as a Fraction of Asset Deposits

Table I
Summary 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 II
Summary Statistics of Based on Optimal Fees — Ad Hoc Approach

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 III
Summary Statistics of Based on Optimal Fees — Estimated Distributions

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.

	optimal fee (in bps)	Days AMMs better	percent mcap deposited	daily savings (bps)	daily savings (in USD)	aggregate savings (in M$)	liquidity supplier costs (bps)	equivalent AMM depth	LOB depth
total	11	93.8	2.1	16.4	9535	8814	-0.2	1132	1259
Quintile									
1	32.8	91.9	4.8	40.6	2023	182	-0.1	2399	1364
2	27	91.1	3	34.3	1697	155	-0.2	1879	1327
3	17.9	90.1	2.9	26	1942	174	-0.2	1475	2670
4	11.6	94.1	2	21	2963	265	-0.2	1234	1376
5	7	95.1	1.6	14.5	4121	380	-0.2	961	1199
6	5.1	94.3	1.5	9.7	5749	534	-0.1	775	1125
7	3.6	95	1.4	7.2	8920	829	-0.2	705	1109
8	2.7	95.3	1.3	5.2	11830	1091	-0.2	639	949
9	1.7	95.4	1.2	3.8	19983	1854	-0.2	627	637
10	0.9	95.8	0.7	1.9	36176	3351	-0.1	626	829

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